

RELEASE 6.10

Infoblox NIOS Administrator Guide



Copyright Statements

© 2014, Infoblox Inc.— All rights reserved.

The contents of this document may not be copied or duplicated in any form, in whole or in part, without the prior written permission of Infoblox, Inc.

The information in this document is subject to change without notice. Infoblox, Inc. shall not be liable for any damages resulting from technical errors or omissions which may be present in this document, or from use of this document.

This document is an unpublished work protected by the United States copyright laws and is proprietary to Infoblox, Inc. Disclosure, copying, reproduction, merger, translation, modification, enhancement, or use of this document by anyone other than authorized employees, authorized users, or licensees of Infoblox, Inc. without the prior written consent of Infoblox, Inc. is prohibited.

For Open Source Copyright information, see [Appendix G, "Open Source Copyright and License Statements"](#), on page 1317.

Trademark Statements

Infoblox, the Infoblox logo, Grid, NIOS, bloxTools, Network Automation and PortIQ are trademarks or registered trademarks of Infoblox Inc.

All other trademarked names used herein are the properties of their respective owners and are used for identification purposes only.

Company Information

<http://www.infoblox.com/contact/>

Product Information

Hardware Models

Trinzic product line: 100, 810, 820, 1410, 1420, 2210, 2220, and Infoblox-4010

Network Insight: ND-800, ND-1400, ND-2200, and ND-4000

Trinzic Reporting: 1400, 2000, 2200, and 4000

Advanced DNS Protection: PT-1400, PT-2200, PT4000

Infoblox-4030 DNS Caching Accelerator Appliance

Infoblox-250-A, -550-A, -1050-A, -1550-A, -1552-A, -1852-A, -2000 and -2000-A

Network Automation: NetMRI-1102-A, NT-1400, NT-2200, and NT-4000

Document Number: 400-0530-000 Rev. E

Document Updated: May 23, 2014

Warranty Information

Your purchase includes a 90-day software warranty and a one year limited warranty on the Infoblox appliance, plus an Infoblox Warranty Support Plan and Technical Support. For more information about Infoblox Warranty information, refer to the Infoblox Web site, or contact Infoblox Technical Support.

Contents

Preface	33
Document Overview	34
Documentation Conventions	34
What's New	36
Related Documentation	39
Customer Care	40
User Accounts	40
Software Upgrades	40
Technical Support	40
 PART 1 APPLIANCE GUI	
 Chapter 1 Infoblox Grid Manager	43
Management System Requirements	46
Supported Browsers	46
Browser Limitations	47
About Grid Manager	48
Admin Permissions for Grid Manager	48
Logging in to the GUI	48
Setting Login Options	49
Specifying the Grid Name and Hostname	49
Creating a Login Banner	49
Changing the Password and Email Address	50
Specifying the Table Size	50
Selecting Your Home Page	51
Setting the Browser Time Zone	51
SSL (Secure Sockets Layer) Protocol	52
Managing Certificates	53
About HTTPS Certificates	53
About Client Certificates	55
About CA Certificates	56
Converting CA Certificates to PEM Format	57
About the Grid Manager Interface	58
System Messages	58
Security and Informational Banners	59
Breadcrumbs Navigation	59
Global Search	59
Finder Panel	59
Toolbar Panel	59
Help Panel	59
Wizards and Editors	60
Tooltips	60
Customizing Tables	60
Selecting Objects in Tables	60
Modifying Data in Tables	62

Finding and Restoring Data	63
Using Bookmarks	63
Using the Recycle Bin	64
Managing Third Party URL Links	66
Using Filters	67
Using Quick Filters	68
Using Global Search	69
Using the Go To Function	71
About Tasks	72
Viewing Tasks	72
Supported Objects for Scheduled and Approval Tasks	73
Guidelines for Upgrading, Backing Up, and Restoring Data	74
Scheduling Tasks	75
Scheduling Additions and Modifications	75
Scheduling Appliance Operations	76
Scheduling Deletions	76
Scheduling Recursive Deletions of Network Containers and Zones	76
Viewing Scheduled Tasks	77
Rescheduling Tasks	78
Canceling Scheduled Tasks	79
Configuring Approval Workflows	80
Supported Tasks for Different Admin Groups	80
Creating Approval Workflows	81
Viewing Approval Workflows	82
Modifying Approval Workflows	83
Deleting Approval Workflows	83
Viewing Approval Tasks	83
Viewing Workflow Notifications	83
Unsupported Operations for Submitters	84
About Long Running Tasks	85
Running Tasks in the Background	85
Monitoring Long Running Tasks	86
About CSV Import	86
CSV Import User Permissions	87
CSV Import Limitations	87
Viewing CSV Import Jobs	88
Creating a Data File for Import	89
Exporting Data to Files	89
Configuring Import Options	89
Managing CSV Imports	90
Exporting Displayed Data	91
Printing from Grid Manager	91
Multilingual Support	92
UTF-8 Supported Fields	92
UTF-8 Support Limitations	92
Support for Internationalized Domain Names	93
Decoding IDNs and Encoding Punycode	93
IDN Supported Fields	93
IDN Support Limitations	94

Chapter 2 Dashboards	97
About Dashboards	99
The Tasks Dashboard	99
About Task Packs	99
The IPAM Task Pack	100
Enabling the Network Automation Tasks	107
Registering Network Automation with NIOS	108
The Network Automation Task Pack	109
Network Automation Task Options	109
Network Provisioning Task	109
Using the Port Activation Automation Task	112
VLAN Reassignment	112
Provision Bare Metal Device	113
Rogue DHCP Server Remediation	113
Using the Task Viewer to View Job Logs and Approve Jobs	114
About Dashboard Templates	114
Adding Dashboard Templates	115
Resetting Dashboard Templates	115
Modifying Dashboard Templates	115
Deleting Dashboard Templates	116
Assigning Dashboard Templates	116
Status Dashboards	116
Adding Widgets to Dashboards	117
Grid Status	121
Grid Upgrade Status	122
Member Status (System Status)	123
DNS Statistics	124
Ranges Over Threshold	125
IPv4 Failover Associations Status	125
DHCP Statistics	126
Network Statistics	127
IPv4 Networks Over Threshold	128
Discovery Status	128
Advanced Discovery Status	129
My Commands	129
DDNS Statistics	130
System Activity Monitor	130
File Distribution Statistics	131
Active WebUI Users	131
Microsoft Servers Status Widget	131
Import Job Manager	132
Load Balancer Status	133
Pending Approvals	133
Response Policy Zone (RPZ) Statistics	134
Infoblox Community	136
Mobile Devices Status	136
Threat Protection Statistics	138
Chapter 3 Smart Folders	139
About Smart Folders	140
Global Smart Folders	141
My Smart Folders	141
Predefined Smart Folders	142

Creating Smart Folders	142
Viewing and Modifying Data in Smart Folders.....	143
Modifying Smart Folders.....	144
Deleting Smart Folders	144
Saving a Copy of a Smart Folder	145
Printing and Exporting Data in Smart Folders	145

PART 2 APPLIANCE ADMINISTRATION

Chapter 4 Managing Administrators 149

About Admin Accounts	152
About Admin Groups	154
Creating Superuser Admin Groups.....	155
Creating Limited-Access Admin Groups	156
About Admin Roles	157
Creating Admin Roles	157
Managing Admin Groups and Admin Roles	158
Modifying Admin Groups and Roles	158
Deleting Admin Groups and Roles	159
Viewing Admin Groups	159
Viewing Admin Roles	159
Viewing Admin Group Assignments	160
About Administrative Permissions	160
Defining Global Permissions.....	161
Defining Object Permissions.....	162
Defining DNS and DHCP Permissions for Grid Members.....	164
Applying Permissions and Managing Overlaps	166
Managing Permissions	167
Authenticating Administrators	169
Creating Local Admins.....	169
Managing Passwords	170
Modifying and Deleting Admin Accounts.....	171
About Remote Admins.....	171
Authenticating Admins Using RADIUS	173
Authentication Protocols.....	174
Accounting Activities Using RADIUS.....	174
Configuring Remote RADIUS Servers	174
Configuring RADIUS Authentication	175
Configuring a RADIUS Authentication Server Group	175
Authenticating Admins Using Active Directory	177
Configuring an Active Directory Authentication Service Group	178
Authenticating Admin Accounts Using TACACS+	179
TACACS+ Accounting	180
Configuring TACACS+	180
Configuring a TACACS+ Authentication Server Group.....	180
Authenticating Admins Using LDAP	182
Authentication Protocols.....	183
Configuring LDAP.....	183
Configuring an LDAP Server Group.....	183
Defining the Authentication Policy.....	185
Configuring a List of Authentication Server Groups	185

Configuring a List of Remote Admin Groups	186
Authenticating Admins Using Two-Factor Authentication	187
Best Practices for Configuring Two-Factor Authentication	189
Configuring the OCSP Authentication Server Group	189
Viewing the OCSP Authentication Server Group	191
Changing Password Length Requirements	191
Notifying Administrators	191
Administrative Permissions for Common Tasks	193
Administrative Permission for the Grid	195
Administrative Permissions for Grid Members	195
Administrative Permissions for Network Discovery	196
Administrative Permissions for Scheduling Tasks	196
Administrative Permissions for Microsoft Servers	197
Administrative Permissions for IPAM Resources	198
Administrative Permissions for IPv4 and IPv6 Networks	198
Administrative Permissions for Hosts	199
Administrative Permissions for DNS Resources	199
Administrative Permissions for DNS Views	200
Administrative Permissions for Zones	201
Administrative Permissions for Resource Records	202
Administrative Permissions for Shared Record Groups	203
Administrative Permissions for DNS64 Synthesis Groups	204
Administrative Permissions for DHCP Resources	205
Administrative Permissions for Network Views	206
Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks	207
Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations	208
Administrative Permissions for IPv4 or IPv6 DHCP Enabled Host Addresses	209
Administrative Permissions for IPv4 and IPv6 DHCP Ranges	210
Administrative Permissions for IPv4 or IPv6 DHCP Templates	211
Administrative Permissions for Roaming Hosts	212
Administrative Permissions for MAC Address Filters	212
Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories	213
Administrative Permissions for File Distribution Services	214
Administrative Permissions for Dashboard Tasks	214
Administrative Permissions for OCSP Server Groups and CA Certificates	215
Administrative Permissions for Load Balancers	216
Administrative Permissions for Named ACLs	217
Administrative Permissions for DNS Threat Protection	217
Chapter 5 Deploying a Grid	221
Introduction to Grids	223
Grid Communications	225
NAT Groups	226
Automatic Software Version Coordination	229
Grid Bandwidth Considerations	231
About HA Pairs	233
Planning for an HA Pair	233
About HA Failover	234
VRRP Advertisements	235
Creating a Grid Master	236
Port Numbers for Grid Communication	238
Grid Setup Wizard	238

Creating an HA Grid Master	238
Creating a Single Grid Master	242
Adding Grid Members	245
Adding a Single Member	245
Adding an HA Member	247
Joining Appliances to the Grid	249
Auto-Provisioning NIOS Appliances	250
Joining Auto-Provisioned Appliances to the Grid	250
Pre-Provisioning NIOS Appliances	252
Guidelines for Pre-provisioning Offline Grid Members	252
Configuring Pre-Provisioned Members	253
About Provisional Licenses	253
Joining Pre-Provisioned Members to the Grid	254
Configuration Example: Configuring a Grid	255
Cable All Appliances to the Network and Turn On Power	257
Create the Grid Master	257
Define Members on the Grid Master	259
Join Appliances to the Grid	260
Import DHCP Data	262
Import DNS Data	263
Using the Wizard	264
After Using the Wizard	266
Managing a Grid	267
Changing Grid Properties	267
Configuring Security Level Banner	268
Configuring Notice and Consent Banner	268
Configuring Informational Level Banner	269
Configuring Recursive Deletions of Networks and Zones	269
Setting the MTU for VPN Tunnels	270
Removing a Grid Member	270
Promoting a Master Candidate	270
About the Master Grid	271

Chapter 6 Deploying Independent Appliances 273

Independent Deployment Overview	275
System Manager GUI	276
Deploying a Single Independent Appliance	276
Method 1 – Using the LCD	277
Method 2 – Using the CLI	277
Method 3 – Using the Infoblox NIOS Startup Wizard	279
Configuration Example: Deploying a NIOS Appliance as a Primary DNS Server	281
Cabling the Appliance to the Network and Turning On Power	282
Specifying Initial Network Settings	282
Specifying Appliance Settings	282
Enabling Zone Transfers on the Legacy Name Server	283
Importing Zone Data on an Independent Appliance	284
Designating the New Primary on the Secondary Name Server (at the ISP Site)	286
Configuring NAT and Policies on the Firewall	287
Deploying an Independent HA Pair	287
Using the Infoblox NIOS Startup Wizard to Configure an HA Pair	287
Configuration Example: Configuring an HA Pair for Internal DNS and DHCP Services	290
Cabling Appliances to the Network and Turning On Power	291
Specifying Initial Network Settings	292

Specifying Appliance Settings	292
Enabling Zone Transfers	294
Importing Zone Data	294
Defining Networks, Reverse-Mapping Zones, DHCP Ranges, and Infoblox Hosts	294
Defining Multiple Forwarders	297
Enabling Recursion on External DNS Servers	297
Modifying the Firewall and Router Configurations	298
Enabling DHCP and Switching Service to the NIOS Appliance	299
Managing and Monitoring	299
Verifying the Deployment	300
Single Independent Appliance	300
Independent HA Pair	300
Infoblox Tools for Migrating Bulk Data	301
 Chapter 7 Managing Appliance Operations	 303
Configuring Access Control	306
Administrative Permissions	306
Operations that Support Access Control	306
Defining Named ACLs	307
Managing Named ACLs	309
Applying Access Control to Operations	311
Managing Time Settings	312
Changing Time and Date Settings	312
Changing Time Zone Settings	312
Monitoring Time Services	313
Using NTP for Time Settings	313
Authenticating NTP	314
NIOS Appliance as NTP Client	315
Configuring a Grid to Use NTP	316
Configuring Grid Members to Use NTP	318
NIOS Appliances as NTP Servers	319
Configuring a NIOS Appliance as an NTP Server	320
Monitoring NTP	322
About Extensible Attributes	322
Adding Extensible Attributes	324
Configuring Inheritable Extensible Attributes	326
Viewing Extensible Attributes	330
Modifying Extensible Attributes	330
Deleting Extensible Attributes	331
Using Extensible Attributes	332
Configuration Examples for Inheritable Extensible Attributes	334
Managing Security Operations	343
Enabling Support Access	343
Enabling Remote Console Access	343
Permanently Disabling Remote Console and Support Access	344
Restricting GUI/API Access	344
Enabling HTTP Redirection	344
Modifying the Session Timeout Setting	344
Disabling the LCD Input Buttons	344
Configuring Security Features	344
Configuring Ethernet Ports	346
About Virtual LANs	346
Implementing Quality of Service Using DSCP	348

Ethernet Port Usage	349
Modifying Ethernet Port Settings	354
Using the LAN2 Port	355
About Port Redundancy	356
Configuring the LAN2 Port	357
Enabling DHCP on LAN2	358
Enabling DNS on LAN2	358
Using the MGMT Port	359
Appliance Management	360
Grid Communications	362
DNS Services	364
About Lights Out Management	366
Enabling LOM	367
Adding LOM User Accounts	368
Configuring the IPMI Network Interface	368
Modifying LOM Settings	369
Viewing LOM Users	369
IPMI Commands and Examples	369
Setting Static Routes	372
Defining IPv6 Static Routes	375
Enabling DNS Resolution	376
Managing Licenses	377
Obtaining and Adding Licenses	377
Obtaining Temporary Licenses	378
Viewing Licenses	378
Backing Up Licenses	379
Removing Licenses	379
Managing the Order of Match Lists	380
Shutting Down, Rebooting, and Resetting a NIOS Appliance	380
Rebooting a NIOS Appliance	380
Shutting Down a NIOS Appliance	380
Resetting a NIOS Appliance	381
Managing the Disk Subsystem on the Infoblox-2000-A and -4010	382
About RAID 10	382
Evaluating the Status of the Disk Subsystem	383
Disk Drive Front Panel LEDs	383
Replacing a Failed Disk Drive	384
Disk Array Guidelines	385
Restarting Services	386
Canceling a Scheduled Restart	388
Chapter 8 File Distribution Services	389
File Distribution Overview	391
Staged Upgrade Limitations	391
File Distribution Storage	392
Usage Threshold Alerts	392
Modifying File Distribution Storage Limits	392
Managing File Distribution Services	393
Configuring the TFTP Service	393
Configuring the FTP Service	393
Configuring the HTTP Service	394
Configuring Access Control for File Distribution	394
Modifying Access Control Lists	395

Starting and Stopping File Distribution Services.....	396
Monitoring File Distribution Services.....	396
Managing Directories	397
Adding Directories.....	397
Modifying Directories	397
Creating a Virtual TFTP Root Directory	397
Viewing Directories From the Files Tab.....	398
Managing Files	399
Uploading Files	399
Enabling Upload to Grid Members	399
Uploading Files using Grid Manager	399
Uploading Files Using TFTP, FTP, or HTTP File Transfer Client	400
Deleting Files From the Grid Master.....	400
Deleting Files From a Member	400
Viewing Files	401
Viewing Files from the Files Tab	401
Viewing Files from the Members Tab	401
Managing Users.....	402
Users Default Home Directory.....	402
Adding FTP Users through Grid Manager.....	402
Adding FTP Users through CSV Import.....	403
Modifying FTP Users	403
Chapter 9 Managing NIOS Software and Configuration Files.....	405
About Upgrades.....	406
Lite Upgrades.....	406
Full Upgrades.....	407
Guidelines for Scheduling Full Upgrades.....	407
Microsoft Management Rules.....	408
Managing Upgrade Groups.....	408
Adding Upgrade Groups	409
Modifying Upgrade Groups.....	410
Viewing Upgrade Groups	410
Deleting Upgrade Groups	411
Viewing Software Versions	411
Upgrading NIOS Software.....	411
Uploading NIOS Software	412
Distributing Software Upgrade Files.....	412
Managing Distributions.....	414
Testing Software Upgrades.....	415
Performing Software Upgrades	416
Managing Upgrades	420
Monitoring Distribution and Upgrade Status.....	421
Downgrading Software	422
Reverting the Grid to the Previously Running Software	423
Backing Up and Restoring Configuration Files.....	423
Backing Up Files	423
Automatically Backing Up Data Files	424
Manually Backing Up Data Files	426
Downloading Backup Files	427
Restoring Backup Files	428
Downloading Backup Files from a Different Appliance	429
Downloading Support Bundles.....	429

Chapter 10 bloxTools Environment	431
About the bloxTools Environment	432
System Requirements	432
Using the bloxTools Environment	433
Configuring the Service	433
Allocating Memory	434
Uploading Files	434
Scheduling Tasks	435
Moving the bloxTools Service	435
Monitoring the Service	435
Viewing the Logs	435
Viewing Detailed Status	436
 Chapter 11 RIR Registration Updates	 437
RIR Address Allocation and Registration Updates	438
About the RIPE Database	438
Requirements and Permissions	438
Configuring RIR Registration Updates	439
Enabling Support for RIR Registration Updates	439
Configuring RIR Communication Settings	440
Managing RIR Data	440
Adding RIR Organizations	441
Modifying RIR Organizations	442
Deleting RIR Organizations	442
Adding and Assigning RIR Networks	442
Viewing RIR networks	443
Modifying RIR Network Data	443
Deleting RIR Networks	444
Managing RIR Attributes	444
RIR Organizational Attributes	445
RIR Network Attributes	447
Monitoring RIR Data	451
Viewing RIR Organizations	451
Previewing Registration Updates	451
 PART 3 IP ADDRESS MANAGEMENT	
 Chapter 12 IP Address Management	 457
About IP Address Management	458
About Network Insight	459
Viewing the Complete List of Discovered Devices	459
About Host Records	459
Assigning Multiple IP Addresses to a Host	461
Adding Host Records	462
Modifying Host Records	463
About Network Containers	464
Adding IPv4 and IPv6 Network Containers and Networks	465
Modifying IPv4 and IPv6 Network Containers and Networks	465
Deleting Network Containers	465

Managing IPv4 Networks	466
IPv4 Network Map	467
Network List	470
Resizing IPv4 Networks	472
Splitting IPv4 Networks into Subnets	472
Joining IPv4 Networks	473
Discovering Networks (Under Network Insight only)	473
Deleting Networks	474
Viewing and Managing IPv4 Addresses	474
IP Map	474
IP Address List	476
Managing IPv4 Addresses	479
Managing IPv6 Networks	480
IPv6 Network Map	480
IPv6 Network List	484
Splitting IPv6 Networks into Subnets	485
Joining IPv6 Networks	485
Viewing IPv6 Data	486
Managing IPv4 and IPv6 Addresses	487
Converting Objects Associated with IP Addresses	487
Reclaiming Objects Associated with IPv4 and IPv6 Addresses	491
Pinging IP Addresses	491
Clearing Active DHCP Leases	491
 Chapter 13 Network Discovery	 493
About Network Discovery	494
Administrative Permissions	495
IP Discovery Process	496
Supported IP Discovery Methods	497
VM Discovery Process	499
About Configuring a Discovery	500
Before Starting a Discovery	501
Selecting a Grid Member	502
Enabling or Disabling the Merging of Discovered Data	502
Updating Discovered Data for Managed Objects	503
Configuring IP Discovery	503
Configuring VM Discovery	504
Guidelines for Starting and Scheduling a Discovery	505
Starting a Discovery Immediately	505
Scheduling a Discovery	506
Configuring a Recurring Discovery	506
Managing a Discovery	507
Monitoring Discovery Status	507
Integrating Data from PortIQ Appliances	508
Integrating Discovered Data From Trinzic Network Automation	509
Viewing Discovered Data	510
Managing Discovered Data	511
Managing Unmanaged Data	511
Resolving Conflicting Addresses	513
Clearing Discovered Data	515

Chapter 14 Network Insight	517
About Network Insight	519
Consolidators and Probes	520
Administrative Permissions	521
Supported Discovery Methods	522
SNMP	523
ICMP	523
TCP	523
Port Scanning	524
NetBIOS	524
Starting and Stopping the Discovery Service	525
Changing the Discovery Member Type	526
Consolidator and Probe Appliance Deployment Guidelines	527
Choosing a Probe Member Interface for Discovery	527
Discovering IPs and Networks	527
Performing Discovery on an Existing Object	529
Smart Folders and Discovered Devices	529
Configuring Grid Properties for Discovery	529
Activating DHCP Routers as Seed Routers	529
Configuring Grid SNMPv1/v2 Properties	530
Configuring Grid SNMPv3 Properties	530
Defining Advanced Discovery Polling Techniques for the Grid	531
Scheduling Switch Port Discovery and Data Collection	532
Configuring Member Properties for Discovery	533
Defining Probe SNMPv1/v2 Properties	533
Configuring Probe SNMPv3 Properties	534
Defining Seed Routers	534
Excluding IP Addresses from Discovery	535
Enabling IPs for Immediate Discovery or for Discovery Exclusion	535
Quick Exclusion of IPs from Discovery	535
Creating a New Fixed Address Object and Excluding it from Discovery	536
Disabling Discovery for a Network	537
Viewing the List of Discovered Devices	537
Viewing Discovery Status	538
Using Discovery Diagnostics	539
Viewing Discovered Interface Information	540
Viewing the List of Networks Associated with a Discovered Device	540
Viewing the Management State of IPs in Discovered Networks	541
Viewing the List of IP Addresses Associated with a Discovered Device	541
Viewing the List of Assets Associated with a Discovered Interface	541
Converting Unmanaged Networks to Managed Networks	542
Adding Discovery Device Support	543
Discovery Data Management	543
Performing VM (Virtual Machine) Discovery	543
Executing a VM Discovery	543
Scheduling a VM Discovery Session	544

PART 4 DNS

Chapter 15 Infoblox DNS Service 547

Configuring DNS Overview	548
DNS Configuration Checklist.....	549
About Inheriting DNS Properties	550
Overriding DNS Properties	551
Understanding DNS for IPv6.....	552
Configuring IPv6 on a Grid Member	553
Configuring DNS for IPv6 Addressing.....	554

Chapter 16 Configuring DNS Services 555

Configuring DNS Service Properties.....	557
Configuring DNS Access Control.....	557
About Time To Live Settings	557
Adding an Email Address to the SOA Record	560
Notifying External Secondary Servers	561
Enabling the Configuration of RRset Orders	561
Specifying Port Settings for DNS	562
Using Extension Mechanisms for DNS (EDNS0).....	564
Specifying Minimal Responses.....	565
Starting and Stopping the DNS Service	565
About DNS Cache.....	565
Clearing DNS Cache.....	565
Clearing Cache for DNS Views.....	566
Clearing Domain Names from Cache	566
Viewing DNS Cache Entries.....	566
Viewing DNS Configuration.....	567
Viewing DNS Cache Details.....	568
Viewing Statistics	568
Using Forwarders	569
Specifying Forwarders.....	569
Controlling DNS Queries	570
Specifying Queriers.....	570
Enabling Recursive Queries	571
Enabling Recursion	571
Restricting Recursive Clients	572
Controlling AAAA Records for IPv4 Clients.....	573
Enabling AAAA Filtering.....	573
About NXDOMAIN Redirection	574
About NXDOMAIN Rulesets.....	575
Examples	576
NXDOMAIN Redirection Guidelines	577
Configuring NXDOMAIN Redirection.....	577
Creating Rulesets.....	577
Enabling NXDOMAIN Redirection	578
About Blacklists.....	579
About Blacklist Rulesets	580
Blacklist Guidelines.....	581
Configuring the Blacklist Feature	581
Enabling Blacklisting.....	582

Enabling Zone Transfers	583
Configuring Zone Transfers	584
Configuring Concurrent Zone Transfers	586
About Root Name Servers	587
Specifying Root Name Servers	587
About Sort Lists	588
Defining a Sort List	588
Configuring a DNS Blackhole List	590
Defining a DNS Blackhole List	590
Specifying Hostname Policies	592
Defining Grid Hostname Policies	592
Defining Hostname Restrictions	593
Obtaining a List of Invalid Record Names	594
About DNS64	594
Configuring DNS64	595
About Synthesis Groups	595
Chapter 17 DNS Views	601
Using Infoblox DNS Views	602
About DNS Views and Network Views	604
Configuring DNS Views	604
Adding a DNS View	605
Defining Match Clients Lists	605
Defining a Match Destinations List	607
Copying Zone Records	608
Managing the DNS Views of a Grid Member	609
Managing Recursive DNS Views	609
Managing the Order of DNS Views	610
Managing DNS Views	611
Configuration Example: Configuring a DNS View	612
Chapter 18 Configuring DNS Zones	615
About Authoritative Zones	616
Configuring Authoritative Zones	616
Creating an Authoritative Forward-Mapping Zone	617
Creating an Authoritative Reverse-Mapping Zone	618
Creating a Root Zone	620
Adding an Authoritative Subzone	620
Locking and Unlocking Zones	621
Enabling and Disabling Zones	621
About Domains and Zones	622
IDN Support For DNS Zones	622
Assigning Zone Authority to Name Servers	623
Specifying a Primary Server	623
Specifying a Secondary Server	626
Using Name Server Groups	629
Adding Name Server Groups	629
Viewing Name Server Groups	630
Applying Name Server Groups	630
Importing Zone Data	631
About Importing Data into a New Zone	632

About Importing Data into an Existing Zone	632
Importing Data into Zones	632
Configuring Authoritative Zone Properties	633
Removing Zones	634
Restoring Zone Data	636
Configuring Delegated, Forward, and Stub Zones	638
Configuring a Delegation	638
Configuring a Forward Zone	640
Configuring Stub Zones	644
Viewing Zones	653
Chapter 19 DNS Resource Records	655
About Bulk Hosts	656
Specifying Bulk Host Name Formats	656
Before Defining Bulk Host Name Formats	656
Adding Bulk Hosts	659
Managing Resource Records	660
Managing A Records	660
Managing NS Records	662
Managing AAAA Records	662
Managing PTR Records	664
Managing MX Records	665
Managing SRV Records	666
Managing TXT Records	668
Managing CNAME Records	669
Managing DNAME Records	671
Managing NAPTR Records	676
Managing LBDN Records	678
Viewing Resource Records	678
Modifying, Disabling, and Deleting Host and Resource Records	679
About Shared Record Groups	680
Shared Records Guidelines	681
Configuring Shared Record Groups	681
Managing Shared Resource Records	683
Managing Associated Zones	685
Configuration Example: Configuring Shared Records	686
Chapter 20 Configuring DDNS Updates from DHCP	689
Understanding DDNS Updates from DHCP	691
Configuring DHCP for DDNS	695
Enabling DDNS for IPv4 and IPv6 DHCP Clients	695
Sending Updates to DNS Servers	696
Configuring DDNS Features	697
Resending DDNS Updates	697
Generating Host Names for DDNS Updates	698
Updating DNS for IPv4 Clients with Fixed Addresses	698
Configuring DDNS Features	698
Replacing Host Names for DDNS Updates	699
About the Client FQDN Option	701
Enabling FQDN Option Support	702
Sending Updates for DHCP Clients Using the FQDN Option	703
Configuring DDNS Update Verification	703

Configuring DNS Servers for DDNS	705
Enabling DNS Servers to Accept DDNS Updates	706
Forwarding Updates	707
Supporting Active Directory	709
Sending DDNS Updates to a DNS Server	709
About GSS-TSIG	710
Sending Secure DDNS Updates to a DNS Server in the Same Domain.	711
Configuring DHCP to Send GSS-TSIG Updates in the Same Domain	712
Sending Secure DDNS Updates to a DNS Server in Another Domain.	719
Configuring DHCP to Send GSS-TSIG Updates to Another Domain.	720
Sending GSS-TSIG Updates to a DNS Server in Another Forest	722
Accepting DDNS Updates from DHCP Clients.	723
Supporting Active Directory and Unauthenticated DDNS Updates.	723
Accepting GSS-TSIG-Authenticated Updates	725
Configuring DNS to Receive GSS-TSIG Updates.	727
 Chapter 21 DNSSEC	 733
About DNSSEC	734
DNSSEC Resource Records	735
DNSKEY Resource Records	735
RRSIG Resource Records	737
NSEC/NSEC3 Resource Records	738
NSEC3PARAM Resource Records	739
DS Resource Records	740
Configuring DNSSEC on a Grid	741
Grid Master as Primary Server	741
Enabling DNSSEC	743
Setting DNSSEC Parameters	743
About the DNSKEY Algorithm	743
About Key Rollovers	744
RRSIG Signatures	745
Configuring DNSSEC Parameters	745
Signing a Zone	746
Managing Signed Zones	747
Importing a Keyset	748
Exporting Trust Anchors	748
Checking Key-Signing Keys	749
Rolling Key-Signing Keys	749
Unsigning a Zone	749
Deleting and Restoring Signed Zones	749
About HSM Signing	750
Configuring a SafeNet HSM Device	750
Monitoring the HSM Group	753
Enabling HSM Signing	753
Testing the HSM Group	753
Synchronizing the HSM Group	754
Configuring Grid Members to Support DNSSEC as Secondary Servers	754
Configuring Recursion and Validation for Signed Zones	754
Enabling Recursion and Validation for Signed Zones	755
Enabling DNSSEC Validation	755

Chapter 22 Configuring IP Routing Options. 757

Using the Loopback Interface	758
Configuring IP Addresses on the Loopback Interface	759
Advertising Loopback Addresses to the Network	760
About Anycast Addressing for DNS	761
Configuring Anycast Addresses	762
Best Practices for Configuring Anycast Addresses	763
IP Routing Options.	763
Anycast and OSPF	764
Configuring OSPF on the NIOS Appliance	765
Anycast and BGP4	767
Configuring BGP in the NIOS Appliance	769

PART 5 DHCP

Chapter 23 Infoblox DHCP Services 773

About Infoblox DHCP Services	774
IPv4 DHCP Protocol Overview	774
IPv6 DHCP Protocol Overview	775
IPv6 Address Structure	776
Configuring DHCP Overview	777
Managing DHCP Data	779
About Networks	779
About Shared Networks	779
About DHCP Ranges	780
About Fixed Addresses	780
About Hosts	780
DHCP Configuration Checklists.	780
About DHCP Inheritance	782
Overriding DHCP Properties	783
Viewing Inherited Values	783
About Network Views	787
Adding Network Views	789
Modifying Network Views	789
Deleting Network Views	790

Chapter 24 Configuring DHCP Properties. 791

About DHCP Properties	793
Configuring IPv4 DHCP Properties	793
Configuring General IPv4 DHCP Properties	793
Specifying Authoritative	794
Defining Lease Times	794
Scavenging Leases	794
Configuring Ping Settings	795
Configuring One Lease per Client	797
Ignoring DHCP Client Identifiers.	797
Limitations of the Ignore Client ID Feature on DHCP Failover Associations	798
Configuring IPv4 BOOTP and PXE Properties	798

About IPv4 DHCP Options	800
DHCP Option Data Types	800
Configuring IPv4 DHCP Options	801
Defining Basic IPv4 Options	801
Defining IPv4 Option Spaces	802
Configuring Custom DHCP Options	803
Applying DHCP Options	804
Configuration Example: Defining a Custom Option	805
Defining Option 60 Match Rules	805
About the DHCP Relay Agent Option (Option 82)	806
Configuring Thresholds for DHCP Ranges	807
Configuring DHCPv6 Properties	809
Defining General IPv6 Properties	809
About DHCPv6 Options	810
Configuring DHCPv6 Options	810
Defining IPv6 Option Spaces	810
Configuring Custom IPv6 DHCP Options	811
Applying DHCPv6 Options	811
Configuring DHCP IPv4 and IPv6 Common Properties	812
Configuring UTF-8 Encoding for Hostnames	812
Associating Networks with Zones	813
Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges	814
Configuring Fixed Address Leases For Display	814
Configuring DHCP Logging	815
Configuring the Lease Logging Member	815
About IF-MAP	816
Configuring a Grid to Support IF-MAP	817
Validating the IF-MAP Server Certificate	818
Configuring Members as IF-MAP Clients	818
Creating IF-MAP Client Certificates	819
Overriding IF-MAP Publishing Settings	820
Deleting Data from the IF-MAP Server	821
Starting DHCP Services on a Member	822
Viewing DHCP Member Status	822
Viewing DHCP Configuration Files	824

Chapter 25 Managing DHCP Templates 825

About DHCP Templates	826
About IPv4 DHCP Templates	826
About IPv4 Range Templates	826
About IPv4 Fixed Address/Reservation Templates	828
About IPv4 Network Templates	829
Configuration Example: Creating an IPv4 Network Using a Template	832
About IPv6 DHCP Templates	834
About IPv6 Range Templates	834
Adding IPv6 Range Templates	834
Modifying IPv6 Range Templates	835
About IPv6 Fixed Address Templates	835
Adding IPv6 Fixed Address Templates	835
Modifying IPv6 Fixed Address Templates	836
About IPv6 Network Templates	837
Adding IPv6 Network Templates	837
Modifying IPv6 Network Templates	838

Viewing Templates.....	839
Deleting Templates	839

Chapter 26 Managing IPv4 DHCP Data..... 841

Configuring DHCP for IPv4	843
About the Next Available Network or IP Address	844
Obtaining the Next Available Network	844
Obtaining the Next Available IP Address	844
Guidelines for the Next Available Network and IP Address.....	844
Configuring IPv4 Networks	845
Adding IPv4 Networks	845
Viewing Networks	848
Modifying IPv4 Networks	851
Deleting IPv4 Networks.....	852
Configuring IPv4 Shared Networks.....	852
Adding IPv4 Shared Networks	852
Viewing Shared Networks	853
Modifying IPv4 Shared Networks	853
Deleting IPv4 Shared Networks	854
Configuring IPv4 Address Ranges	854
Adding IPv4 Address Ranges	854
Modifying IPv4 Address Ranges.....	855
Controlling Lease Assignments	856
Deleting IPv4 Address Ranges	857
Configuring IPv4 Fixed Addresses	857
Adding IPv4 Fixed Addresses	858
Modifying IPv4 Fixed Addresses.....	859
Deleting Fixed Addresses	860
Configuring IPv4 Reservations	860
Adding IPv4 Reservations	861
Modifying Reservations	862
Viewing IPv4 DHCP Objects	862
About Roaming Hosts	863
Configuring Roaming Hosts	863
Enabling Support for Roaming Hosts	864
Adding IPv4 Roaming Hosts	864
Adding IPv6 Roaming Hosts	865
Adding IPv4/IPv6 Roaming Hosts.....	865
Viewing Roaming Hosts	866
Setting Properties for Roaming Hosts	867
Deleting Roaming Hosts	868

Chapter 27 Managing IPv6 DHCP Data..... 869

Configuring IPv6 Networks	870
Defining Global IPv6 Prefixes	870
Managing IPv6 Networks	871
Adding IPv6 Networks.....	871
Modifying IPv6 Networks	874
Deleting IPv6 Networks.....	874
About IPv6 Shared Networks	875
Adding IPv6 Shared Networks	875

Modifying IPv6 Shared Networks	875
Deleting IPv6 Shared Networks	876
Configuring IPv6 Address Ranges	876
Adding IPv6 Address Ranges	876
Setting the Priority of IPv6 Address Ranges	877
Modifying IPv6 Address Ranges	878
Deleting IPv6 Address Ranges	878
Configuring IPv6 Fixed Addresses	878
Adding IPv6 Fixed Addresses	878
Modifying IPv6 Fixed Addresses	880
Deleting IPv6 Fixed Addresses	880
Viewing IPv6 DHCP Objects	880
 Chapter 28 DHCP Failover	883
About DHCP Failover	884
Failover Association Operations	884
Configuring Failover Associations	885
Adding Failover Associations	886
Managing Failover Associations	887
Modifying Failover Associations	887
Monitoring Failover Associations	888
Deleting Failover Associations	889
Setting a Peer in the Partner-Down State	889
Performing a Force Recovery	890
 Chapter 29 Configuring IPv4 DHCP Filters	891
About IPv4 DHCP Filters	892
IP Address Allocation	892
IP Address Allocation Using Filters	895
About MAC Address Filters	897
Defining MAC Address Filters	898
Adding MAC Address Filter Items	899
About Relay Agent Filters	899
Defining Relay Agent Filters	900
About Option Filters	901
DHCP Hardware Operator	901
Defining Option Filters	902
Using Match Rules in Option Filters	903
Configuring User Class Filters	904
Configuration Example: Using Option Filters	905
About DHCP Fingerprint Filters	906
Defining DHCP Fingerprint Filters	907
Applying Filters to DHCP Address Ranges	907
Adding Filters to the Class Filter List	907
Adding Filters to the Logic Filter List	908
Configuration Example: Using the Class and Logic Filter Lists	909
Managing DHCP Filters	912
Modifying DHCP Filters	912
Viewing DHCP Filters	914
Deleting Filters	914

Chapter 30 Authenticated DHCP..... 915

About Authenticated DHCP	917
DHCP Authentication Process	917
Configuring DHCP Authentication	921
About Authentication Server Groups	921
Configuring a RADIUS Authentication Server Group	921
Configuring an Active Directory Authentication Server Group	923
About the Captive Portal	923
Configuring Captive Portal Properties	924
Customizing the Captive Portal Interface	925
Managing Captive Portal Certificates	926
Starting the Captive Portal Service	927
Defining the IPv4 Network and DHCP Ranges	928
Defining MAC Address Filters	929
Using the Captive Portal Wizard	929
Adding and Modifying the Filters and Associations	930
Monitoring DHCP Authentication	931
Viewing DHCP Ranges and Filters	931
Configuration Example: Configuring Authenticated DHCP	931
NAC Integration	937
Configuring NAC with RADIUS Servers	938
About Authentication Servers	938
Adding a Server Group	938
Associating a Server Group with a Member	940
Managing Server Groups	940
Clearing the Authentication Cache	940
Configuring DHCP Ranges	940
Listing DHCP Ranges	940
About NAC Filters	941
Defining NAC Filters	942

Chapter 31 Managing Leases 945

About DHCP Leases	946
Viewing Current Leases	946
Viewing Detailed Lease Information	948
Viewing Lease History	949
Viewing Lease Event Detailed Information	949
Exporting Lease Records	950
Clearing Leases	950

PART 6 MANAGING MICROSOFT WINDOWS SERVERS

Chapter 32 Managing Microsoft Windows Servers..... 953

About Managing Microsoft Servers	954
Requirements	955
Deployment Guidelines	956
Configuring Members to Manage Microsoft Servers	957
Setting Microsoft Server Credentials	957
Configuring a Managing Member	958

Managing Microsoft Servers.....	961
Setting Microsoft Server Properties.....	961
Changing the Managing Member or Management Mode.....	961
Backing Up Synchronized Data.....	962
Disabling Synchronization.....	962
Removing a Managed Microsoft Server.....	962
Monitoring Managed Microsoft Servers.....	962
Viewing the Status of Servers.....	963
Viewing Detailed Status Information.....	964
Viewing Synchronization Logs.....	965

Chapter 33 Managing Microsoft DNS Services 967

Managing Microsoft DNS Servers.....	968
Synchronizing DNS Data.....	968
IDN Support for Synchronized DNS Data.....	969
Synchronizing with Multiple Servers.....	970
Managing Synchronized DNS Data.....	970
Adding Zones to Microsoft Servers.....	971
Setting Zone Properties.....	971
Deleting and Restoring Synchronized Zones.....	973
Managing Resource Records in Synchronized Zones.....	973
Synchronizing Updates.....	974
Synchronizing Delegations.....	977
Synchronizing AD-Integrated Zones.....	980
Resolving Conflicts.....	981
Viewing Members and Managed Servers.....	981
Specifying Forwarders for Microsoft Servers.....	982
Disabling and Removing Microsoft DNS Servers.....	982

Chapter 34 Managing Microsoft DHCP Services 983

About Microsoft DHCP Management.....	984
Synchronizing DHCP Data from Microsoft Servers.....	984
Viewing Synchronized Leases.....	986
Managing Synchronized DHCP Data.....	986
Adding and Managing Scopes.....	987
Viewing Scopes.....	992
Adding Fixed Addresses/Microsoft Reservations.....	993
About Superscopes.....	995
Synchronizing Updates.....	997
Managing Microsoft DHCP Servers.....	998
Viewing Members and Managed DHCP Servers.....	998
Setting Microsoft DHCP Server Properties.....	999
Controlling the DHCP Service of a Microsoft Server.....	1000
Disabling and Removing Microsoft DHCP Servers.....	1000
Modifying DHCP Server Assignments.....	1000

PART 7 MONITORING AND REPORTING

Chapter 35 Monitoring the Appliance 1003

Viewing Status.....	1004
Grid Status.....	1004
Member Status	1004
Viewing the Grid Node Tree.....	1009
Viewing Hardware Status	1010
Monitoring Services	1011
Service Status	1011
Monitoring Grid Services	1011
Monitoring Member Services	1012
Using a Syslog Server	1012
Specifying Syslog Servers	1013
Configuring Syslog for Grid Members	1014
Setting DNS Logging Categories.....	1015
Viewing the Syslog	1016
Searching in the Syslog	1017
Downloading the Syslog File.....	1017
Monitoring Tools	1018
Using the Audit Log.....	1018
Viewing the Replication Status.....	1020
Using the Traffic Capture Tool.....	1021
Using the Capacity Report.....	1022
Participating in the Customer Experience Improvement Program	1023
Monitoring DNS Transactions.....	1024
Viewing DNS Alert Indicator Status	1026
Configuring DNS Alert Thresholds	1026

Chapter 36 DHCP Fingerprint Detection..... 1031

Infoblox DHCP Fingerprint Detection	1032
About DHCP Fingerprints.....	1033
Standard and Custom DHCP Fingerprints	1033
Administrative Permissions	1034
Enabling and Disabling DHCP Fingerprint Detection.....	1034
Configuring DHCP Fingerprints.....	1034
Adding New DHCP Fingerprints.....	1035
Modifying Custom DHCP Fingerprints	1035
Deleting Custom DHCP Fingerprints.....	1036
Viewing DHCP Fingerprint Information.....	1036

Chapter 37 Monitoring with SNMP..... 1037

Understanding SNMP	1038
About SNMPv1 and SNMPv2	1039
About User-Based Security Model in SNMPv3.....	1039
Configuring SNMP	1039
Configuring SNMPv3 Users.....	1040
Modifying SNMPv3 Users	1041
Deleting SNMPv3 Users	1041
Accepting Queries	1041

Adding Trap Receivers	1042
Setting SNMP System Information	1043
Defining Thresholds for Traps	1043
Setting SNMP and Email Notifications	1044
Testing the SNMP Configuration	1047
SNMP MIB Hierarchy	1048
MIB Objects	1049
System Object IDs	1049
Infoblox MIBs	1051
Loading the Infoblox MIBs	1051
ibTrap MIB	1053
ibPlatformOne MIB	1086
ibDHCPone MIB	1100
ibDHCPv6Module	1104
ibDNSOne MIB	1107
IB-DNSSERV-MIB	1111
IB-DNSHITRATIO-MIB	1111
IB-DNSQUERYRATE-MIB	1111
IB-DHCPSEV-MIB	1111
Chapter 38 Infoblox Reporting Solution	1113
Infoblox Reporting Solution	1116
Supported Platforms for Reporting	1119
Infoblox-4030 Supported Reports	1120
Grid Reporting Properties	1121
Configuring General Grid Reporting Properties	1121
Setting Network Port for Reporting	1122
Setting Email Properties	1122
Defining PDF Settings	1122
Defining DNS Query Settings	1122
Configuring the Capture of DNS Queries and Responses	1123
Monitoring Client Queries	1128
Scheduling Report Deliveries	1129
About Reports	1130
Adding New Reports	1131
Adding New Panels to Reports	1131
Cloning Predefined Reports	1132
Modifying User-Defined Reports	1132
Deleting User-Defined Reports	1132
About Searches	1132
Guidelines for Customizing Searches	1133
Reporting Indexes and Update Time Intervals	1134
Cloning Searches	1137
Modifying Searches	1138
Scheduling Searches	1138
Exporting Searches	1139
Scheduling Exports of Search Results	1139
About Alerts	1140
Alerting Logic	1140
Defining Alerts	1141
Previewing Alerting Logic	1141
Defining Advanced Alert Settings	1142

About IP Blocks and IP Block Groups	1142
Adding IP Block Groups	1143
Modifying IP Block Groups	1143
Adding IP Blocks	1143
Modifying IP Blocks	1144
Deleting IP Block Groups and IP Blocks	1144
Exporting IP Block Groups and IP Blocks	1144
Printing IP Block Groups and IP Blocks	1144
Predefined Reports	1145
Predefined Report Categories	1145
Applying Time Filters	1147
Changing Report Formats	1147
DDNS Query Reports	1147
DHCP Lease Reports	1148
DHCP Fingerprint Reports	1149
DHCP Performance Reports	1153
DNS Performance	1155
DNS Query Reports	1156
DNS Last Queried	1160
IPAMv4 Utilization Reports	1162
System Utilization Reports	1165
Security	1166
Managing Reports	1172
Printing Reports	1172
Backing Up Reporting Data	1172
Restoring the Reporting Database	1174

PART 8 GLOBAL LOAD BALANCER INTEGRATION

Chapter 39 Managing Global Load Balancers 1177

Integrating Global Load Balancers	1178
About Load Balancer Synchronization Groups	1180
Requirements and Permissions	1181
Deployment Guidelines	1182
Configuring the Management of GLBs	1183
Setting Usernames and Permissions on GLBs	1183
Setting the Management Mode	1183
Configuring the Synchronization Interface	1184
Assigning Grid Members to Load Balancers	1184
Validating Load Balancer Connection	1186
Managing Global Load Balancers	1187
Setting Load Balancer Properties	1187
Modifying Load Balancer Synchronization Group Properties	1188
Setting Priorities for Managed Load Balancers	1188
Changing the Managing Member or Management Mode	1189
Replicating DNS Data	1189
Backing Up Synchronized Data	1189
Disabling Synchronization	1189
Removing a Load Balancer	1190
Removing a Load Balancer Synchronization Group	1190
Monitoring Global Load Balancers	1190
Viewing Global Load Balancers	1191

Viewing Detailed Status Information	1192
Viewing Detailed Status Information of a Synchronization Group	1192

Chapter 40 Managing Global Load Balancer Data.....1193

About Global Load Balancer Data	1195
Supported Global Load Balancer Objects	1195
Viewing Global Load Balancer Objects	1196
Viewing Traffic Management Structures	1197
Managing Synchronized DNS Data	1198
Mapping DNS Views	1198
Resolving DESYNC Conflicts	1199
Configuring Delegations to Zones on Load Balancers	1199
Modifying Delegations to Zones on Load Balancers	1200
Managing LBDN on GLBs	1201
Associating GLBs with Authoritative Zones	1201
Adding LBDNs	1201
Modifying LBDNs	1202
Viewing LBDNs and LBDN Records	1203
Managing Listeners	1203
Adding Listeners	1203
Modifying Listeners	1204
Managing Data Centers	1204
Adding Data Centers	1204
Modifying Data Centers	1205
Managing Global Load Balancing Pools	1205
Creating Pools and Adding Pool Members	1205
Modifying Load Balancing Pools	1206
Managing Global Load Balancer Servers	1207
Adding GLB Servers	1207
Modifying GLB Servers	1208
Managing Global Load Balancer Virtual Servers	1209
Adding GLB Virtual Servers	1209
Modifying GLB Virtual Servers	1210

PART 9 INFOBLOX INFRASTRUCTURE SECURITY

Chapter 41 Infoblox Advanced DNS Protection.....1213

About Infoblox Advanced DNS Protection	1214
Configuring Advanced DNS Protection	1215
License Requirements	1216
Administrative Permissions	1216
Starting and Stopping Threat Protection Service	1216
Understanding Threat Protection Rules	1217
System and Auto Rules	1217
Custom Rules	1218
About Rule Versions	1220
Using the Events Per Second Rule Setting	1220
Configuring Grid Security Properties	1221
Enabling Multiple DNS Requests through a Single TCP Session	1222
Creating Custom Rules	1222

Managing Threat Protection Rules	1223
Viewing Threat Protection Rules	1223
Enabling and Disabling Rules	1224
Manually Uploading Rule Updates	1224
Publishing Rule Updates	1225
Modifying System and Auto Rules	1226
Modifying Custom Rules	1226
Monitoring Threat Protection Events	1227
Monitoring through Syslog	1227
Threat Protection Statistics Widget	1228
Threat Protection Reports	1229
DNS and Network-Flood Threats	1229

Chapter 42 Infoblox DNS Firewall 1231

About Infoblox DNS Firewall	1233
Setting Up Infoblox DNS Firewall	1234
License Requirements and Admin Permissions	1235
For Local RPZs and RPZ Feeds	1235
For FireEye Integrated RPZs	1235
Best Practices for Configuring RPZs	1236
General RPZ Best Practices	1236
Best Practices For FireEye Integrated RPZs	1236
Enabling Recursion for RPZs	1237
Configuring RPZs for All Recursive Servers	1237
Configuring Local RPZs	1238
Configuring Rules for RPZs	1239
Managing Passthru Rules	1239
Managing Block (No Such Domain) Rules	1240
Managing Block (No Data) Rules	1241
Managing Substitute (Domain Name) Rules	1242
Managing Substitute (Record) Rules	1243
Configuring RPZ Feeds	1248
Configuring the Infoblox RPZ Feed	1249
Infoblox RPZ feeds	1251
Downloading Rules of an RPZ Feed	1252
Testing RPZ Feed Rules	1253
About FireEye Integrated RPZs	1254
Configuring FireEye RPZs	1254
Configuring Rules for FireEye RPZs	1257
Configuring the FireEye appliance	1257
Handling Alerts from the FireEye appliance	1259
Logging FireEye Integrated RPZ messages	1259
Configuration Examples	1259
Managing RPZs	1261
Viewing RPZs	1261
Modifying RPZs	1262
Reordering RPZs	1263
Locking and Unlocking RPZs	1263
Deleting RPZs	1263
Managing RPZ Rules	1264
Viewing RPZ Rules	1264
Modifying RPZ Rules	1265
Deleting RPZ Rules	1265

Copying RPZ Rules	1265
Importing RPZ Rules	1266
Verifying RPZ Configuration	1266
Viewing RPZ in the Syslog	1266
Viewing the Last Updated RPZs	1267

PART 10 REFERENCE

Appendix A Glossary of Terms.....	1271
--	-------------

Appendix B Grid Manager Icons	1279
--	-------------

Appendix C Guidance Documentation Supplement.....	1285
--	-------------

Pre-Requisites	1286
Verifying the Hardware	1286
Security Guidelines	1286
Installation and Configuration	1286
Administration.....	1287
Setting Password Restrictions for Local Admins	1287
Enabling/Disabling Common Criteria Mode	1288
Using the CLI	1288
Licenses and Services.....	1289
WebUI Settings	1290
Creating a Login Banner	1290
Modifying the Session Timeout Setting.....	1290
Managing Certificates	1290
DNS	1290
DNSSEC	1291
Backing Up and Restoring the Database	1291
Audit Log	1292
Syslog	1296

Appendix D Regular Expressions	1303
---	-------------

Supported Expressions for Search Parameters.....	1303
--	------

Appendix E vNIOS Appliance Limitations.....	1305
--	-------------

vNIOS for Riverbed	1306
vNIOS for VMware	1306
vNIOS for Hyper-V	1307

Appendix F Product Compliance	1309
--	-------------

Power Safety Information	1310
AC	1310
DC.....	1310
Agency Compliance	1311
FCC	1311

Canadian Compliance	1311
VCCI	1312
RFC Compliance	1313
DNS RFC Compliance	1313
DHCP RFC Compliance	1315
DHCPv6 RFC Compliance	1316
IDN (Internationalized Domain Names) RFC Compliance	1316

Appendix G Open Source Copyright and License Statements 1317

GNU General Public License	1319
GNU Lesser General Public License	1322
Apache Software License, Version 2.0	1328
perl Artistic License	1333
ISC BIND Copyright	1334
ISC DHCP Copyright	1335
Julian Seward Copyright	1336
Carnegie Mellon University Copyright	1336
Thai Open Source Software Center Copyright	1337
Ian F. Darwin Copyright	1338
Lawrence Berkeley Copyright	1339
MIT Kerberos Copyright	1339
BSD License	1340
David L. Mills Copyright	1341
OpenLDAP License	1341
OpenSSL License	1342
VIM License	1343
ZLIB License	1345
Wietse Venema Copyright	1345
ECLIPSE SOFTWARE	1346
Eclipse Public License - v 1.0	1346
AOP Alliance (Java/J2EE AOP standards)	1350
ASM	1350
Distributed Computing Laboratory, Emory University	1351
COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)	1351
The FreeType Project LICENSE	1355
The Independent JPEG Group's JPEG software	1359
Net-SNMP	1361
The PHP License, version 3.01	1367
INFO-ZIP	1368
MIT License	1370
Ehcache	1370

Appendix H Threat Protection Rules 1371

DNS Cache Poisoning	1372
DNS Message Type	1373
General DDoS	1378
Reconnaissance	1379
DNS Malware	1379
DNS Protocol Anomalies	1380

TCP/UDP Flood.....1381

DNS Tunneling.....1382

DNS Amplification and Reflection.....1383

NTP.....1384

BGP.....1385

OSPF.....1386

ICMP.....1387

Default Pass/Drop.....1392

Index1393



Preface

This preface describes the document conventions of this guide, and provides information about how to find additional product information, including accessing Infoblox Technical Support. It includes the following sections:

- [Document Overview](#) on page 34
 - [Documentation Conventions](#) on page 34
- [What's New](#) on page 36
- [Related Documentation](#) on page 39
- [Customer Care](#) on page 40
 - [User Accounts](#) on page 40
 - [Software Upgrades](#) on page 40
 - [Technical Support](#) on page 40

DOCUMENT OVERVIEW

This guide describes how to configure and manage NIOS appliances using the NIOS 6.10. It was last updated on May 23, 2014. For updated documentation, visit our Support site at <https://support.infoblox.com>.

Documentation Conventions

The text in this guide follows the following style conventions.

Style	Usage
bold	<ul style="list-style-type: none"> Indicates anything that you input in the user interface, by clicking, choosing, selecting, typing, or by pressing on the keyboard. Indicates the field names in the user interface.
input	Signifies command line entries that you type.
<i>variable</i>	Signifies variables typed into the user interface that you need to modify specifically for your configuration. These can be command line variables, file names, and keyboard characters. Indicates the names of the wizards, editors, and dialog boxes in Grid Manager, such as the <i>Add Network</i> wizard or the <i>DHCP Network</i> editor.

Variables

Infoblox uses the following variables to represent values that you type, such as file names and IP addresses.

Variable	Value
<i>a_record</i>	A record
<i>aaaa_record</i>	AAAA record
<i>admin_group</i>	Name of a group of administrators
<i>admin_name</i>	Name of the appliance administrator
<i>addr_range</i>	IP address range
<i>dhcp_template</i>	DHCP template
<i>domain_name</i>	Domain name
<i>directory</i>	Directory name
<i>failover_association</i>	Failover association
<i>filter_name</i>	Name of a DHCP filter
<i>fingerprint</i>	DHCP fingerprint
<i>fixed_address</i>	Fixed address
<i>fixed_address_template</i>	Fixed address template
<i>glb</i>	Global Load Balancer
<i>Grid</i>	Grid name
<i>Grid_master</i>	Grid Master
<i>Grid_member</i>	Grid Member

Variable	Value
<i>hostname</i>	Host name of an independent appliance
<i>host_record</i>	Host record
<i>ifmap_client</i>	IF-MAP client
<i>ip_addr</i>	IPv4 address
<i>lease</i>	IP address of a lease
<i>mac_filter</i>	Name of a MAC filter
<i>match_rule</i>	Name of a match rule
<i>member</i>	Grid member name
<i>ms_server</i>	Microsoft server
<i>named_acl</i>	Named ACL (access control list)
<i>netmask</i>	Subnet mask
<i>network</i>	IP address of a network
<i>network_access_server</i>	Name of a NAS
<i>network_template</i>	Network template
<i>network_view</i>	Network view
<i>option_space</i>	DHCP option space
<i>policy</i>	Name of a policy on RADIUSone
<i>policy_group</i>	Name of a Policy Group
<i>port</i>	Number of a port; predefined for certain protocols
<i>ptr_record</i>	PTR record
<i>reservation</i>	Reservation
<i>roaming_host</i>	Roaming host
<i>scheduled_task</i>	Scheduled task
<i>server_group</i>	Name of a group of servers
<i>shared_network</i>	Shared network
<i>service</i>	One of the services available from Grid Manager
<i>template</i>	DHCP template
<i>dns_view</i>	DNS view
<i>zone</i>	DNS zone

Navigation

Infoblox technical documentation uses an arrow “->” to represent navigation through the user interface. For example, to edit a fixed address, the description is as follows:

From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed_address* check box, and then click the Edit icon.

WHAT'S NEW

The following sections have been updated in Rev. A through Rev. E of this guide:

- **Next Available IP Enhancements**

This release adds the following enhancements for the Next Available IP functionality:

- When multiple users simultaneously request for the next available IP address, only the user who first saves the configuration gets the IP address. Other users can request another IP address or enter a new one.
- To avoid IP address conflicts, the appliance validates the next available IP address to ensure that it is not used for other objects or associated with another operation, such as a scheduled task or an approval workflow.

For more information, see [About the Next Available Network or IP Address](#) on page 844.

- **TR-800 Appliance and vNIOS Virtual Appliances for Reporting**

Infoblox adds support for the Trinzic 800 reporting appliance (TR-800 with 2 GB daily reporting limit) and the following vNIOS reporting virtual appliances: IB-VM-800 (with 1 GB or 2 GB daily reporting limit) and IB-VM-1400 (with 5 GB daily reporting limit). You can deploy the vNIOS virtual appliances as reporting servers in an Infoblox Grid. They are designed to run on VMware ESX or ESXi 4.1, 5.0, and 5.1 servers. For more information, refer to the *Installation Guide for Infoblox 800 Series Appliances* and the *Installation Guide for vNIOS Reporting Virtual Appliances*.

- **Pre-Provisioning Enhancements**

The pre-provisioning feature now supports RPZs (Response Policy Zones) and FireEye integrated zones. It also enhances the hardware selection options in Grid Manager. For more information, see [Pre-Provisioning NIOS Appliances](#) on page 252.

- **Microsoft Windows Server 2012 R2 Support**

Infoblox adds support for Microsoft Management and GSS-TSIG on Microsoft Windows Server 2012. For more information, see [Managing Microsoft Windows Servers](#) on page 953 and [About GSS-TSIG](#) on page 710.

- **F5 TMOS 11.4 Support**

Infoblox adds support for Global Load Balancers that run TMOS version 11.4.x. For more information, see [Managing Global Load Balancers](#) on page 1177.

- **Auto-Provisioning NIOS Appliances**

You can now set up a NIOS appliance using the auto-provisioning feature, which allows a DHCP server to automatically assign an IP address to the appliance. You can then configure and join the auto-provisioned appliance to the Grid. For more information, see [Auto-Provisioning NIOS Appliances](#) on page 250.

- **Pre-Provisioning NIOS Appliances**

Before joining a member to the Grid, you can now enable provisional licenses and make necessary configurations on the offline member and associate DNS and DHCP data with the member prior to its deployment. For more information, see [Pre-Provisioning NIOS Appliances](#) on page 252.

- **Support for NTP Anycast**

When you configure DNS anycast on an appliance and use it as an NTP server, the appliance can answer NTP requests through the anycast IP address. For more information, see [About Anycast Addressing for DNS](#) on page 761.

- **Infoblox Advanced DNS Protection**

The Infoblox Advanced DNS solution employs hardware-accelerated security rules to detect, report upon, and stop DoS (Denial of Service), DDoS (Distributed Denial of Service) and other network attacks targeting DNS caching and authoritative applications. This feature helps minimize “false positives” and ensures that your mission-critical DNS services continue to function even when under attack. Advanced DNS Protection is designed to provide visibility and protection against network floods and DNS threats. It detects DNS attacks through predefined and custom threat protection rules, and mitigates DNS threats by dropping problematic packets while responding only to legitimate traffic. With valid licenses installed, you can subscribe to automatic rule updates that deliver near real-time protection against new and emerging attacks. You may also manually perform the rule update process based on your configuration. Advanced DNS Protection runs on Infoblox

Advanced Appliances that support subscriber-facing DNS caching and external DNS authoritative applications. It supports both IPv4 and IPv6 and can be enabled on the Infoblox-4030 Rev-2 appliance and the following Infoblox Advanced Appliances: PT-1400, PT-2200 and PT-4000. For more information, see [About Infoblox Advanced DNS Protection](#) on page 1214.

- **Infoblox Network Insight**

Network Insight delivers network intelligence by integrating (in real-time) DNS, DHCP, and IPAM data with network infrastructure data, providing visibility across your entire network. The collection and correlation of this data enables network administrators to easily gather necessary information, analyze it, then take appropriate actions to better manage their networks, validate designs, effectively provision, troubleshoot and deliver network services. Network Insight improves decision making, reduces security and service interruption risk, and breaks down operational silos. For more information, see [About Network Insight](#) on page 519.

- **FireEye Adapter for Infoblox DNS Firewall**

Infoblox DNS firewall now provides a mechanism to further protect your network from malware and Advanced Persistent Threats through the integration of FireEye appliances. When your NIOS appliance is properly integrated with a FireEye appliance, it receives periodic alerts from the FireEye appliance when it identifies such threats. Based on your configuration, the NIOS appliance translates these alerts into RPZ rules that not only further protect your network from malicious attacks, but also aid in identifying clients that have been compromised. For more information, see [About FireEye Integrated RPZs](#) on page 1254.

- **Infoblox DNS Firewall Enhancements**

This release supports the following DNS Firewall enhancements:

- New *Response Policy Zone (RPZ) Statistics* widget in the Dashboard
- New *DNS Top RPZ Hits* report
- New fields added to the Local RPZ Rulesets panel

For more information, see [Infoblox DNS Firewall](#) on page 1231.

- **Support for VLAN Tagging**

You can now assign VLANs (Virtual Local Area Networks) to the LAN1, LAN2, and VIP (for HA pairs) interfaces on the appliance so the appliance can provide DNS service to different subnetworks on the same interface. You can set up VLANs on these interfaces to provide segmentation services to address issues such as scalability, security, and network management. This feature is currently supported on the following Infoblox appliances: Trinzic 2210, Trinzic 2220, and Infoblox-4010. For information about these appliances, refer to the respective installation guides on the Infoblox Support web site. For more information about VLAN tagging, see [About Virtual LANs](#) on page 346.

- **Support for Disabling EDNS0**

The NIOS appliance supports EDNS0 (Extension Mechanisms for DNS) by default. As defined in RFC 2671, EDNS0 provides extended UDP packet size that supports additional DNS functionality, such as DNSSEC. On occasions in which you have end servers that do not support EDNS0 and you want to ensure that they respond to recursive queries from the NIOS appliance while improving DNS performance, you can disable EDNS0 for the Grid and individual Grid members. For more information, see [Using Extension Mechanisms for DNS \(EDNS0\)](#) on page 564.

Note: When you disable EDNS0, all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled.

- **Max Cache and Max Negative Cache TTL Configuration**

For recursive DNS servers, you can now specify the maximum cache TTL value that establishes the time limit for the name server to cache positive responses, and the maximum negative cache TTL value that specifies the time limit for the name server to cache negative responses. For more information, see [Specifying Max Cache TTL and Max Negative Cache TTL Settings](#) on page 559.

- **Global NAC Filters**

You can now disable all NAC filters that specify authentication results from a remote, backend RADIUS server so you can perform maintenance on it. When you disable NAC filters for the Grid, the appliance bypasses evaluations of all NAC filters, and there are no configuration changes, service down times, or service restarts on the Grid. For more information, see [About NAC Filters](#) on page 941.

- **DNS Top RPZ Hits by Client Report**

This release adds the *DNS Top RPZ Hits by Client* report, which lists the total number of RPZ hits from a client during a specified time interval, irrespective of the rules and mitigation actions. You can view the IP address of the client, total RPZ hits, and the date and time during which the hits were received. For more information, see [DNS Top RPZ Hits by Clients](#) on page 1167.

- **DHCP Enabled Host Address Permission**

This release supports new admin permissions for IPv4 and IPv6 host addresses. Limited-access users now have the ability to create, modify, and delete IPv4 and IPv6 DHCP enabled host addresses in a specified network when granted read-write permission to IPv4 Host Address or IPv6 Host Address. For more information, see [Administrative Permissions for IPv4 or IPv6 DHCP Enabled Host Addresses](#) on page 209.

- **Support for Scheduled Local Backups**

To avoid missing a backup when a remote server is unavailable during a scheduled automatic backup, you can now select to save a local copy of the backup on your appliance while backing up to the remote server. For more information, see [Backing Up Files](#) on page 423.

- **Support for Link Local Address as the IPv6 Default Gateway**

You can now define a link-local address as the default IPv6 gateway and isolate the LAN segment so the local router can provide global addressing and access to the network and Internet. This is supported for both LAN1 and LAN2 interfaces as well as LAN1 and LAN2 in the failover mode. For more information, see [Configuring VLANs](#) on page 347.

- **Captive Portal Modifications**

This release supports the following Captive Portal enhancements:

- Support for mobile browsers.
- Addition of detected devices in IPv4 filters.

For more information about Captive Portal, see [About the Captive Portal](#) on page 923.

- **Infoblox RESTful Web API**

This release adds newly supported objects to the Infoblox RESTful Web API. For more information, refer to the *Infoblox API Documentation*.

- **DHCP Fingerprint Enhancements**

This release supports the following DHCP fingerprint enhancements:

- Support for IPv6 mobile devices and DHCP fingerprint configuration for IPv6.
- New reports to further identify top devices and device trends.
- Reporting on devices whose fingerprints have changed, which could indicate potential MAC spoofing attacks.
- A new status dashboard (Mobile Device Status) to track active leases of mobile devices.

For more information, see [Infoblox DHCP Fingerprint Detection](#) on page 1032.

RELATED DOCUMENTATION

Other Infoblox appliance documentation:

- *Infoblox CLI Guide*
- *Infoblox API Documentation*
- *Infoblox CSV Import Reference*
- *Infoblox Installation Guide for the Trinzic 100 Appliance*
- *Infoblox Installation Guide for the Trinzic 800 Series and Network Insight ND-800 Platforms*
- *Infoblox Installation Guide for the Trinzic 1400 Series and Network Insight ND-1400 Platforms*
- *Infoblox Installation Guide for the Trinzic 2200 Series and Network Insight ND-2200 Platforms*
- *Infoblox Installation Guide for the Infoblox-4010 and Network Insight ND-4000 Platforms*
- *Infoblox Installation Guide for Infoblox Advanced Appliance PT-1400*
- *Infoblox Installation Guide for Infoblox Advanced Appliance PT-2200*
- *Infoblox Installation Guide for Infoblox Advanced Appliance PT-4000*
- *Infoblox Installation Guide for the Infoblox-4030 Appliance*
- *Infoblox DNS Caching Acceleration Application Guide*
- *Infoblox Installation Guide for the Trinzic Reporting 1400 Appliance*
- *Infoblox Installation Guide for the Trinzic Reporting 2000 Appliance*
- *Infoblox Installation Guide for the Trinzic Reporting 2200 Appliance*
- *Infoblox Installation Guide for the Trinzic Reporting 4000 Appliance*
- *Infoblox Installation Guide for the Infoblox-250-A Appliance*
- *Infoblox Installation Guide for the Infoblox-550-A, -1050-A, -1550-A, and -1552-A Appliances*
- *Infoblox Installation Guide for the Infoblox-1852-A Appliance*
- *Infoblox Installation Guide for the Infoblox-2000-A Appliance*
- *Infoblox Installation Guide for vNIOS Software on Riverbed Services Platforms*
- *Infoblox Installation Guide for Installing vNIOS Software on Cisco Platforms*
- *Infoblox Installation Guide for vNIOS Software on VMware*
- *Infoblox Installation Guide for vNIOS on Microsoft 2008 R2 for Hyper-V*
- *Quick Start Guide for Installing vIBOS Software on VMware Platforms*
- *Infoblox IBOS Administrator Guide*
- *Infoblox Safety Guide*

To provide feedback on any of the Infoblox technical documents, please e-mail techpubs@infoblox.com.

CUSTOMER CARE

This section addresses user accounts, software upgrades, licenses and warranties, and technical support.

User Accounts

The Infoblox appliance ships with a default user name and password. Change the default `admin` account password immediately after the system is installed to safeguard its use. Make sure that the NIOS appliance has at least one administrator account with superuser privileges at all times, and keep a record of your account information in a safe place. If you lose the `admin` account password, and did not already create another superuser account, the system will need to be reset to factory defaults, causing you to lose all existing data on the NIOS appliance. You can create new administrator accounts, with or without superuser privileges. For more information, see *Managing Administrators* on page 149.

Software Upgrades

Software upgrades are available according to the Terms of Sale for your system. Infoblox notifies you when an upgrade is available. Register immediately with Infoblox Technical Support at <http://www.infoblox.com/support/customer/evaluation-and-registration> to maximize your Technical Support.

Technical Support

Infoblox Technical Support provides assistance via the Web, e-mail, and telephone. The Infoblox Support web site at <https://support.infoblox.com> provides access to product documentation and release notes, but requires the user ID and password you receive when you register your product online at: <http://www.infoblox.com/support/customer/evaluation-and-registration>.



PART 1 APPLIANCE GUI

This section introduces you to Grid Manager, the web interface through which you can manage your DNS, DHCP and IP address management (IPAM) infrastructure. It also describes the Task and Status Dashboards, your home page on Grid Manager, and Smart Folders, which you can use to organize your data.

It includes the following chapters:

- [Chapter 1, *Infoblox Grid Manager*](#), on page 43
- [Chapter 2, *Dashboards*](#), on page 97
- [Chapter 3, *Smart Folders*](#), on page 139



Chapter 1 Infoblox Grid Manager

This chapter lists requirements for the management system you use to access the NIOS appliance. It also explains how to access the Grid Manager web interface and describes its major components. This chapter includes the following sections:

- [*Management System Requirements*](#) on page 46
 - [*Supported Browsers*](#) on page 46
 - [*Browser Limitations*](#) on page 47
- [*About Grid Manager*](#) on page 48
 - [*Admin Permissions for Grid Manager*](#) on page 48
 - [*Logging in to the GUI*](#) on page 48
- [*Setting Login Options*](#) on page 49
 - [*Specifying the Grid Name and Hostname*](#) on page 49
 - [*Creating a Login Banner*](#) on page 49
 - [*Changing the Password and Email Address*](#) on page 50
 - [*Specifying the Table Size*](#) on page 50
 - [*Selecting Your Home Page*](#) on page 51
 - [*Setting the Browser Time Zone*](#) on page 51
- [*SSL \(Secure Sockets Layer\) Protocol*](#) on page 52
- [*Managing Certificates*](#) on page 53
 - [*About HTTPS Certificates*](#) on page 53
 - [*About Client Certificates*](#) on page 55
- [*About the Grid Manager Interface*](#) on page 58
 - [*System Messages*](#) on page 58
 - [*Security and Informational Banners*](#) on page 59
 - [*Breadcrumbs Navigation*](#) on page 59
 - [*Global Search*](#) on page 59
 - [*Finder Panel*](#) on page 59
 - [*Toolbar Panel*](#) on page 59
 - [*Help Panel*](#) on page 59
 - [*Wizards and Editors*](#) on page 60
 - [*Tooltips*](#) on page 60
 - [*Customizing Tables*](#) on page 60

- [Selecting Objects in Tables](#) on page 60
- [Modifying Data in Tables](#) on page 62
- [Finding and Restoring Data](#) on page 63
 - [Using Bookmarks](#) on page 63
 - [Using the Recycle Bin](#) on page 64
 - [Managing Third Party URL Links](#) on page 66
 - [Using Filters](#) on page 67
 - [Using Quick Filters](#) on page 68
 - [Using Global Search](#) on page 69
 - [Using the Go To Function](#) on page 71
- [About Tasks](#) on page 72
 - [Viewing Tasks](#) on page 72
 - [Supported Objects for Scheduled and Approval Tasks](#) on page 73
 - [Guidelines for Upgrading, Backing Up, and Restoring Data](#) on page 74
- [Scheduling Tasks](#) on page 75
 - [Scheduling Additions and Modifications](#) on page 75
 - [Scheduling Appliance Operations](#) on page 76
 - [Scheduling Deletions](#) on page 76
 - [Scheduling Recursive Deletions of Network Containers and Zones](#) on page 76
 - [Viewing Scheduled Tasks](#) on page 77
 - [Rescheduling Tasks](#) on page 78
 - [Canceling Scheduled Tasks](#) on page 79
- [Configuring Approval Workflows](#)
 - [Creating Approval Workflows](#) on page 81
 - [Viewing Approval Workflows](#) on page 82
 - [Modifying Approval Workflows](#) on page 83
 - [Deleting Approval Workflows](#) on page 83
 - [Viewing Approval Tasks](#) on page 83
 - [Viewing Workflow Notifications](#) on page 83
 - [Unsupported Operations for Submitters](#) on page 84
- [About Long Running Tasks](#) on page 85
 - [Running Tasks in the Background](#) on page 85
 - [Monitoring Long Running Tasks](#) on page 86
- [About CSV Import](#) on page 86
 - [CSV Import Limitations](#) on page 87
 - [Creating a Data File for Import](#) on page 89
 - [Exporting Data to Files](#) on page 89
 - [Configuring Import Options](#) on page 89
 - [Managing CSV Imports](#) on page 90
- [Exporting Displayed Data](#) on page 91
- [Printing from Grid Manager](#) on page 91
- [Multilingual Support](#) on page 92
 - [UTF-8 Supported Fields](#) on page 92
 - [UTF-8 Support Limitations](#) on page 92

-
- [*Support for Internationalized Domain Names*](#) on page 93
 - [*Decoding IDNs and Encoding Punycode*](#) on page 93
 - [*IDN Supported Fields*](#) on page 93
 - [*IDN Support Limitations*](#) on page 94

MANAGEMENT SYSTEM REQUIREMENTS

The management system is the computer from which you configure and manage the NIOS appliance. The management system must meet the following requirements.

Figure 1.1 Software and Hardware Requirements for the Management System

Management System Software Requirements	Management System Hardware Requirements
GUI ACCESS <ul style="list-style-type: none"> See Supported Browsers on page 46 for details. CLI ACCESS <ul style="list-style-type: none"> Secure Socket Shell (SSH) client that supports SSHv2 Terminal emulation program, such as minicom or Hilgraeve Hyperterminal® 	<ul style="list-style-type: none"> Minimum System: 1.4 GHz CPU with 1 GB RAM available to the product GUI, and 256 Kbps connectivity to NIOS appliance Recommended System: 2.0 GHz (or higher) dual core CPU with 2 GB RAM available for the product GUI, and network connectivity to NIOS appliance Monitor Resolution: 1024 x 768 (minimum) 1280 x 800 or better (recommended)

Supported Browsers

Grid Manager supports the following operating systems and browsers. You must install and enable Javascript for Grid Manager to function properly. Grid Manager supports only SSL version 3 and TLS version 1 connections.

Infoblox recommends that you use a computer that has a 2 GHz CPU and at least 1 GB of RAM.

Infoblox supports the following browsers for Grid Manager:

Operating System	Supported Browser
Microsoft Windows 8®	Microsoft Internet Explorer® 11.x*, 10.x* Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x
Microsoft Windows 7®	Microsoft Internet Explorer® 11.x*, 10.x*, 9.x, and 8.x Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x
Microsoft Windows XP®(SP2+)	Microsoft Internet Explorer 7.x and 8.x Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x
Red Hat® Enterprise Linux® 6.x	Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x
Red Hat® Enterprise Linux 5.x	Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x
Apple® Mac OS X 10.8.x	Safari 6.x Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x

Operating System	Supported Browser
Apple® Mac OS X 10.7.x	Safari 5.x Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x
Apple® Mac OS X 10.6.x	Safari 5.x Mozilla Firefox 25.x, 21.x, 16.x, and 10.x Google Chrome 30.x, 27.x, 22.x, and 16.x

Note: * Grid Manager fully supports Microsoft Internet Explorer® 11.x and 10.x when you enable compatibility view in the browser. Features in the **Reporting** tab may not function properly if you disable compatibility view. In the browser, go to **Tools** -> **Compatibility View** to enable the feature.

Infoblox recommends using the latest release of the supported versions of Internet Explorer, Mozilla Firefox or Google Chrome for best performance.

Browser Limitations

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

ABOUT GRID MANAGER

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

- Use Smart Folders to organize your data based on criteria you specify. For information, see [Smart Folders](#) on page 139.
- The network and IP address maps and lists provide views of your networks and IP addresses, so you can quickly evaluate IP address usage and understand how your network resources are being utilized. You can quickly determine which IP addresses are in use, when they were allocated, and to which devices they were assigned. For information, see [Chapter 12, IP Address Management](#), on page 457.
- Customize the Dashboard to monitor your Grid and networks. The Dashboard also provides access to frequently-used commands and the network discovery feature. You can run network discoveries to identify IP address conflicts and troubleshoot network issues. For information, see [Dashboards](#) on page 97.
- Tools such as the *Finder* panel, filters, and global search help you quickly find the information you need. For information, see [About the Grid Manager Interface](#) on page 58.
- Use wizards to quickly create new networks and resource records. Editors allow you to configure additional operational parameters. For information, see [Wizards and Editors](#) on page 60.

Before you can use Grid Manager, you must install and configure the NIOS appliance as described in the installation guide that shipped with your product. You can then access Grid Manager using one of the supported browsers. For information, see [Supported Browsers](#) on page 46.

Admin Permissions for Grid Manager

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

Note that superusers must configure admin groups and accounts in the Grid Manager application of the NIOS appliance. In Grid Manager, superusers can set and change permissions for specific objects, such as IPv4 networks, IPv6 networks, and resource records. For information about user accounts and administrative permissions, see [Managing Administrators](#) on page 149.

Logging in to the GUI

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

To log in to Grid Manager:

1. Open an Internet browser window and enter **https://<IP address or hostname of your NIOS appliance>**. The Grid Manager login page appears. For information, see [Supported Browsers](#) on page 46.
2. Enter your user name and password, and then click **Login** or press Enter. The default user name is **admin** and the default password is **infoblox**. Note that if your password expired or was reset by a superuser, you may be required to enter a new password.

If you are a smart card user and two-factor authentication is enabled on the appliance, your user name, which is the same as your CN (common name) in the client certificate, appears automatically. Enter the password you use to log in to the user account. For information about two-factor authentication, see [Authenticating Admins Using Two-Factor Authentication](#) on page 187.

3. Read the Infoblox End-User License Agreement. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.

Click **I Accept**. The *Grid Setup* wizard appears.

SETTING LOGIN OPTIONS

Grid Manager provides several options that you can set to facilitate the login process. Additionally, you can manage CA (Certificate Authority) and server certificates on the NIOS appliance. You can import certificates, select and view their details, or remove them. To manage certificates, see [Managing Certificates](#) on page 53.

Specifying the Grid Name and Hostname

To define the default hostname that appears when the login prompt displays:

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

Creating a Login Banner

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

To create a login banner:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Grid Properties -> Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Security** tab, and then select **Enable Login Banner**. In the text field, enter the text that you want displayed on the login screen.
3. Save the configuration and click **Restart** when it appears at the top of the screen.

Changing the Password and Email Address

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

Note that when multiple users log in to Grid Manager using the same admin account, they share the same user profile and preference settings, such as the widget, table size and column settings, independent of their browser settings. Instead of using the same admin account for multiple users, you can add multiple users to the same admin group so they can share the same permissions. For information about configuring admin accounts and admin groups, see [Managing Administrators](#) on page 149.

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

To change your password and email address:

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

Specifying the Table Size

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

To specify table size:

1. Select any tab in Grid Manager, and then click **User Profile** from the Toolbar.
2. In the *User Profile* editor, complete the following:
 - **Table Size:** Specify the number of lines of data you want a table or a single list view to contain. You can set the number of lines from 10 to 256. The default is 20.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Selecting Your Home Page

When you first log in to Grid Manager, the Tasks Dashboard is your home page. You can change your home page for subsequent logins. You can specify the maximum number of widgets that can be configured per dashboard. You can set up to 20 widgets per dashboard.

To change your home page:

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

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

Setting the Browser Time Zone

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

To manually set the time zone of your browser:

1. Select any tab in Grid Manager, and then click **User Profile** from the Toolbar.
The *User Profile* editor displays your user name, user type, and admin group.
2. In the *User Profile* editor, complete the following:
 - **Time Zone:** Select the time zone Grid Manager uses to convert all displayed time values. The default is **Auto-detect time zone**. You must select a specific time zone when Grid Manager cannot automatically detect the time zone of your browser.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

SSL (SECURE SOCKETS LAYER) PROTOCOL

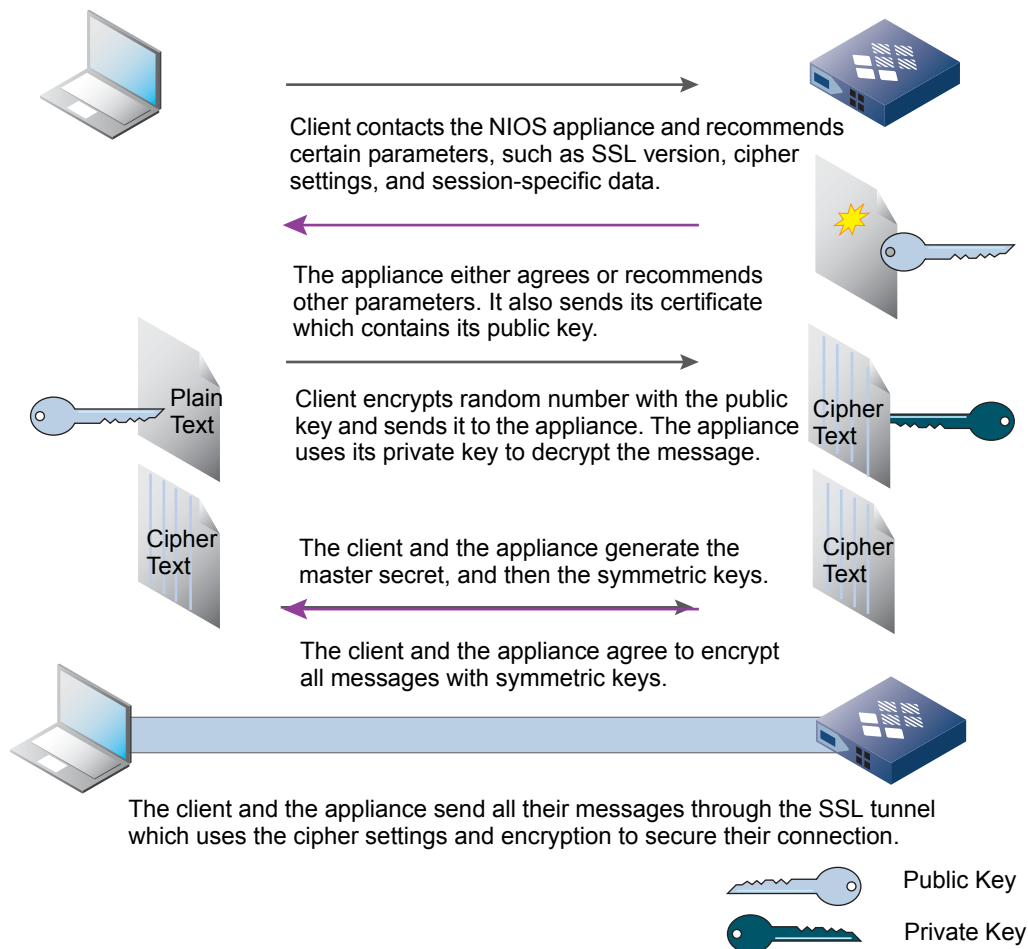
When you log in to the NIOS appliance, your computer makes an HTTPS (Hypertext Transfer Protocol over Secure Sockets Layer protocol) connection to the NIOS appliance. HTTPS is the secure version of HTTP, the client-server protocol used to send and receive communications throughout the Web. HTTPS uses SSL (Secure Sockets Layer) to secure the connection between a client and server. SSL provides server authentication and encryption. The NIOS appliance supports SSL versions 2 and 3.

When a client first connects to a server, it starts a series of message exchanges, called the SSL handshake. During this exchange, the server authenticates itself to the client by sending its server certificate. A certificate is an electronic form that verifies the identity and public key of the subject of the certificate. (In SSL, the subject of the certificate is the server.) Certificates are typically issued and digitally signed by a trusted third party, the Certificate Authority (CA). A certificate contains the following information: the dates it is valid, the issuing CA, the server name, and the public key of the server. For information about certificates, see [Managing Certificates](#) on page 53.

A server generates two distinct but related keys: a public key and a private key. During the SSL handshake, the server sends its public key to the client. Once the client validates the certificate, it encrypts a random value with the public key and sends it to the server. The server decrypts the random value with its private key.

The server and the client use the random value to generate the master secret, which they in turn use to generate symmetric keys. The client and server end the handshake when they exchange messages indicating that they are using the symmetric keys to encrypt further communications.

Figure 1.2 SSL Handshake



To avoid possible attacks in which HTTP or HTTPS connections are made to a web server and stay open much longer than they should be, Infoblox provides the `set connection_limit` and `show connection_limit` CLI commands that you can use to mitigate these attacks. In general, these attacks can result in the web server reaching its maximum number of concurrent connections, and thus denying connections from legitimate sources. You can use the CLI commands to limit the number of concurrent HTTP and HTTPS connections from a given client that corresponds to a particular IP address. For information about the CLI commands and how to use them, refer to the *Infoblox CLI Guide*.

MANAGING CERTIFICATES

About HTTPS Certificates

The NIOS appliance generates a self-signed certificate when it first starts. A self-signed certificate is signed by the subject of the certificate, and not by a CA (Certificate Authority). This is the default certificate. When your computer first connects to the NIOS appliance, it sends this certificate to authenticate itself to your browser.

Because the default certificate is self-signed, your browser does not have a trusted CA certificate or a cached NIOS appliance server certificate (saved from an earlier connection) to authenticate the NIOS appliance certificate. Also, the hostname in the default certificate is `www.infoblox.com`, which is unlikely to match the hostname of your NIOS appliance. Consequently, messages appear warning that the certificate is not from a trusted certifying authority and that the hostname on the certificate is either invalid or does not match the name of the site that sent the certificate. Either accept the certificate just for this session or save it to the certificate store of your browser.

To eliminate certificate warnings, you can replace the default self-signed certificate with a different certificate that has the hostname of your NIOS appliance. The NIOS appliance supports X.509 certificates in .PEM format. After the initial login, you can do one of the following:

- Generate another self-signed certificate with the correct hostname and save it to the certificate store of your browser.
- Request a CA-signed certificate with the correct hostname and load it on the NIOS appliance. For more information, see [Generating Certificate Signing Requests](#) on page 54.
- When you receive the certificate from the CA, import it to the appliance, as described in [Uploading Certificates](#) on page 54.
- Download the certificate from a trusted CA, as described in [Downloading Certificates](#) on page 55.

Generating Self-Signed Certificates

You can replace the default certificate with a self-signed certificate that you generate. When you generate a self-signed certificate, you can specify the correct hostname and change the public/private key size, enter valid dates and specify additional information specific to the NIOS appliance. If you have multiple appliances, you can generate a certificate for each appliance with the appropriate hostname.

To generate a self-signed certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click **Certificates** -> **HTTPS Cert** -> **Generate Self-signed Certificate** from the Toolbar. In a Grid, ensure that you select the Grid Master when generating a self-signed certificate.
or
Infoblox Orchestration Server with the dedicated certificate feature enabled: From the **Data Management** tab, select the **IF-MAP** tab, and then click **IF-MAP Service Certificate** -> **Generate Self-signed Certificate** from the Toolbar.
2. In the *Generate Self-Signed Certificate* dialog box, complete the following:
 - **Key Size:** Select either **2048** or **1024** for the length of the public key.
 - **Days Valid:** Specify the validity period of the certificate.
 - **Common Name:** Specify the domain name of the NIOS appliance. You can enter the FQDN (fully qualified domain name) of the appliance.

- **Organization:** Enter the name of your company.
 - **Organizational Unit:** Enter the name of your department.
 - **Locality:** Enter a location, such as the city or town of your company.
 - **State or Province:** Enter the state or province.
 - **Country Code:** Enter the two-letter code that identifies the country, such as US.
 - **Admin E-mail Address:** Enter the email address of the appliance administrator.
 - **Comment:** Enter information about the certificate.
3. Click **OK**.
 4. If the appliance already has an existing HTTPS certificate, the new certificate replaces the existing one. In the *Replace HTTPS Certificate Confirmation* dialog box, click **Yes**. The appliance logs you out, or you can manually log out. When you log in to the appliance again, it uses the new certificate you generated.

Generating Certificate Signing Requests

You can generate a CSR (certificate signing request) that you can use to obtain a signed certificate from your own trusted CA. Once you receive the signed certificate, you can import it in to the NIOS appliance, as described in [Uploading Certificates](#) on page 54.

To generate a CSR:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click **Certificates** -> **HTTPS Cert** -> **Create Signing Request** from the Toolbar.
or
Infoblox Orchestration Server with the dedicated certificate feature enabled: From the **Data Management** tab, select the **IF-MAP** tab, and then click **IF-MAP Service Certificate** -> **Create Signing Request** from the Toolbar.
2. In the *Create Certificate Signing Request* dialog box, enter the following:
 - **Key Size:** Select either 2048 or 1024 for the length of the public/private key pair.
 - **Common Name:** Specify the domain name of the NIOS appliance. You can enter the FQDN of the appliance.
 - **Organization:** Enter the name of your company.
 - **Organizational Unit:** Enter the name of your department.
 - **Locality:** Enter a location, such as the city or town of your company.
 - **State or Province:** Enter the state or province.
 - **Country Code:** Enter the two-letter code that identifies the country, such as US.
 - **Admin E-mail Address:** Enter the email address of the appliance administrator.
 - **Comment:** Enter information about the certificate.
3. Click **OK**.

Uploading Certificates

When you receive the certificate from the CA, and import it to the appliance, the NIOS appliance finds the matching CSR and takes the private key associated with the CSR and associates it with the newly imported certificate. The appliance then automatically deletes the CSR.

If the CA sends an intermediate certificate that must be installed along with the server certificate, you can upload both certificates to the appliance. The appliance supports the use of intermediate certificates to complete the chain of trust from the server certificate to a trusted root CA. This eliminates intermediate certificate security warnings that appear when you open a web browser and try to connect to an Infoblox appliance.

To import a certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click **Certificates** -> **HTTPS Cert** -> **Upload Certificate** from the Toolbar.
or
Infoblox Orchestration Server with the dedicated certificate feature enabled: From the **Data Management** tab, select the **IF-MAP** tab, and then click **IF-MAP Service Certificate** -> **Upload Certificate** from the Toolbar.
2. Navigate to where the certificate is located and click **Open**.
3. If the appliance already has an existing HTTPS certificate, the new certificate replaces the existing one. In the *Replace HTTPS Certificate Confirmation* dialog box, click **Yes**.
The appliance imports the certificate and logs you out. When you log in to the appliance again, it uses the certificate you imported.

Downloading Certificates

You can download the current certificate or a self-signed certificate, as described in [Generating Self-Signed Certificates](#) on page 53.

To download a certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click **Certificates** -> **HTTPS Cert** -> **Download Certificate** from the Toolbar.
or
Infoblox Orchestration Server with the dedicated certificate feature enabled: From the **Data Management** tab, select the **IF-MAP** tab, and then click **IF-MAP Service Certificate** -> **Download Certificate** from the Toolbar.
2. Navigate to where you want to save the certificate, enter the file name, and then click **Save**.

About Client Certificates

You can generate client certificates for a Grid Master or a Grid Master candidate, and then send it to another server, such as a Hardware Security Module (HSM).

Generating a Client Certificate

To generate a client certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.
Grid Master Candidate: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box.
2. From the Toolbar, click **Certificates** -> **Client Cert** -> **Generate Client Certificate**, and select either **RSASHA1** or **RSASHA256**.
 - If you are generating a certificate for an HSM group with SafeNet Luna SA 4 devices, you must select **RSASHA1**; and if the certificate is for an HSM group with SafeNet Luna SA 5 devices, select **RSASHA256**.

The appliance displays a confirmation dialog after it generates the certificate. If a certificate had been previously generated, the appliance displays a dialog warning that if the previous certificate was registered with a server, then the new certificate must be registered with the server.

Viewing Client Certificates

To view the client certificates that were generated:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.
Grid Master Candidate: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box.
2. From the Toolbar, click **Certificates** -> **Client Cert** -> **View Client Certificate**, and select either **RSASHA1** or **RSASHA256**.

The appliance displays the selected certificate.

Downloading Client Certificates

To download a client certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.
Grid Master Candidate: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box.
2. From the Toolbar, click **Certificates** -> **Client Cert** -> **Download Client Certificate**, and select either **RSASHA1** or **RSASHA256**.
3. Save the certificate.

About CA Certificates

If the CA sends an intermediate certificate that must be installed along with the server certificate, you can upload both certificates to the appliance. The appliance supports the use of intermediate certificates to complete the chain of trust from the server certificate to a trusted root CA. This eliminates intermediate certificate security warnings that appear when you open a web browser and try to connect to an Infoblox appliance.

When you configure two-factor authentication for smart card users, ensure that you upload the required CA certificates before you enable the OSCP service. For information about two factor authentication and how to configure it, see [Defining the Authentication Policy](#) on page 185. Only superusers and limited-access users with the required permissions can manage CA certificates. For information about admin permissions, see [Administrative Permissions for OSCP Server Groups and CA Certificates](#) on page 215.

Uploading CA Certificates

To upload a CA-signed certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box.
2. Select **Certificates** -> **Manage CA Certificates** from the Toolbar.
3. In the *CA Certificates* editor, click the Add icon.
4. In the *Upload* dialog box, click **Select** and navigate to the certificate you want to upload.
5. Select the file and click **Upload**.

Note: NIOS can only upload certificates that are in PEM format. A.PEM file can contain more than one certificate. For information about how to convert CA certificates to .PEM format, see [Converting CA Certificates to PEM Format](#) on page 57.

Repeat the steps to add additional CA-signed certificates.

The *CA Certificates* dialog box displays the following information about the intermediate certificates:

- **Issuer:** The name of the trusted CA that issued the certificate.
- **Valid From:** The date from which the certificate becomes valid.
- **Valid To:** The date until which the certificate is valid.
- **Subject:** The name of the certificate.

You can also do the following:

- Select a certificate and click the Delete icon to delete it.
- Print the data or export it in .csv format.

Converting CA Certificates to PEM Format

NIOS can only upload certificates that are in PEM format. PEM files are Base64 encoded ASCII files. You can use OpenSSL to convert other certificate formats, such as P7B and DER, into PEM format.

You can run OpenSSL on Linux and Windows systems. For Linux, OpenSSL is pre-installed. For Windows, you can manually install an OpenSSL for Windows. For information about OpenSSL, visit its web site at <http://www.openssl.org/>.

To convert a P7B file to PEM format using OpenSSL:

1. Download and unzip the CA certificate file in P7B format.
2. Navigate to the directory where you unzip the CA certificate file.
3. Identify the PKCS7 directory.
4. Use the following OpenSSL command to convert the P7B file to PEM format:

```
$ openssl pkcs7 -in xxxx.p7b -print_certs -out yyyy.pem
```

where xxxx is the name of the P7B file and yyyy is the name of the converted PEM file.

To convert a DER file to PEM format using OpenSSL:

1. Download and unzip the CA certificate file in DER format.
2. Navigate to the directory where you unzip the CA certificate file.
3. Use the following OpenSSL command to convert the DER file to PEM format:

```
$ openssl x509 -inform DER -outform PEM -in xxxx.cer -out yyyy.pem
```

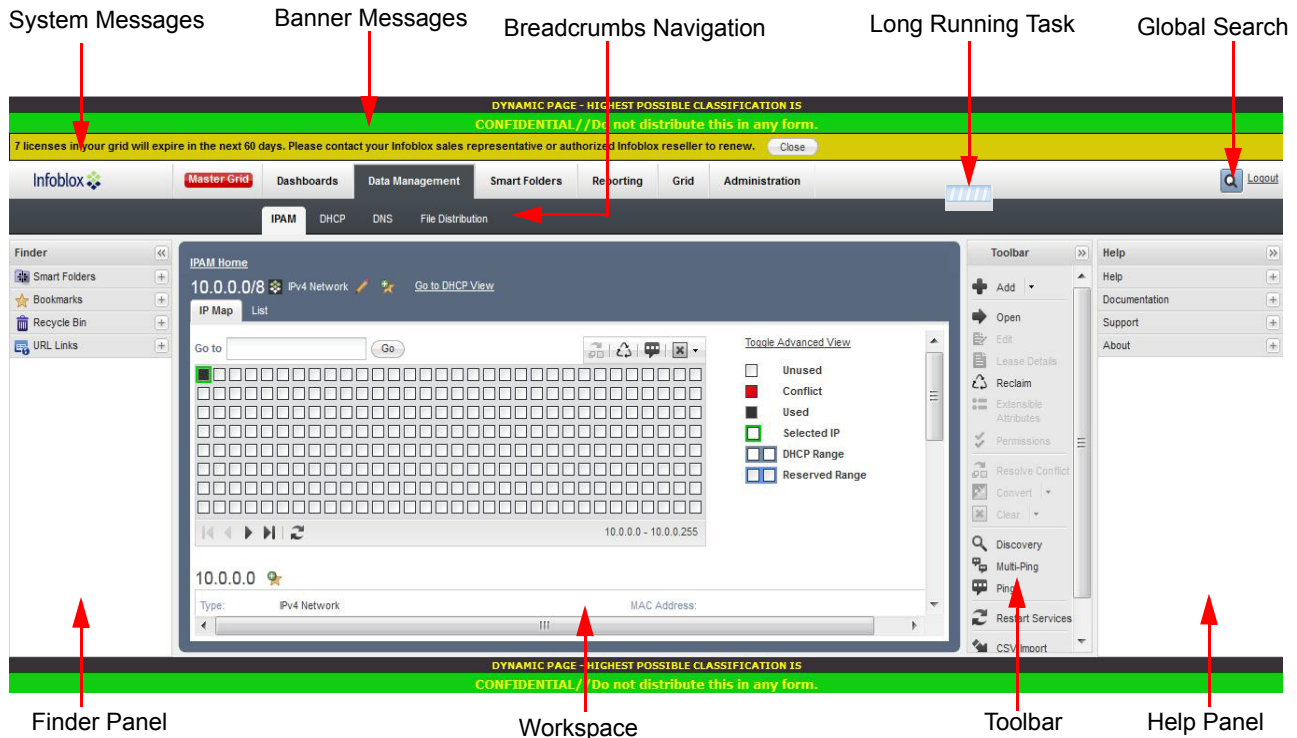
where xxxx is the name of the DER file and yyyy is the name of the converted PEM file.

ABOUT THE GRID MANAGER INTERFACE

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

[Figure 1.3](#) illustrates the typical layout of Grid Manager. It identifies common elements of the interface and features that you can use:

Figure 1.3 Grid Manager Interface



System Messages

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

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

Security and Informational Banners

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

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

For more information, see [Configuring Security Level Banner](#) on page 268 and [Configuring Informational Level Banner](#) on page 269.

Breadcrumbs Navigation

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

Global Search

Use Global Search to find data. Grid Manager searches the entire NIOS database for data that matches the criteria you specify. For additional information on Global Search, see [Using Global Search](#) on page 69.

Finder Panel

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

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

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

Toolbar Panel

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

Help Panel

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

- **Help:** Expand this section to view information about the window currently displayed.
- **Documentation:** Expand this section to download the latest versions of the *Infoblox Administrator Guide* and *Infoblox API Documentation*.
- **Support:** Expand this section to view links to the Infoblox web site and Technical Support site.
- **About:** Expand this section to view information about the NIOS software version.

You can resize, collapse, and expand the *Help* panel. In addition, each dialog box also provides a *Help* panel that contains information specific to the dialog box. You can expand and collapse the *Help* panel in dialog boxes as well.

Wizards and Editors

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

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

Tooltips

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

Customizing Tables

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

To resize columns in a table:

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

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

1. In the table, mouse over to a header title and click the down arrow key.
2. Select **Sort Ascending** or **Sort Descending**.

To edit columns:

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

Grid Manager displays the selected column in the table.

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

Selecting Objects in Tables

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

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

- Use **SHIFT+click** to select multiple contiguous rows.
- Use **CTRL+click** to select multiple non-contiguous rows.
- Click the check box in the table header to select all rows on a page, as shown in [Figure 1.4](#).

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

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

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

Figure 1.4 Select All in a Table

The screenshot shows the Grid Manager interface with the following components and annotations:

- Navigation Bar:** Dashboard, Data Management, Smart Folders, Grid, Administration.
- Sub-Menu:** IPAM, DHCP, DNS, File Distribution.
- Page Header:** default NetworkView, Off, Toggle Filter On, Show Filters.
- Search Bar:** Go to [input field] Go.
- Table:**

Network	Comment	IPAM Utilization	Site
<input checked="" type="checkbox"/> 65.0.0.0/8		78%	Info_site_25
<input checked="" type="checkbox"/> 66.0.0.0/8		78%	Info_site_3
<input checked="" type="checkbox"/> 67.0.0.0/8		78%	Info_site_39
<input checked="" type="checkbox"/> 68.0.0.0/8		78%	Info_site_8
<input checked="" type="checkbox"/> 69.0.0.0/8		78%	Info_site_20
<input checked="" type="checkbox"/> 70.0.0.0/8		0%	Info_site_68
<input checked="" type="checkbox"/> 71.0.0.0/8	default	39%	Info_site_13
<input checked="" type="checkbox"/> 72.0.0.0/8	default	39%	Info_site_72
<input checked="" type="checkbox"/> 73.0.0.0/8	default	39%	Info_site_35
<input checked="" type="checkbox"/> 74.0.0.0/8	default	39%	Info_site_93
<input checked="" type="checkbox"/> 75.0.0.0/8	default	39%	Info_site_92
<input checked="" type="checkbox"/> 76.0.0.0/8	default	39%	Info_site_47
<input checked="" type="checkbox"/> 77.0.0.0/8	default	39%	Info_site_9
<input checked="" type="checkbox"/> 78.0.0.0/8	default	39%	Info_site_58
<input checked="" type="checkbox"/> 79.0.0.0/8	default	39%	Info_site_31
<input checked="" type="checkbox"/> 80.0.0.0/8	default	0%	Info_site_88
<input checked="" type="checkbox"/> 81.0.0.0/8		78%	Info_site_73
<input checked="" type="checkbox"/> 82.0.0.0/8		78%	Info_site_0
<input checked="" type="checkbox"/> 83.0.0.0/8		78%	Info_site_46
- Annotations:**
 - Click this link to select all rows on all pages. (Points to "Select all objects in this dataset")
 - Click this check box to select all rows on this page only. (Points to the first row's check box)
 - Use these navigational buttons to page through the display. (Points to the pagination controls at the bottom)

Modifying Data in Tables

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

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

To modify data in a table:

1. From any panel that supports inline editing, double click the row of data that you want to modify. The appliance displays the inline editing editor in the selected row, as shown in [Figure 1.5](#).
2. Depending on the data type, enter the new data in the field or select an item from the drop-down list. Note that some fields are read-only.
3. Click **Save** to save the changes, or click **Cancel** to discard them.

Figure 1.5 Inline Editing

ices to apply the change.

Dashboard

Data Management

Smart Folders

Grid

Administration

IPAM

DHCP

DNS

File Distribution

default

NetworkView

+

Off

Toggle Filter On

Show Filter

Go to

Go

➡

+

✖

🔍

🔄

🖨

Network	Comment	IPAM Utilization	Disabled	Leaf Network	Building	Country	Region
10.0.0.0/8	comment	0%	No	Yes			
11.0.0.0/8	default	3%		No	Infoblox_build_94	USA_1	South
12.0.0.0/8	default	78%		No	Infoblox_build_17	USA_1	South
13.0.0.0/8				No	Infoblox_build_34	USA_1	South
14.0.0.0/8				No	Infoblox_build_66	USA_1	South
15.0.0.0/8	chnage com...	78%		No	Infoblox_build_94	USA_1	South
16.0.0.0/8		3%		No	Infoblox_build_74	USA_1	South
17.0.0.0/8	default	78%		No	Infoblox_build_36	USA_1	South
18.0.0.0/8		3%		No	Infoblox_build_80	USA_1	South
19.0.0.0/8	default	78%		No	Infoblox_build_65	USA_1	South
20.0.0.0/8		0%	No	Yes	Infoblox_build_23	USA_1	South
21.0.0.0/8		7%		No	Infoblox_build_33	USA_1	South
22.0.0.0/9		0%	No	Yes	Infoblox_build_92	USA_1	South
24.0.0.0/11		0%	No	Yes	Infoblox_build_65	ASIA_1	West

Save

Cancel

FINDING AND RESTORING DATA

Grid Manager provides tools for organizing and quickly retrieving your DNS, DHCP and IP address management data. The *Finder* panel, which appears on all pages in Grid Manager, provides tools for organizing your data. The *Finder* panel provides easy access to the following:

- **Smart Folders:** Contains a hierarchical list of smart folders that are available in My Smart Folders. For information, see [My Smart Folders](#) on page 141.
- **Bookmarks:** Contains bookmarked objects, such as networks and IP addresses. For information, see [Using Bookmarks](#).
- **Recycle Bin:** Contains deleted objects that can be restored or permanently removed. For information, see [Using the Recycle Bin](#) on page 64.
- **URL Links:** Contains a list of third party URLs that you previously added. You can add more URL links, and modify and delete existing URL links. For information, see [Managing Third Party URL Links](#) on page 66.

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

In addition, Grid Manager also provides the following:

- **Filters** to customize data displays. For more information, see [Using Filters](#) on page 67 and [Using Quick Filters](#) on page 68.
- **Global Search** to search the NIOS database for objects that match your criteria. For more information, see [Using Global Search](#) on page 69.
- **Go To** function to quickly locate an object. For more information, see [Using the Go To Function](#) on page 71.

Using Bookmarks

The Bookmarks section displays objects for which you have created bookmarks. You can create bookmarks for objects such as networks, DNS zones, and admin groups. To bookmark an object, navigate to its page and click the Bookmark icon at the top of the page. If you have more than one network view, Grid Manager displays the name of the bookmark with the network view to which the object belongs. For example, when you bookmark IP address 10.128.0.10 in the default network view, Grid Manager displays the bookmark as `default > 12.128.0.10`. However, if you have only one network view, Grid Manager displays only the object name `12.128.0.10`. If you create a bookmark before adding more network views, the bookmark name (without the network view) remains the same. You can rename the bookmark at anytime. You can create only one bookmark for each object, up to 500 objects. When your bookmarks are close to 500, you may want to remove some to create room for new ones.

You can do the following in Bookmarks:

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

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

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

To edit the name of a bookmark:

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

To delete a bookmark:

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

Using the Recycle Bin

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

You must enable the Recycle Bin in Grid Manager to store and restore deleted objects. For information about how to enable and disable the Recycle Bin, see [Enabling and Disabling the Recycle Bin](#). When you use the Recycle Bin, you can restore deleted objects to the active configuration. You can also permanently remove the objects from the Recycle Bin. If you do not enable the Recycle Bin, the appliance immediately removes objects from the database when you delete them using Grid Manager.

On a NIOS appliance, only superusers have permissions to fully manage the Recycle Bin. If you have limited-access permissions, you can view, restore, and permanently remove only the objects that you deleted. On an Infoblox IF-MAP server, only superusers can fully manage the Recycle Bin. Limited-access admins cannot view or restore IF-MAP clients from the Recycle Bin.

You can do the following in the Recycle Bin:

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

Enabling and Disabling the Recycle Bin

To enable or disable the Recycle Bin:

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

Viewing Objects in the Recycle Bin

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

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

To view detailed information about deleted objects:

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

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

- **Name:** The short name of the object. For example, the short names for fixed addresses and reservations are their IP addresses.
- **Type:** The object type.
- **Parent/Container:** The parent object or parent container to which the object belongs.
- **Admin:** The admin name of the user who deleted the object.
- **Data:** The data that the object contains, if any.

- **Network View:** The network view to which the object belongs.
- **Time:** The time stamp when the object was deleted.

To close the dialog box, click **Close**.

Restoring Objects from the Recycle Bin

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

To restore items from the Recycle Bin:

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

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

Deleting Objects in the Recycle Bin

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

To delete objects in the Recycle Bin:

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

Emptying the Recycle Bin

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

To empty the Recycle Bin:

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

Grid Manager displays the *Confirm Empty Recycle Bin* dialog box to confirm that you wish to empty the Recycle Bin.

3. Click **Yes**.

Managing Third Party URL Links

In the URL Links section, you can add the URL links of frequently used third party portals and destination pages. For example, you can add the URL of a trouble ticket system and quickly access the portal once you are logged in to the Infoblox GUI. When you click an existing URL link, Grid Manager displays the destination page in a new browser window. You can also modify and delete existing URL links in this section.

On the appliance, only superusers have permissions to fully manage the URL links. Superusers can create URL links and make them globally available to all users. If you have limited-access permissions, you can only add URL links for your own use. You cannot share the links with other users.

You can do the following in the URL Links section:

- Add new URL links, as described in [Adding URL Links](#).
- Modify URL links, as described in [Modifying URL Links](#).
- Delete URL links, as described in [Deleting URL Links](#).

Adding URL Links

1. In the *Finder* panel, expand **URL Links**.
2. Click **Add**.
3. In the *URL Configuration* dialog box, complete the following:
 - **URL:** Enter the URL of the destination page you want to add. The appliance supports valid URL entries that contain up to 2000 characters. When you enter the URL, the appliance validates the entry. You cannot save the entry if the URL is not in a valid format.
 - **Name:** Enter a name that represents the portal or site of the URL.
 - **Set as global parameter:** This field appears only if you log in as a superuser. Select this check box to make the URL link globally available to all users.
 - **Logo:** Click **Upload** to add a logo to the URL. The appliance displays the logo in 16x16 pixels. Click **Reset to Default** to use the default logo.
4. Save the configuration.

Modifying URL Links

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

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

Deleting URL Links

To permanently delete an URL link:

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

Using Filters

You can control the amount and the kind of data displayed in a specific panel by adding filter criteria. When you add filter criteria, the appliance screens the data based on your filter rules and displays only the information that matches the rules. To narrow your search for specific information, you can add up to 10 filter rules. In some panels, such as the DHCP Networks tab, you can switch between viewing information with and without the filter criteria by toggling the filter on and off. You can save filter criteria as quick filters so you can reuse the same filter rules to obtain updated information without redefining them each time you log in to the appliance. For information about quick filters, see [Using Quick Filters](#) on page 68.

You can also use filters to find objects that have failed an operation. When you try to modify multiple objects with the same extensible attribute, the appliance may not modify all of the selected objects. For information, see [About Extensible Attributes](#) on page 322. For example, after you modify the extensible attribute “Building” with new value “West”, you can find the objects that are not updated by defining a filter with “Building” “does not equal” “West”.

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

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

To use a filter:

1. In a panel, click **Show Filter** to enable the function.
2. In the filter section, complete the following:
 - In the first drop-down list, select a field such as an object name, comment, or an extensible attribute (fields with a gray background) as the filter criterion. Grid Manager displays only the supported fields.
 - In the operator drop-down list, select an operator for the filter criterion. Depending on what you select in the first filter field, this list displays the relevant operators for the selection. The operator **Inheritance State equals** is displayed only when you select an inheritable extensible attribute from the **Type** drop-down list. This operator is not displayed if the extensible attribute is not inheritable.
 - In the value field, enter or select the attribute value for the first filter field. Depending on what you select for the first two filter fields, you can either enter a value or select an attribute from a drop-down list. For example, if you select an extensible attribute in the first filter field, you can enter the attribute value here. If you select an inheritable extensible attribute from the **Type** drop-down list, and select **Inheritance State equals** in the operator drop-down list, the value field displays a drop-down list with these values: **Inherited** and **Overridden/No Parent**. When you select **Inherited**, extensible attributes that are inherited by the descendants are listed. When you select **Overridden/No Parent**, extensible attributes which are overridden or do not have a parent are listed.
3. Optionally, click the **+** icon to add another filter rule. You can add up to 10 filter rules.
4. Click **Apply** to apply the rules
 - or
 - Click **Reset** to clear the filter criteria.

To view information with or without the filter criteria:

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

Using Quick Filters

A quick filter saves filter rules that you define in a specific panel. You can reuse a quick filter to find updated information in a panel without specifying the same rules each time. Superusers can define quick filters and share them with local users. Limited-access users can only create quick filters for their own use. You can create up to 10 global and 10 local quick filters in each panel that supports filters. For information about filters, see [Using Filters](#) on page 67.

The appliance supports the following quick filters:

- **System quick filters:** These are predefined filters. You cannot modify the criteria of these filters. System quick filters are prefixed with [S] in the quick filter list. Infoblox currently supports the following system quick filters in the DNS data panels:
 - **All Forward Mapping Zones:** This quick filter displays all forward mapping zones in lexicographical order.
 - **All Reverse Mapping Zones:** This quick filter displays all IPv4 and IPv6 reverse mapping zones in numerical order. The appliance displays IPv4 zones before IPv6 zones.
 - **All IPv4 Reverse Mapping Zones:** This quick filter displays only the IPv4 reverse mapping zones in numerical order.
 - **All IPv6 Reverse Mapping Zones:** This quick filter displays only the IPv6 reverse mapping zones in numerical order.
 - **RPZ Logs:** This quick filter displays only the RPZ syslog messages in CEF format. This option is displayed only in the **Syslog** when RPZ license is enabled.

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

- **Global quick filters:** Only superusers can define global quick filters. You can make global filters available to all users. Limited-access users can use global quick filters, but they cannot modify them. Global filters are prefixed with [G] in the filter list.
- **Local quick filters:** Limited-access users can create local quick filters for their own use. You cannot share local quick filters with other users in the Grid. Local filters are prefixed with [L] in the filter list.

Adding Quick Filters

1. In a panel that supports filters, click **Show Filters**.
 2. In the filter section, define filter criteria for the quick filter, as described in [Using Filters](#) on page 67.
 3. Click **Save**.
 4. In the *Save Quick Filter* dialog box, complete the following:
 - **Name:** Enter a name for the quick filter. The name must be 20 characters or longer. Ensure that you use a unique name for each quick filter in a particular filter category. For example, you can use the same filter name for both a global and local filter, but you cannot do so for two local filters.
 - **Set as a global quick filter:** This displays only if you log in as a superuser. Select this check box to make the quick filter globally available to all users.
 5. Save the configuration.
- The appliance adds the quick filter to the quick filter drop-down list in the specified panel.

Modifying Quick Filters

1. In a panel that supports filters, click **Show Filters**, and then select the quick filter you want to modify from the **Quick Filter** drop-down list.
2. In the filter section, click the Edit icon next to the filter name.
3. Modify the filter criteria, as described in [Using Filters](#) on page 67.
4. Click **Save**.

5. In the *Save Quick Filter* dialog box, you can click **Save** to save the modified filter criteria under the same quick filter name. You can also modify the quick filter name, as described in [Adding Quick Filters](#), and save the entry as a new quick filter.
6. Save the configuration.

Applying Quick Filters

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

Turning Off Quick Filters

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

- Select **None** from the quick filter drop-down list.
- Click **Toggle Filter Off** or **Reset** in the filter section.
- Delete a quick filter, as described in [Deleting Quick Filters](#).

Deleting Quick Filters

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

Using Global Search

You can use the global search function to search the entire NIOS database for data that matches a specific value and filter criteria. You can enter a search value and define filter criteria to refine the search. Grid Manager supports regular expressions in global search. Grid Manager can display up to 500 search results. When search results exceed 500, a warning message appears and you may want to refine your search. Search results remain in the *Search* dialog box until you reset the search parameters or log out of Grid Manager. You can search for DNS zones and resource records that contain IDNs. For information about IDNs, see [Support for Internationalized Domain Names](#) on page 93.

Note: Depending on the size of your database, global search may take a long time to complete. Grid Manager times out when queries or searches take longer than 120 seconds. To expedite searches, use filters to refine the search criteria.

To search globally:

1. Click the global search icon on the navigation bar.
2. In the *Search* dialog box, do the following:
 - In the first field, enter the value that you want your search results to match. For example, if you want to search for hostnames that contain “Infoblox”, enter Infoblox in this field. You can also specify the value of an inheritable extensible attribute. You can use regular expressions in the search value. For information, see [Regular Expressions](#) on page 1303.
 - In the **Type** drop-down list, select an object type, comment, or an extensible attribute (fields with a gray background) as the filter criterion. Grid Manager displays all the supported fields in the drop-down list. The default is **Type**. Grid Manager searches all objects when you use the default. You can narrow down the search and improve the search performance by selecting an object type. Extensible attributes are displayed with a gray background.

- In the operator drop-down list, select an operator for the filter criterion. Depending on what you select in the first filter field, this list displays the relevant operators for the selection. The operator **Inheritance State equals** is displayed only when you select an inheritable extensible attribute from the **Type** drop-down list. This operator is not displayed if the extensible attribute is not inheritable.
- In the value field, enter or select the attribute value for the first filter field. Depending on what you select for the first two filter fields, you can either enter a value or select an attribute from a drop-down list. For example, if you select an extensible attribute in the first filter field, you can enter the attribute value here. If you use the default **Type** in the first filter field, you can select an object or record type from the drop-down list. The default is **ALL**. Grid Manager searches all object types when you use the default. If you select an inheritable extensible attribute from the **Type** drop-down list, and select **Inheritance State equals** in the operator drop-down list, the value field displays a drop-down list with these values: **Inherited** and **Overridden/No Parent**. When you select **Inherited**, extensible attributes that are inherited by the descendants are listed. When you select **Overridden/No Parent**, extensible attributes which are overridden or do not have a parent are listed.

3. Optionally, click the **+** icon to add another filter. You can add up to 10 filter rules.

4. After you finish defining filters, click **Apply** or press **Enter**.

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

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

You can click **Reset** to clear the search results and start a new search. You can also click the Refresh icon to refresh the search results. Grid Manager stores the search results until you reset the search parameters or log out.

Editing Matching Objects in Search Results

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

To edit an object in the Results table:

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

Deleting Matching Objects in Search Results

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

To delete a matching object:

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

Grid Manager deletes the selected objects from the database. Most deleted objects are stored in the Recycle Bin. For information, see [Using the Recycle Bin](#) on page 64.

You can print search results. You can also export search results in CSV (comma separated value) format. For information, see [About CSV Import](#) on page 86 and [Exporting Displayed Data](#) on page 91.

Editing Multiple Extensible Attributes in Search Results

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

To edit multiple extensible attributes:

1. In the Results table, select the object check box. You can edit multiple extensible attribute values.
2. Click the Extensible Attributes icon.
3. In the *Multi-Select Edit Extensible Attributes* editor, click on the **Value** column to edit the value of the respective extensible attribute. For information about which values you can edit, see [Editing Multiple Extensible Attribute Values](#) on page 333.

Using the Go To Function

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

To use the Go to function:

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

Grid Manager completes the operation based on the selected object.

ABOUT TASKS

When you perform a task, such as adding a DNS zone or modifying a DHCP range, you can execute it immediately or schedule it for a future date and time, depending on your permissions. For information about how to schedule a task, see [Scheduling Tasks](#) on page 75.

Certain tasks, scheduled or not, may be subject to approvals if approval workflows are defined for specific admin groups. For information about how to define submitters and approvers for an approval workflow, see [Configuring Approval Workflows](#) on page 80.

Note that not all tasks can be scheduled or routed for approval. For a list of supported objects, see [Supported Objects for Scheduled and Approval Tasks](#) on page 73.

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

- The appliance cannot execute a scheduled or approval task that is associated with an extensible attribute, if you delete the extensible attribute after you have scheduled the task or submitted it for approval. For information about extensible attributes, see [About Extensible Attributes](#) on page 322.
- The appliance cannot execute, reschedule, or delete a task that is associated with a child object (such as a DHCP range) if you delete the parent object (such as a network) after you have scheduled the task or submitted it for approval.
- There are certain guidelines about scheduled and approval tasks when you upgrade the software, back up the database, and restore data. For information, see [Guidelines for Upgrading, Backing Up, and Restoring Data](#) on page 74.

Viewing Tasks

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

The appliance logs all tasks in the audit log and associates each with a task ID. By default, Grid Manager sorts tasks by Task ID in **Task Manager**. You can view tasks that you are allowed to see based on your permissions. For information about admin permissions, see [About Administrative Permissions](#) on page 160.

To view tasks:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab.
2. Grid Manager displays the following information for each task:
 - **Task ID:** The ID associated with the task. The appliance assigns an ID to a task in chronological order. By default, the appliance sorts tasks by **Task ID**.
 - **Affected Object:** The name or value of the object that is associated with the task. For example, if the task involves an A record, this field displays the domain name of the record. If it is a fixed address, it displays the IP address of the fixed address.
 - **Scheduled Time:** The date, time, and time zone when the appliance executes the task.
 - **Submitted Time:** The date, time, and time zone when the task was submitted.
 - **Submitter:** The username of the admin who scheduled or submitted the task.
 - **Ticket Number:** For an approval workflow, this number may be entered by the submitter to associate the task with a help desk ticket number or a reference number.
 - **Submitter Comment:** Comments entered by the submitter.
 - **Approval Status:** The current approval status. Possible values are **Approved**, **Not Applicable**, **Pending**, and **Rejected**.
 - **Execution Status:** The execution status of the task. Possible values are **Completed**, **Failed**, **Pending**, and **Executing**.

- **Executed Time:** The date, time, and time zone the task was executed.
- **Action:** The operation the appliance performs in this task. The can be: **Add, Modify, Delete, Network Discovery, Lock/Unlock Zone, or Restart Services.**
- **Task Details:** Detailed information about the task. This message also appears in the audit log.
- **Approver:** The username of the admin who has approved this task.
- **Approver Comment:** Comments entered by the approver.
- **Object Type:** The object type. For example, the appliance can display A Record or Fixed Address.

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

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

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

- Create a quick filter to save frequently used filter criteria. Grid Manager provides the following default quick filters that you can select from the Quick Filter drop-down list: **Pending Approvals, Rejected Tasks, and Scheduled Tasks**. For more information, see [Using Quick Filters](#) on page 68.
- Export and print the information in the table.
- Control the display of information in the panel by toggling between a single-line view and a multi-line view.
- Reschedule a task, cancel a scheduled task, or execute a task immediately.
- For approvers, select a task and click the Approve icon to approve the task, or click the Reject icon to disapprove the task. You can also reschedule the task while approving it.

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

Supported Objects for Scheduled and Approval Tasks

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

- IPv4 DHCP filters (MAC, option, NAC, relay agent, and DHCP Fingerprint)

Note: Only IPv4 MAC filters support approval workflows.

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

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

- Network discoveries
- Service restarts (for scheduled tasks only)

Note: Service restarts are not subject to approvals.

Guidelines for Upgrading, Backing Up, and Restoring Data

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

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

SCHEDULING TASKS

You can schedule tasks, such as adding DNS zones, modifying fixed addresses, and restarting services, for a future date and time. The scheduling feature is useful when you want to add, modify, or delete a record, or schedule a network discovery at a desired date and time. Using this feature, you can streamline your day-to-day operations. For example, you can schedule the deletion of records that you use for testing when the test time is up. You can also reassign an IP address to a fixed address when the location of the server to which the fixed address is assigned changes from one network to another.

You can schedule the addition, modification, and deletion of certain objects. For a list of the supported objects, see [Supported Objects for Scheduled and Approval Tasks](#) on page 73.

Depending on your permissions and the admin group to which you belong, your scheduled tasks may be subject to approvals by other admins in your organization. You may or may not receive email notifications about the status of your scheduled tasks depending on the configuration of the approval workflows. Approvers can reschedule your tasks after they have approved the tasks, if they have scheduling permissions. When you schedule and submit a task, you may need to enter a ticket number associated with the task or a comment about the task. For more information about approval workflows, see [Configuring Approval Workflows](#) on page 80.

Only superusers can view, reschedule, and delete all scheduled tasks. Limited-access admins can reschedule and delete only their scheduled tasks. If your scheduled tasks require approvals, the approvers who have scheduling permissions may reschedule your tasks to a different date and time after they have approved the tasks. Depending on your admin permissions, there are certain scheduled and approval tasks that you may or may not be able to perform. For more information, see [Supported Tasks for Different Admin Groups](#) on page 80.

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

- A superuser schedules a task, and another superuser reschedules or deletes the task.
- A limited-access admin schedules a task, and a superuser reschedules or deletes the task.
- A superuser or a limited-access admin schedules a task, and the task fails.
- An admin is configured to receive notifications based on the configuration of an approval workflow. For information about approval workflows, see [Configuring Approval Workflows](#) on page 80.

Superusers can also grant scheduling permissions to other admin groups. When the scheduling permission is added or inherited from an admin role, limited-access admin groups can schedule tasks. For information, see [Administrative Permissions for Network Discovery](#) on page 196.

Scheduling Additions and Modifications

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

To schedule an addition or a modification:

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

- **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Scheduling Appliance Operations

The appliance supports the scheduling of the following operations:

- Network discoveries—For information, see [Chapter 13, Network Discovery](#), on page 493.
- Service restarts—For information, see [Restarting Services](#) on page 386.

Scheduling Deletions

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

To schedule a deletion:

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

Scheduling Recursive Deletions of Network Containers and Zones

Superusers can determine which group of users are allowed to schedule the deletion of a network container and its child objects as well as a zone and its child objects. For information about how to configure the recursive deletion of network containers and zones, see [Configuring Recursive Deletions of Networks and Zones](#) on page 269.

To schedule the recursive deletion of network containers and zones:

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

- Select one of the following:
 - **Delete only the parent container:** Select this to delete only the parent objects and re-parent the child objects.
 - **Delete the parent container and its children:** Select this to delete the parent objects and all its child objects.

4. Click **Schedule Deletion**.

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

Viewing Scheduled Tasks

After you schedule a task, you can view the pending task in the **Administration** tab -> **Workflow** tab -> **Task Manager** tab. For more information, see [Viewing Tasks](#) on page 72. Superusers can view all scheduled tasks, and limited-access admins can view their own scheduled tasks.

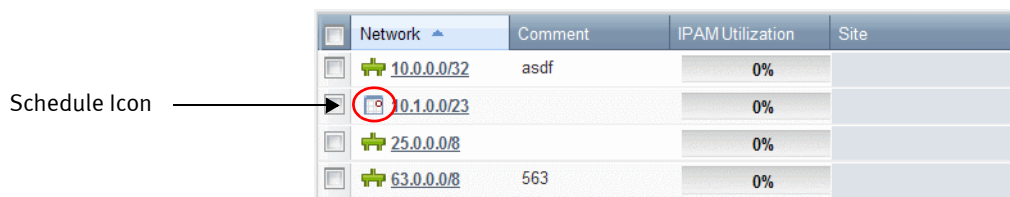
In certain panels such as the Network list panel and Smart Folders, Grid Manager displays a calendar icon next to objects that are associated with scheduled tasks, except for the addition of an object. You can click the icon to view the configuration and schedule. You can also reschedule the task if you are the owner of the task, a superuser, or an approver of the task (after you have approved it). In the corresponding editor, the Schedule icon is green when there is a pending scheduled task. For information, see [Icons for Scheduled Tasks](#) on page 77.

Icons for Scheduled Tasks

Grid Manager displays a scheduled task icon next to an object that is associated with a scheduled task (except for the addition of an object), as shown in [Figure 1.6](#). When you mouse over the icon, an informational dialog box appears displaying the type of action, the date and time of the scheduled task, and the person who scheduled the task.

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

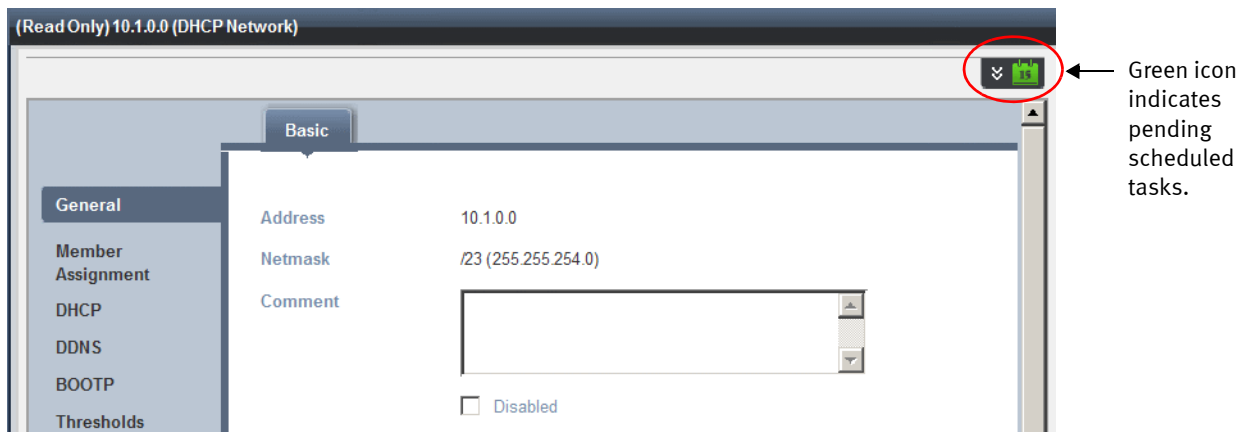
Figure 1.6 Icon for a Scheduled Task



Network	Comment	IPAM Utilization	Site
10.0.0.0/32	asdf	0%	
10.1.0.0/23		0%	
25.0.0.0/8		0%	
63.0.0.0/8	563	0%	

In the editor, Grid Manager displays the Schedule icon in green to indicate a pending scheduled task associated with the corresponding object, as shown in [Figure 1.7](#). You can click the Schedule icon to view the date and time of the scheduled task. You can also reschedule the task if you have the applicable permissions. For information, see [Rescheduling Tasks](#) on page 78.

Figure 1.7 Scheduling Icon Indicating a Pending Task



Pending Tasks for Operations

You can view all pending tasks for a network discovery or service restart in **Task Manager** if you have the applicable permissions. For information, see [Viewing Tasks](#) on page 72. You can also view the pending tasks in their corresponding dialogs.

To view the pending tasks in an editor:

1. **Network Discovery:** From the **Data Management** tab, select the **IPAM** tab, and then click **Discovery** from the Toolbar.
Service Restarts for the Grid: From the **Data Management** tab, select the **IPAM**, **DHCP** or **DNS** tab, and then click **Restart Services** from the Toolbar, or from the **Grid** tab, click **Restart Services** from the Toolbar.
Service Restarts for Grid members: From the **Data Management** tab, select the **DHCP** or **DNS** tab -> **Members** tab, select a member check box, and then click **Restart Services** from the Toolbar.
2. Click the Schedule icon at the top of the wizard, and then select **Click here to view/manage the scheduled items**. Note that this link appears only when you have one or more scheduled tasks.
3. Grid Manager displays the following information in the *Scheduled Tasks*:
 - **Scheduled Time:** The date, time, and time zone when the appliance executes the task.
 - **Submitted Time:** The date, time, and time zone when the task was submitted.
 - **Submitter:** The admin who scheduled the task.
 - **Task Details:** The message that appears in the audit log.

By default, the appliance sorts the tasks by **Scheduled Time** starting with the earliest scheduled start time.

You can do the following in this viewer:

- Sort the tasks in ascending or descending order by column, except for **Task Details**.
- Reschedule a selected task. For information, see [Rescheduling Tasks Associated with Operations](#) on page 79.
- Delete a selected task by selecting the task check box and clicking the Delete icon.
- Export and print the information in the table.

Rescheduling Tasks

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

To reschedule tasks associated with objects, see [Rescheduling Tasks Associated With Objects](#).

To reschedule tasks associated with operations, see [Rescheduling Tasks Associated with Operations](#).

Rescheduling Tasks Associated With Objects

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

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

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

To reschedule a task in an editor:

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

Rescheduling Tasks Associated with Operations

To reschedule a network discovery or a service restart:

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

Canceling Scheduled Tasks

To cancel a scheduled task:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab -> *scheduled_task* check box, and then click the Delete icon.
2. In the *Confirm Delete Request* dialog box, click **Yes**.
The appliance deletes the scheduled task and does not perform the scheduled operation. Therefore, no change is made to any record after you delete a scheduled task.

CONFIGURING APPROVAL WORKFLOWS

Approval workflows support routing certain core network service tasks submitted by an admin group to another for approval. You can add an admin group to an approval workflow and define the group as a submitter or approver group. Note that only superusers can create approval workflows. For information about how to set up admin groups, see [About Admin Groups](#) on page 154.

In an approval workflow, you can add a submitter group and an approver admin group that you have previously defined. You can also define when and to whom email notifications are sent, and configure options such as whether submitters or approvers must enter a comment or a ticket number when they submit tasks for approval. Approval workflows are useful when you want to control tasks that require reviews. For example, if you have a group of help desk users who can add, modify, and delete hosts and you want members of an operation group to review these tasks, you can define the help desk users as submitters, and then set up members of the operation group as approvers. You can then add the submitter and approver groups to an approval workflow and configure notifications options and other configurations, such as allowing the approvers to reschedule the submitted tasks.

Not all core network service tasks can be routed for approval. You can configure approval tasks associated with certain objects. For a list of supported objects, see [Supported Objects for Scheduled and Approval Tasks](#).

Note: When an admin group is defined as a submitter group, there are certain operations the submitters cannot perform even though they may have the permissions to do so. For information about such operations, see [Unsupported Operations for Submitters](#) on page 84.

To create an approval workflow, complete the following:

1. If you have not already done so, set up admin groups that you can configure as submitter groups and approver groups in an approval workflow, as described in [About Admin Groups](#) on page 154.
2. Create an approval workflow and configure email notifications and other options, as described in [Creating Approval Workflows](#) on page 81.

You can do the following after you have created approval workflows:

- View a list of approval workflows, as described in [Viewing Approval Workflows](#) on page 82.
- Modify approval workflows, as described in [Modifying Approval Workflows](#) on page 83.
- Delete approval workflows, as described in [Deleting Approval Workflows](#) on page 83.
- View a list of approval tasks, as described in [Viewing Approval Tasks](#) on page 83.
- View approval notifications, as described in [Viewing Approval Tasks](#) on page 83.

Supported Tasks for Different Admin Groups

Depending on your admin permissions, you may or may not be able to perform certain tasks that are subject to approvals. [Table 1.1](#) lists specific tasks and indicates which admin group can perform the tasks.

Table 1.1 Supported Tasks for Admin Groups

Tasks related to approval workflows	Admin groups that can perform the task		
	Submitters	Approvers	Superusers
Change the schedule of a task when it is pending approval	Yes	No	Yes
Change the schedule of a task after it has been approved	Yes (Task is re-submitted for approval)	Yes	Yes
Execute the task now when it is pending approval	No	No	No

Tasks related to approval workflows	Admin groups that can perform the task		
	Submitters	Approvers	Superusers
Execute the task after it has been approved	No	Yes	Yes
Delete a task when it is pending approval	Yes	No	Yes
Delete a task after it has been approved but pending execution	No	No	Yes
Delete a task after it failed or has been executed	No	No	Yes
Delete tasks by selecting the Select all objects in this dataset option	Yes	Yes	Yes

Note: Not all tasks are deleted, depending on the task status and the admin who performs the deletion.

Creating Approval Workflows

Before you create an approval workflow, ensure that you have admin groups that you can define as submitters and approvers. Note that a submitter group can be added to only one approval workflow, and approver groups can be added to multiple workflows. An approver can choose to approve a task and either keep or change the date and time when the task is executed. For information about scheduling and rescheduling tasks, see [Scheduling Tasks](#) on page 75. An approver can also reject a submitted task.

All submitted tasks are executed based on submitter permissions. When an admin submits a task, the appliance logs the task in the audit log and associates it with a task ID. You can view your tasks in **Task Manager**, as described in [Viewing Tasks](#) on page 72. Depending on your configuration, you can control when and to whom email notifications are sent. For example, you can configure the appliance to send notifications to only the approver each time when a task requires approval, or send notifications to both the submitter and approver group each time when a task is disapproved.

To create an approval workflow:

- From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab, and then click the Add icon.
- In the *Add Approval Workflow* wizard, complete the following:
 - Submitter Group:** From the drop-down list, select the admin group whose submitted tasks require approvals. Note that performing CSV imports do not require approvals. If there is a warning that the submitter group has CSV import permission, you may want to remove the permission.
 - Approver Group:** From the drop-down list, select the group that can approve tasks submitted by admins of the submitter group. If the approver group you select does not have the permission to schedule tasks, the approvers cannot reschedule the execution dates and times of the tasks when they approve them.
 - Ticket Number:** From the drop-down list, select one of the following to determine whether a ticket number is required when a submitter submits a task for approval.
 - Required:** The submitter must enter a ticket number when submitting a task.
 - Optional:** The submitter can choose to enter a ticket number or not when submitting a task.
 - Not Used:** The **Ticket Number** field does not appear when the submitter creates a task.
 - Submitter Comment:** From the drop-down list, select whether the submitter must enter a comment or not when submitting a task for approval. You can select **Required**, **Optional**, or **Not Used**.
 - Approver Comment:** From the drop-down list, select whether the approver must enter a comment or not when approving a task. You can select **Required**, **Optional**, or **Not Used**.

3. Click **Next** and complete the following to specify notification options for the workflow:
 - **Approver Notification Address(es)**: Select one of the following to specify to which approver email addresses the appliance sends workflow notifications. The default is **Group Email Address(es)**.
 - **Group Email Address(es)**: Select this if you want the appliance to send notifications to the list of email addresses configured for the admin group. For information about how to configure this list, see [About Admin Groups](#) on page 154.
 - **User Email Address(es)**: Select this if you want the appliance to send notifications to individual email addresses of the admin group.
 - **Notifications sent on**: Select the operations that can trigger email notifications. When you select an operation, the appliance sends a notification each time that operation occurs. By default, all operations are selected.
 - **Approval Required**: The appliance sends an email notification each time an approval is required.
 - **Task Approved**: The appliance sends an email notification each time a task is approved.
 - **Task Rejected**: The appliance sends an email notification each time a task is rejected.
 - **Task Succeeded**: The appliance sends an email notification each time a task is completed successfully.
 - **Task Failed**: The appliance sends an email notification each time the execution of a task fails.
 - **Task Rescheduled**: The appliance sends an email notification each time a task is being rescheduled.
 - **Notifications sent to**: For each operation, select whether the **Approver**, **Submitter**, or **Both** are notified when the operation occurs. The default value is **Both** for all operations. For information about email notifications, see [Viewing Approval Tasks](#) on page 83.
4. Optionally, click **Next** to add extensible attributes to the approval workflow. For information, see [About Extensible Attributes](#) on page 322.
5. Save the configuration.

Viewing Approval Workflows

Grid Manager lists all approval workflows in the **Approval Workflows** tab. Only superusers can view approval workflows defined for the Grid. Limited-access users cannot view approval workflows.

To view approval workflows:

1. From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab.
2. Grid Manager displays the following for each approval workflow:
 - **Submitter Group**: The name of the admin group whose tasks require approvals.
 - **Approver Group**: The name of the admin group that can approve tasks submitted by members of the submitter group.
 - **Ticket Number**: Displays whether the submitter is required to enter a ticket number when submitting tasks that require approvals. Possible values are **Not Used**, **Optional**, and **Required**.
 - **Submitter Comment**: Displays whether the submitter is required to enter a comment when submitting tasks that require approvals. Possible values are **Not Used**, **Optional**, and **Required**.
 - **Approver Comment**: Displays whether the approver is required to enter a comment when approving tasks. Possible values are **Not Used**, **Optional**, and **Required**.
 - **Site**: Values that were entered for this predefined extensible attribute.

You can do the following in this tab:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read-only.
- Sort the data in ascending or descending order by column.
- Select an approval workflow and click the Edit icon to modify data, or click the Delete icon to delete it.
- Use filters and the Go to function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the Go to field and select the object from the possible matches.

- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print and export the data in this tab.

Modifying Approval Workflows

You can modify information in an approval workflow, except for the submitter group.

To modify approval workflow configuration:

1. From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab.
2. Select an approval workflow and click the Edit icon.
3. Grid Manager provides the following tabs from which you can modify information:
 - **General** tab: You can modify the approver group and decide whether the ticket number, submitter comment, and approver comment are required, but you cannot change the submitter group. For information, see [Creating Approval Workflows](#) on page 81.
 - **Approval Notifications** tab: You can modify when and to whom email notifications are sent. For information, see [Creating Approval Workflows](#) on page 81.
 - **Extensible Attributes** tab: You can add or modify values of extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
4. Save the configuration.

Deleting Approval Workflows

You can delete an approval workflow any time after you have created it. Note that when you delete a workflow that has associated tasks that are pending approvals, the tasks will be rejected after you delete the workflow.

To delete an approval workflow:

1. From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab.
2. Select an approval workflow and click the Delete icon.
3. Click **Yes** in the *Delete Confirmation* dialog.

Viewing Approval Tasks

If you belong to an approver admin group, you can view, approve, or reject tasks that are pending your approval in the **Task Manager** tab. For information, see [Viewing Tasks](#) on page 72. Submitters can view all pending and completed tasks they have submitted.

Viewing Workflow Notifications

When a submitter and approver receives an email notification about their tasks, the appliance lists the approval status and workflow related information such as task ID, submitter name, execution time, object type and action in the email notification.

Following is a sample email notification:

```
Notification:
=====
Message: Task 32 submitted by subm has been approved
The following task has been approved:
Task details

Task ID: 32
Submitter: subm
```


Approver: jdoe
 Submit time: 2012-10-09 05:55:01 (UTC) Coordinated Universal Time
 Execution time: N/A
 Object type: NS Record
 Action: Add
 Affected object: corp1.com
 Ticket number: MKTG245
 Submitter comment: Create an NS record.
 Approver comment: Approved.

Click here to go to the task management tab -
<https://192.168.1.2/ui/?contextId=taskmanager>

Note: When you can click the hyperlink displayed in the notification, you can log in to **Grid Manager** and access the **Task Manager** tab in a separate browser tab or window.

Unsupported Operations for Submitters

When admins are part of a submitter group in an approval workflow, there are certain operations they cannot perform even though they may have the permissions to do so. Following is the list of operations that submitters cannot perform:

- Reclaim IPv4 or IPv6 addresses
- Expand networks
- Resize networks
- Split networks
- Sign (DNSSEC) zones
- Unsign DNSSEC signed zones
- Import DS to DNSSEC signed zones
- Perform KSK rollovers on a DNSSEC signed zones
- Copy records from one DNS zone to another
- Clear all discovered data
- Clear discovered timestamps
- Clear unmanaged addresses
- Resolve discovery conflicts
- Update extensible attributes on multiple objects at a the same time
- Delete or modify several objects at a time (using the “Select all objects in this dataset” option from Grid Manager)
- Order DHCP Ranges inside a network (feature is available only when used with Sophos)
- Configure member DHCP Captive Portal through the wizard
- Restore objects from the Recycle Bin
- Delete non-native NIOS DNS resource records. These objects can only be synchronized from a Microsoft DNS server.
- Copy rules from one Response Policy Zone to another
- Order Response Policy Zones

ABOUT LONG RUNNING TASKS

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

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

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

The appliance supports the following long running tasks:

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

Running Tasks in the Background

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

You can run the following tasks in the background:

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

To run a task in the background:

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

You can view the progress of the task by clicking the progress bar at the top of the interface. For information, see [Monitoring Long Running Tasks](#) on page 86.

Monitoring Long Running Tasks

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

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

ABOUT CSV IMPORT

Use **CSV Import** to import DNS, DHCP, and IPAM data through Grid Manager. You can use this feature to migrate or add new data, overwrite existing data, and merge new data with existing data.

To import new data, you must first prepare a data file (include all required fields and follow the proper syntax), and then start an import through Grid Manager. You can also export existing data to a data file, modify the data, and then import the modified data to the database. You can either overwrite existing data with the modified data or merge new data with the existing data.

The appliance supports CSV import for most record types. You can use IDNs and punycode for the domain name field for most of the DNS object types. For information about IDNs and punycode, see [Support for Internationalized Domain Names](#) on page 93. For each supported record type, you must include all required fields in the header row of the dataset that you want to import. For a list of supported record types and specific guidelines for creating a data file, refer to the *Infoblox CSV Import Reference*.

To import a data file:

1. Create a data file if you do not already have one. Follow the guidelines for the supported objects to ensure that you include all the required fields in the file. For more information, refer to the *Infoblox CSV Import Reference*. You can also export existing data and update the file for re-import. For information, see [Exporting Data to Files](#) on page 89.
2. Configure import options. For information, see [Configuring Import Options](#) on page 89.
3. Start a CSV import. For information, see [Managing CSV Imports](#) on page 90.

When you submit multiple CSV imports, the appliance puts the import jobs in queue and executes them one at a time in the order they are submitted. When a job is being executed, it is in **RUNNING** state. When a job is in queue for execution, it is in **PENDING** state. Each user can have only one import job at a time. When you upload a data file, you cannot upload another one until you start the import of the first uploaded file. You can view the status of each import job through **Import Job Manager**. Superusers can view all import jobs while limited-access users can view only the jobs they submitted.

To access **Import Job Manager**, from the **Data Management** tab, click **CSV Import** from the Toolbar, or from the **Tasks Dashboard**, click **CSV Import** in the IPAM Task Pack. Superusers and limited-access users that have applicable configurations and permissions can perform CSV imports and exports. For information about user permissions for CSV imports and exports, see [CSV Import User Permissions](#).

You can do the following in **Import Job Manager**:

- View the list of CSV import jobs, as described in [Viewing CSV Import Jobs](#) on page 88.
- Add and start CSV import jobs, upload data files, or stop CSV imports, as described in [Managing CSV Imports](#) on page 90.

- Select a job and click the Status icon to view and modify its current status or settings, as described in [Managing CSV Imports](#) on page 90.
- Select a PENDING or UPLOADED job and click the Cancel icon to cancel the job. You can only cancel a pending job or an uploaded file.
- Click the Refresh icon to refresh the Job Viewer.

Note that superusers can view any jobs in the Job Viewer, and limited-access users can only view jobs they submitted.

Note: The list of CSV import jobs are not restored when you restore a backup file or when you promote a master candidate.

CSV Import User Permissions

Superusers can perform any CSV import tasks. You must assign limited-access users the correct configurations and permissions so they can perform CSV imports and exports. For information about how to configure the CSV Import task for limited-access users, see [About Dashboard Templates](#) on page 114. Limited-access users can import data to which they have proper permissions. For information about admin permissions, see [About Administrative Permissions](#) on page 160.

Changes you make to user permissions can affect CSV import and export behaviors. The following table lists actions performed on user permissions and the corresponding effects on CSV imports and exports:

Table 1.2

Actions taken on user permissions	CSV import and export behaviors associated with the affected user account
Delete a user account	<ul style="list-style-type: none"> • CSV import jobs remain in the system and are accessible by superusers only. • All PENDING import jobs cannot be executed due to authentication failures. • If the action is taken during a RUNNING import job, the rest of the import will fail. • All STOPPED and COMPLETED jobs are available to superusers only.
Modify a user account	<ul style="list-style-type: none"> • If the action is taken during a RUNNING import job, the rest of the import will fail.
Remove user permissions in a user account	PENDING jobs may be completed with errors.

CSV Import Limitations

Ensure that you understand the following limitations before you start an import:

- You can import only one CSV file at a time.
- Do not use UTF-8 characters in the CSV file name.
- When you perform a CSV import that includes objects that have scheduled changes or updates associated with them, the import fails. Only superusers can cancel the scheduled changes.
- When you stop an import, the appliance completes the import of the data row it is currently processing before it stops the import. You cannot resume the import from where it stopped.
- You cannot roll back to previous data.
- The following data cannot be imported: Microsoft management, DNSSEC, and GSS-TSIG data.
- CSV import does not support DNSSEC zones, though resource records added for a signed zone are supported.

- Only editable data can be imported. Discovered data cannot be imported or manipulated.
- When you promote a new Grid Master during an import, the import stops; and it does not restart on the new Grid Master. When a failover occurs during an import, the import stops on the old active node, and it does not restart on the new active node.
- It may take longer than expected to import a large number of DHCP ranges that are associated with a single MAC address filter.
- When a CSV import starts, the appliance validates the first 100,000 rows of data in the CSV file. If the file contains more than 100,000 rows of data, the appliance validates the rest of the data as the import progresses.
- The appliance supports up to one million rows of data in each CSV import.
- You cannot import network containers.
- To successfully import RIR (Regional Internet Registries) organizations, you must also specify the maintainer password. Note that the password field is not exported during a CSV export. For information about RIR updates, see [RIR Registration Updates](#) on page 437.

Viewing CSV Import Jobs

1. From the **Data Management** tab, click **CSV Import** from the Toolbar.
or
From the **Tasks Dashboard**, click **CSV Import** in the IPAM Task Pack.
2. In **Import Job Manager**, Grid Manager displays the following information about the import jobs that were submitted in the past 14 days:
 - **User Name:** The admin user who submitted the CSV import. Only superusers can view this column.
 - **Status:** The current status of the import job. The status can be one of the following:
 - **RUNNING:** The job is being executed.
 - **PENDING:** The job is in queue for execution.
 - **COMPLETED:** The import is completed. Check the **Message** field for information about the import.
 - **UPLOADED:** The data file has been uploaded, but import is not started.
 - **STOPPED:** The job has been stopped. You can select the job and restart the import.

Note: After a product restart, which can be caused by a failover, all **RUNNING** jobs go into **STOPPED** state; all **PENDING** jobs continue to be queued for execution.

- **Submitted:** The timestamp when the job was submitted.
 - **Completed:** The timestamp when the job was completed. This field is blank if the job has not been completed yet.
 - **File Name:** The CSV data file name.
 - **Message:** This field displays the number of rows of data that have been processed and the number of rows of data the import has detected errors. Depending on the import options, Grid Manager displays the row number at which it stops the import when it encounters an error, or the total number of rows it has processed by skipping over the erroneous data. For example, if there are 100 rows of data and you select “On error: Stop importing,” and there is an error in row 90, the appliance displays **90 of 100 completed, 1 error**. If you select “On error: Skip to the next row and continue,” the appliance displays **100 of 100 completed, 1 error**.
-

Note: Superusers can view all CSV import jobs and limited-access users can view only the jobs they submitted.

Creating a Data File for Import

If you are migrating new data into the database, you must prepare the data file using the correct format and syntax before you can import it successfully. You must include all the required fields and understand the dependencies among some of the fields. For detailed information about the guidelines, supported record types, and interdependencies among fields, refer to the *Infoblox CSV Import Reference*.

Exporting Data to Files

You can export existing data to a CSV file. The appliance marks all required fields with an asterisk (*) in the exported file. It also adds a `_new_XXXX` field to each required field so you can use this field to update data. You cannot stop an export once you start it.

Note: Limited-access users can export up to 2,000 rows (2,000 objects) of data. For performance reasons, NIOS has limited the objects to 2,000. If data exceeds 2,000 rows, the CSV file contains the first 2,000 rows of data and a message at the end that indicates the report is not complete. Only superusers can export data that exceeds 2,000 rows. For example, consider that you are exporting **Infoblox::DNS::Host** objects. The NIOS appliance displays an output file with 4,000 lines, because the **Host** record consists of two lines, one for the **HostRecord** and another for the **HostAddress**; however, the total count of objects will still remain 2,000.

To export all data to a CSV file:

1. From Grid Manager, navigate to the panel that contains the data you want to export. For example, if you want to export data for all DNS zones, select the **Data Management** tab -> **DNS** tab -> **Zones** tab.
2. In the panel, select **Export data in Infoblox CSV Import format** from the **Export** drop-down menu.
3. In the *Export* dialog box, complete the following:
 - **Separator:** Select the separator used in the data file. The default is Comma.
 - Click **Export**.

The appliance exports all the fields of the records that are displayed in this panel based on your filter criteria. You can either open the data file or save it to your computer. The appliance uses a default file name depending on the panel from which you perform the export. For example, when you export the data from the **IPAM** tab, the default file name is *Allnetworks.csv*. When you export data from the **DNS** tab, the default file name is *Allzones.csv*. The file contains a header row that includes all the fields of the corresponding record type. You can update this data file, and then re-import the data in to the database.

You can also export the displayed fields in a panel. For information, see [Exporting Displayed Data](#) on page 91.

Configuring Import Options

Before you import a data file, configure the import options, as follows:

1. From the **Data Management** tab, click **CSV Import** from the Toolbar.
or
From the **Tasks Dashboard**, click **CSV Import** in the IPAM Task Pack.
2. From **Import Job Manager**, click the Add icon.
3. In the *Import Manager* editor, select the **Options** tab, and then complete the following:
 - **Add - Create a new object from each row in the uploaded file:** Select this to add new information to the database.
 - **Edit - Modify existing object:** Select one of the following to update the existing data:
 - **Overwrite - Values from the uploaded file replace existing values:** When you select this option, the appliance overwrites the existing data with the data from the uploaded file. If you want to overwrite values in the required fields, you must include the required fields and the corresponding `_NEW_XXXX` fields in the data file.

- **Merge - Update missing values with values from the uploaded file:** When you select this option, the appliance adds only the data that is not currently in the database. It does not overwrite the existing data, even if the data file contains new values for certain fields. If you want to overwrite values in the required fields, you must include the required fields and the corresponding `_NEW_XXXX` fields in the data file.
- **On error:** Select one of the following to tell the appliance what to do when it encounters an error during an import:
 - **Stop importing:** The appliance stops the data import once it encounters an error in the uploaded file.
 - **Skip to the next row and continue:** The appliance skips over errors and continues the data import. You can download an error report to identify the erroneous data. For information, see [Managing CSV Imports](#) on page 90.

Managing CSV Imports

After you configure the import options, you can select a data file and start an import or upload a data file.

1. From the **Data Management** tab, click **CSV Import** from the Toolbar.
or
From the **Tasks Dashboard**, click **CSV Import** in the IPAM Task Pack.
2. Click **CSV Import** from the Toolbar.
3. From **Import Job Manager**, click the Add icon. Note that the Add icon is disabled if you have one RUNNING or PENDING job and an uploaded file in the Job Viewer.
4. In the *Import Manager* editor, select the **File Selection** tab, and then complete the following:
 - **File:** Click **Select**. In the *Upload* dialog box, click **Select** to navigate to the file you want to import, and then click **Upload**.
 - **Separator:** From the drop-down list, select the separator you use in the data file. The default is **Comma**.

In the File Preview table, Grid Manager displays the header row, the first six rows, and up to 15 columns of the imported data. Field names with asterisks (*) indicate required fields. Note that you must define these fields in the imported file. If any of the required fields are missing, the appliance generates an error during the import. Grid Manager also displays the following information:

 - **Current Status:** If an import is in progress, this field displays its current status. Otherwise, it displays the date and time of the last import.
 - **Last Action:** Displays the last operation and the admin who initiated it.
 - **Rows Completed:** The number of rows of data the import has processed. Depending on the import options, Grid Manager displays either the row number at which it stops an import when it encounters an error or the total number of rows it has processed by skipping over the erroneous data. For example, if there are 100 rows of data and you select “On error: Stop importing,” and there is an error in row 90, Grid Manager displays **90 of 100** here. If you select “On error: Skip to the next row and continue,” Grid Manager displays **100 of 100** here and displays **1** in **Rows with Errors**.
 - **Rows with Errors:** The number of rows of data the import has detected errors. Click **Download Errors** to download the CSV file that contains the fields and the rows of erroneous data. You can use this report as a reference to update the data file before you import the file again.
5. Click **Start** to start the CSV import or click **Save & Close** to upload the selected file.

Note: You can also start and stop an import, and review the import status from the *CSV Import Status* widget on the Dashboard. For information, see [Import Job Manager](#) on page 132.

EXPORTING DISPLAYED DATA

You can export visible information, such as global search results and the syslog file, in CSV format from panels and pages that support the Export function, and then easily convert the file to PDF and other file formats. You can also export all data in a specific panel. For information, see [Exporting Data to Files](#) on page 89.

To export displayed data:

1. From Grid Manager, navigate to the panel that contains the data you want to export. For example, if you want to export data for DNS zones, select the **Data Management** tab -> **DNS** tab -> **Zones** tab.
2. In the panel, select **Export visible data** from the **Export** drop-down menu.
3. In the *Export* dialog box, click **Start**. Grid Manager displays a message about the time required to export data could be long depending on the amount of data.
4. Click **Download** when the export is finished.
5. Depending on your browser and operating system, you may need to do one of the following in the *Opening .csv* dialog box:
 - **Open with:** Select a program with which you want to open the .csv file.
 - **Save to Disk:** Select this if you want to save the .csv file to your local computer.
 - **Do this automatically for files like this from now on:** Select this check box if you want Grid Manager to use the same method for future exports. When you select this check box, Grid Manager does not display the *Opening .csv* dialog box in the future.
6. Click **OK**.
Depending on the selected option, Grid Manager opens the file using the program you select, or saves the file to your local computer.

PRINTING FROM GRID MANAGER

In Grid Manager, you can print information from panels and pages that support the Print function. Grid Manager prints data one page at a time. The amount of data that is displayed in a specific panel depends on the table size configuration that you set in your user profile. For information, see [Specifying the Table Size](#) on page 50.

To print:

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

MULTILINGUAL SUPPORT

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

- Hostnames for Microsoft Windows clients that support Microsoft Windows code pages. For information, see [Configuring UTF-8 Encoding for Hostnames](#) on page 812.
- Input fields through Grid Manager. For information, see [UTF-8 Supported Fields](#).

UTF-8 is a variable-length character encoding standard for Unicode characters. Unicode is a code table that lists the numerous scripts used by all possible characters in all possible languages. It also has a large number of technical symbols and special characters used in publishing. UTF-8 encodes each Unicode character as a variable number of one to four octets (8-bit bytes), where the number of octets depends on the integer value assigned to the Unicode character. For information about UTF-8 encoding, refer to [RFC 3629](#) (*UTF-8, a transformation format of ISO 10646*) and the *ISO/IEC 10646-1:2000 Annex D*. For information about Unicode, refer to *The Unicode Standard*.

Depending on the OS (operating system) your management system uses, you must install the appropriate language files in order to enter information in a specific language. For information about how to install language files, refer to the documentation that comes with your management system.

UTF-8 Supported Fields

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

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

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

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

UTF-8 Support Limitations

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

- Object names that have data restrictions due to their usage outside of the Infoblox database do not support UTF-8 encoding. For example, IP addresses and Active Directory domain names.
- When importing a database, most of the ASCII control characters cannot be encoded. This might cause failures in upgrades or database restore operations.
- Search is based on the Unicode standard. Depending on the language, you might not be able to perform a case-sensitive search.
- Binary data is encoded as text.
- UTF-8 encoding does not fully support regular expressions. It matches constant strings. However, it does not encode characters that are inside square brackets or followed by regular expressions such as *, ?, or +.
- You can use UTF-8 characters to authenticate both the User Name and Password through the Infoblox GUI, but not through the Infoblox CLI.
- You cannot use UTF-8 characters to name a custom DHCP fingerprint. For information about DHCP fingerprint detection, see [Infoblox DHCP Fingerprint Detection](#) on page 1032.

SUPPORT FOR INTERNATIONALIZED DOMAIN NAMES

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

An IDN is a domain name that contains a language-specific script or alphabet, such as Arabic, Chinese, Russian, Devanagari, or the Latin alphabet-based characters with diacritics, such as French. IDNs are encoded in multi-byte Unicode and are decoded into ASCII strings using a standardized mechanism known as Punycode transcription. For example, DNS Zone 'инфоблокс.рф' (IDN in Russian) can be written as 'xn--90anhdigczv.xn--p1ai' in the punycode representation. In addition, the appliance has a built-in conversion tool to assist you in identifying and troubleshooting an IDN or the punycode representation of an IDN. For information about how to decode IDNs, see [Decoding IDNs and Encoding Punycode](#) on page 93.

The appliance supports IDNs in certain fields. For more information, see [IDN Supported Fields](#) on page 93. There are certain guidelines and limitations about IDN support. For more information, see [IDN Support Limitations](#) on page 94.

Decoding IDNs and Encoding Punycode

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

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

1. Select any tab in Grid Manager, and then click **IDN Converter** from the Toolbar.
2. In the *IDN Converter* wizard, complete the following:
 - Specify the domain name in the **Unicode** text box and click **Convert to Punycode**. The **Punycode** field displays the punycode representation of the domain name.
 - Specify the punycode representation of a domain name in the **Punycode** field and click **Convert to Unicode**. The **Unicode** field displays the domain name in its original character set.

Note: You can use special characters in the **Unicode** and **Punycode** fields.

- Click **Clear** to clear the entries. Note that when you click **Clear** for a specific conversion, the appliance clears only the error message that corresponds to that conversion.
3. Click **Close**.

IDN Supported Fields

The NIOS appliance supports IDNs in all domain name fields. For information, see [IDN Support For DNS Zones](#) on page 622. You can enter non-English characters in the domain name fields through Grid Manager and the Infoblox API. The NIOS appliance does not support IDNs for data that is configurable through the Infoblox CLI commands. You can use the punycode representation to configure data through the CLI commands.

The appliance supports IDNs in the following:

- You can use UTF-8 characters when defining your own hostname checking policy. For information, see [Specifying Hostname Policies](#) on page 592.
- You can use both IDNs and punycode to search for IDN data through Global Search. For information, see [Global Search](#) on page 59.
- Use smart folders to organize and monitor IDN data. However, if the content in a smart folder contains IDNs, then the punycode representation is not available. For information, see [About Smart Folders](#) on page 140.
- You can import data that contains IDNs in CSV format for the supported fields and objects using CSV import. For more information, see [About CSV Import](#) on page 86. For a list of supported record types and specific guidelines for creating a data file, refer to the *Infoblox CSV Import Reference*.
- The IPAM tab displays IDNs for DNS resource records associated with IP addresses, such as A records, AAAA records, hosts, and PTR records. For information, see [About IP Address Management](#) on page 458.

- The *DNS Zones Last Queried* report and *DNS Resource Records Last Queried* report support IDNs. For information, see [Predefined Report Categories](#) on page 1145.
- The audit log entries are displayed in their original characters. The audit log contains IDN data as received by the appliance and as specified by the administrators. Note that the punycode representation generated by NIOS is not displayed in the audit log.
- When you upgrade from a previous NIOS release, the appliance converts all punycode to IDNs. If the conversion fails, the appliance retains the punycode representation to avoid upgrade failure. For information about upgrades, see [Guidelines for Upgrading, Backing Up, and Restoring Data](#) on page 74.
- When you restore a backup file from a previous NIOS release, the appliance converts all punycode to IDNs. If the conversion fails, the appliance retains the punycode representation to avoid failure to restore the database. For information, see [Guidelines for Upgrading, Backing Up, and Restoring Data](#) on page 74.
- If synchronized data between the appliance and Microsoft servers contains IDNs, the IDNs are preserved. For information, see [Managing Microsoft DNS Servers](#) on page 968.

IDN Support Limitations

The appliance has the following IDN support limitations:

- F5® load balancers does not support IDNs. The NIOS appliance does not encode punycode to IDNs for F5 load balancer related objects. Only the punycode representation is available.
- Multi-Grid configuration does not support IDNs.
- The Infoblox CLI does not support IDNs.
- If a resource record containing an IDN is added to the Infoblox Grid through DDNS updates, the domain name field displays the record name in UTF-8 encoded format. For more information, see [Managing Resource Records](#) on page 660.
- The following FQDNs does not support IDNs:
 - FQDN of an external DNS Server (direct or via name server group)
 - FQDN of a DNS root server
 - FQDN of a Microsoft server
 - FQDN of an Infoblox Grid Member
 - FQDN of an external authentication source (Active Directory, LDAP, OCSP, RADIUS, TACACS+)
 - FQDN of an NTP server
 - FQDN of a HSM SafeNet Module
 - FQDN of an email relay server
 - FQDN of a vSphere/ESX server
 - FQDN of a Kerberos Key Distribution Center

Using IDNs for Unsupported Objects

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

Use the punycode representation of IDNs for the following:

- When you configure domain names in forwarder servers, NXDOMAIN rulesets, blacklist rules, and DNS resolver search lists.
- When you configure domain names for DHCP and DHCPv6 services, including DDNS domain name, any DHCP options that accept domain names (host-name (12) string) or lists of domain names (domain-search (119) domain-list), and DHCPv6 options that accept domain names (dhcp6.fqdn (39) string) or lists of domain names (dhcp6.domain-search (24)) domain-list.

- When you add domains in the Inclusion list and Exclusion list. For information, see [Excluding Domains From Query and Response Capture](#) on page 1127.
- When you configure rules for a local RPZ and RPZ feed. For information, see [Configuring Local RPZs](#) on page 1238 and [Configuring RPZ Feeds](#) on page 1248.

Displaying IDN Entries in Punycode

The appliance displays IDN entries in punycode for the following:

- The data of a zone for which an Infoblox Grid member is the secondary server.
- The CLI commands `dig`, `ddns_add`, `ddns_delete`, `show dns`, and `set dns` support punycode only. For information about CLI commands, refer to the *Infoblox CLI Guide*.
- All syslog entries generated by DNS.
- IDN data in database files is stored in punycode.
- The DNS cache of a Grid member that contains IDNs.
- The **Reporting** tab displays all report data that contains IDNs in punycode. For information, see [Predefined Report Categories](#) on page 1145.



Chapter 2 Dashboards

Dashboards provide easy access to tasks and a quick overview to the status of your Grid and DNS, DHCP and IPAM services. This chapter contains the following sections:

- [About Dashboards](#) on page 99
- [The Tasks Dashboard](#) on page 99
 - [About Task Packs](#) on page 99
 - [The IPAM Task Pack](#) on page 100
 - [The Network Automation Task Pack](#) on page 109
- [About Dashboard Templates](#) on page 114
 - [Adding Dashboard Templates](#) on page 115
 - [Modifying Dashboard Templates](#) on page 115
 - [Deleting Dashboard Templates](#) on page 116
 - [Assigning Dashboard Templates](#) on page 116
- [Status Dashboards](#) on page 116
 - [Adding Widgets to Dashboards](#) on page 117
 - [Grid Status](#) on page 121
 - [Grid Upgrade Status](#) on page 122
 - [Member Status \(System Status\)](#) on page 123
 - [DNS Statistics](#) on page 124
 - [Ranges Over Threshold](#) on page 125
 - [IPv4 Failover Associations Status](#) on page 125
 - [DHCP Statistics](#) on page 126
 - [Network Statistics](#) on page 127
 - [IPv4 Networks Over Threshold](#) on page 128
 - [Discovery Status](#) on page 128
 - [My Commands](#) on page 129
 - [DDNS Statistics](#) on page 130
 - [System Activity Monitor](#) on page 130
 - [File Distribution Statistics](#) on page 131
 - [Active WebUI Users](#) on page 131
 - [Microsoft Servers Status Widget](#) on page 131
 - [Import Job Manager](#) on page 132
 - [Load Balancer Status](#) on page 133

- [*Pending Approvals*](#) on page 133
- [*Response Policy Zone \(RPZ\) Statistics*](#) on page 134
- [*Infoblox Community*](#) on page 136
- [*Mobile Devices Status*](#) on page 136
- [*Threat Protection Statistics*](#) on page 138

ABOUT DASHBOARDS

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

- **Tasks:** The Tasks dashboard contains task packs that provide easy access to commonly performed tasks. A task pack is a collection of tasks that belong to a specific service or function, such as IPAM or Automation. For information, see [The Tasks Dashboard](#) on page 99.
- **Status:** A status dashboard contains widgets from which you can view and manage DNS, DHCP, and IPAM status and data. You can configure multiple status dashboards for managing a large number of Grid members. For information, see [Status Dashboards](#) on page 116.

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

To change your home page:

1. Navigate to any tab in Grid Manager (except for the **Dashboards** tab).
 2. Click **User Profile** from the Toolbar and complete the following in the *User Profile* dialog box:
 - **Default Dashboard:** Select **Status** or **Task** from the drop-down list.
 3. Save the configuration.
- Grid Manager displays the selected dashboard as your home page when you log in the next time.

THE TASKS DASHBOARD

The Tasks Dashboard provides easy access to commonly performed tasks, such as adding networks and adding host records. Tasks are grouped by service-specific task packs. You must install valid licenses on the appliance to see and perform specific tasks on the Tasks Dashboard. For information about the required licenses for IPAM tasks, see [Table 2.1](#) on page 100.

You must also have at least read-only permission to a task-related object to add or hide the task in its task pack. To execute a task, you must have the appropriate permissions to the member and objects that are related to the tasks. For example, to add a host record from the Tasks Dashboard, you must have at least read-only permission to the host records task and read/write permission to the zone and network in which the host records are created. For information about permissions, see [Administrative Permissions for Dashboard Tasks](#) on page 214.

About Task Packs

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

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

To use the Automation Task Pack, you must enable the Network Automation Tasks feature set and establish a working connection between the NIOS appliance and an Infoblox Network Automation appliance. See [Enabling the Network Automation Tasks](#) on page 107 for details.

Each task in a task pack opens a workflow dialog in which you can create task-related objects without navigating through other tabs and editors in Grid Manager. Depending on the task you perform, Grid Manager displays task results in the Result page from which you can access newly created objects, such as networks and host records. Note that when a task takes longer than usual to complete, it becomes a long running task. For information about long running tasks, see [About Tasks](#) on page 72.

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

- [The IPAM Task Pack](#) on page 100
- [The Network Automation Task Pack](#) on page 109

The IPAM Task Pack

The IPAM task pack contains the following tasks:

- [Add Networks](#)
- [Add Hosts](#)
- [Add Fixed Addresses](#)
- [Add CNAME Record](#)
- [Add TXT Record](#)
- [Add MX Record](#)
- [CSV Import](#)

Depending on your administrative permissions and the Dashboard template configuration, Grid Manager displays tasks you can access in specific task packs. You can configure your task packs by adding or hiding certain tasks. For information about Dashboard templates, see [About Dashboard Templates](#) on page 114.

To hide tasks in a task pack:

1. Click the Configure icon at the upper right corner of the task pack.
2. In the configuration panel, select the tasks you want to hide from the Active Tasks table. You can use SHIFT+click and CTRL+click to select multiple tasks.
3. Click the left arrow to move the selected tasks to the Available Tasks table.

Click the Configuration icon again to hide the configuration panel after you complete the modification.

Required Licenses for IPAM Tasks

[Table 2.1](#) lists the required licenses for viewing and performing IPAM tasks on the Tasks Dashboard.

Table 2.1 Required Licenses for IPAM Tasks

Task	Required Licenses
Add Networks	DHCP or MSMGMT license
Add Hosts	DNS or DHCP license
Add Fixed Addresses	DHCP or MSMGMT license
Add CNAME Record	DNS or MSMGMT license
Add TXT Record	DNS or MSMGMT license
Add MX Record	DNS or MSMGMT license

For information about how to install licenses, see [Managing Licenses](#) on page 377.

Add Networks

You can create IPv4 and IPv6 networks from the Tasks Dashboard (either from scratch or from a network template that contains predefined properties). You can also create networks from the **Data Management** tab. For more information about IPv4 and IPv6 networks, see [Configuring IPv4 Networks](#) on page 845 and [Configuring IPv6 Networks](#) on page 870.

To add networks from the Tasks Dashboard:

1. Click **Add Networks** in the IPAM task pack and complete the following in the *Add Networks* wizard:
 - **Regional Internet Registry:** This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#) on page 437. Complete the following to create an RIR IPv4 network container or network:
 - **Internet Registry:** Select the RIR from the drop-down list. The default is **None**, which means that the network is not associated with an RIR organization. When you select **RIPE**, the appliance displays **Organization ID** field where you can select an RIR organization.
 - **Organization ID:** Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
 - **Registration Status:** The default is **Not Registered**. When adding an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** check box below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**. The updated status and timestamp are displayed in the **Status of last update** field in the *IPv4 /IPv6 Network Container* or *IPv4/IPv6 Network* editor.
 - **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 or IPv6 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you are adding an existing RIR allocated network to NIOS, select **None**. When you are adding networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
 - **Do not update registrations:** Select this check box if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
 - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
 - **Protocol:** Select **IPv4** to add IPv4 networks and **IPv6** to add IPv6 networks.
 - **Netmask:** Enter the netmask or use the netmask slider to select the appropriate number of subnet mask bits for the network.
 - **Template:** Click **Select Template** to select a network template. When you use a template to create a network, the configuration of the template applies to the new network. If the template specifies a fixed netmask, you cannot edit the netmask in this dialog. You can click **Clear** to remove the template. For information about templates, see [About IPv4 Network Templates](#) on page 829 and [About IPv6 Network Templates](#) on page 837.

NIOS may execute discovery on the newly created network after you save your settings. When you create a network in NIOS, it inherits its discovery capabilities (whether or not it is immediately discovered, its polling settings, and any possible exclusions from discovery), from its parent network (if it has one) or its network container. If the new network is a parent network, it inherits its polling settings from the Grid and its discovery member selection and Enable Discovery action must be defined by the user.

- **Networks:** Do one of the following to add new networks:
 - Click the Add icon to create a new network.

- For IPv4 networks: Grid Manager adds a row to the table. Enter the network address in the **Network** field. Click the Add icon to add another network. You can also select a network and click the **Delete** icon to delete it.
- For IPv6 networks: If you are adding a network for a previously defined global IPv6 prefix, you can select the prefix from the **IPv6 Prefix** drop-down list. The default is **None**, which means that you are not creating an IPv6 network for a previously defined subnet route. If you have defined a global prefix at the Grid level, the default is the global prefix value. Click **Add** and Grid Manager adds a row to the table. Enter the network address in the **Network** field. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2001:0db8:0000:0000:0000:0102:0304 can be shortened to 2001:db8::0102:0304. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Click **Add** again to add another network. You can also select a network and click the **Delete** icon to delete it.

or

Click the Next Available icon to have the appliance search for the next available network. For more information about the next available network, see [About the Next Available Network or IP Address](#) on page 844. Complete the following in the Next Available Networks section:

- **Create new network(s) under:** Enter the network container in which you want to create the new network. When you enter a network that does not exist, the appliance adds it as a network container. When you enter a network that is part of a parent network, the parent network is converted into a network container if it does not have a member assignment or does not contain address ranges, fixed addresses, reservations, shared networks, and host records that are served by DHCP. When you enter a network that has a lower CIDR than an existing network, the appliance creates the network as a parent network and displays a message indicating that the newly created network overlaps an existing network. You can also click **Select Network** to select a specific network in the *Network Selector* dialog box. For information about how the appliance searches for the next available network, see [Obtaining the Next Available Network](#) on page 844.
- **Number of new networks:** Enter the number of networks you want to add to the selected network container. Note that if there is not enough network space in the selected network to create the number of networks specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing networks in the table and you select one, the number you enter here includes the selected network.
- Click **Add Next** to add the networks. Grid Manager lists the networks in the table. You can click **Cancel** to reset the values.

Note: You must click **Add Next** to add the network container you enter in the Next Available Networks section. If you enter a network in the Next Available Networks section and then use the Add icon to add another network, the appliance does not save the network you enter in the Next Available Networks section until you click **Add Next**.

- **Extensible Attributes:** Click the Add icon to enter extensible attributes. Grid Manager adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value. All inheritance attributes which can be inherited from a parent object will be automatically inherited when you add a network. Inheritable extensible attributes that are required are automatically displayed. Optional extensible attributes that are not inheritable are not automatically displayed. For more information about extensible attributes, see [Using Extensible Attributes](#) on page 332.
- If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#) on page 447.

Preview RIR Submissions: Click this to view the updates before the appliance submits them to the RIPE database. This button is enabled only when the registration action is **Create**, **Modify**, or **Delete**, and the **Do not update registrations** check box is not selected.

2. Save the configuration.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

The appliance saves the networks you just created, and Grid Manager displays them in the Result page. When you click a newly created network on this page, Grid Manager displays the **IP Map** panel from which you can view detailed information about the network. For information about the IP Map panel, see [IP Map](#) on page 474.

You can also add and modify other information about the networks you just created. For information about modifying network information, see [Managing IPv4 DHCP Data](#) on page 841 and [Managing IPv6 DHCP Data](#) on page 869.

Add Hosts

Host records provide a unique approach to the management of DNS, DHCP, and IPAM data. By using host records, you can manage multiple DNS records and DHCP and IPAM data collectively, as one object on the appliance. You can add IPv4 and IPv6 addresses to host records from the Tasks Dashboard or the **Data Management** tab. Note that when you add a host record from the Tasks Dashboard, they are configured only for DNS. For more information about Infoblox host records, see [About Host Records](#) on page 459.

To add host records from the Tasks Dashboard:

1. Click **Add Hosts** in the IPAM Task Pack and complete the following in the *Add Hosts* wizard:
 - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the host record.
 - **Zone Name:** Click **Select** to select a DNS zone from the *Zone Selector* dialog box.
 - **Exclude from Network Discovery** and **Immediate Discovery.** When creating the new Host record, you can direct NIOS to immediately discover the host, or to exclude it from network discovery. By default, the Add Hosts task enables immediate discovery.
 - **DNS View:** Displays the DNS view of the selected zone.
 - **Hosts:** Do one of the following to add a host record:
 - Click the Add icon and the appliance adds a row to the table. Complete the following in the table to add a new host record:
 - **Name:** Enter the name of the host record.
 - **Zone:** Displays the DNS zone you select in **Zone Name**. When you enter a different zone here, the appliance displays an error message.
 - **Address:** Enter the IP address you want to associate with this host record.

or

Click the Next Available icon to have the appliance search for the next available IP address for the host record. For information about the next available IP address, see [About the Next Available Network or IP Address](#) on page 844. Complete the following in the Next Available IP section:

- **Create new host addresses under:** Click **Select** to select the network or address range in the *Network/Range Selector* dialog box from which you want the appliance to search for the next available IP address for this host record.
- **Number of new host addresses:** Enter the number of host addresses. Note that if there is not enough space in the selected network or address range to create the number of host addresses specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing host addresses in the table and you select one, the number you enter here includes the selected host address.
- Click **Add Next** to add the IP addresses to their corresponding hosts. Grid Manager lists the host addresses in the table. Ensure that you enter a name for each host record.

Note: When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

- **Extensible Attributes** table: Click the Add icon to enter extensible attributes. The appliance adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value. All inheritance attributes which can be inherited from a parent object will be automatically inherited when you add a host. Inheritable extensible attributes that are required are automatically displayed. Optional extensible attributes that are not inheritable are not automatically displayed. For more information about extensible attributes, see [About Extensible Attributes](#) on page 322.

2. Save the configuration.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information about how to schedule a task, see [Scheduling Tasks](#) on page 75.

The appliance saves the host records you just created, and Grid Manager displays them in the Result page.

When you click a newly created host on this page, Grid Manager displays the **Data Management** -> **DNS** -> **Zones** tab from which you can view information about the host record.

You can also add and modify other information about the host records. For information about modifying host information, see [About Host Records](#) on page 459.

Add Fixed Addresses

You can add IPv4 and IPv6 fixed addresses from the Tasks Dashboard or from the **Data Management** tab. For more information about fixed addresses, see [Configuring IPv4 Fixed Addresses](#) on page 857 and [Configuring IPv6 Fixed Addresses](#) on page 878.

To add fixed addresses from the Tasks Dashboard:

1. Click **Add Fixed Addresses** in the IPAM task pack and complete the following in the *Add Fixed Addresses* wizard:

- **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the fixed address.
- **Protocol:** Select **IPv4** to add IPv4 addresses and **IPv6** to add IPv6 addresses.
- **Template:** Click **Select Template** to select a fixed address template. When you use a template to create a fixed address, the configuration of the template applies to the new fixed address. You can also click **Clear** to remove the template. For information about templates, see [About DHCP Templates](#) on page 826.
- **Exclude from Network Discovery** and **Immediate Discovery.** When creating the new fixed address, you can direct NIOS to immediately discover the device associated with the fixed address, or to exclude it from network discovery. By default, the Add Fixed Addresses task enables immediate discovery.
- **Addresses:** Do one of the following to add fixed addresses:

Click the Add icon and Grid Manager adds a row to the table. Complete the following to create fixed addresses:

- For IPv4 fixed addresses: Enter the IPv4 address and MAC address. Click the Add icon to add another fixed address.
- For IPv6 fixed addresses: Enter the IPv6 address and DUID. Click the Add icon again to add another fixed address.

or

Click the Next Available icon to have the appliance search for the next available address. Complete the following:

- **Create new fixed addresses under:** Click **Select** to select the network or address range in the *Network/Range Selector* dialog box from which you want the appliance to search for the next available IP address for this fixed address.
- **Number of new fixed addresses:** Enter the number of fixed addresses you want to add to the selected network or address range. Note that if there is not enough space in the selected network or address range to create the number of fixed addresses specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing fixed addresses in the table and you select one, the number you enter here includes the selected fixed address.

- Click **Add Next** to add the fixed addresses. The appliance lists the fixed addresses to the table. Ensure that you enter the MAC address or DUID for each fixed address.

Note: When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

- **Extensible Attributes** table: Click the Add icon to enter extensible attributes. The appliance adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value. All inheritance attributes which can be inherited from a parent object will be automatically inherited when you add a fixed address. Inheritable extensible attributes that are required are automatically displayed. Optional extensible attributes that are not inheritable are not automatically displayed. For more information about extensible attributes, see [About Extensible Attributes](#) on page 322.

2. Save the configuration.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

The appliance saves the fixed addresses you just created, and Grid Manager displays them in the Result page. When you click a newly created fixed address on this page, Grid Manager displays the **Data Management -> IPAM -> IP Map** or **List** tab from which you can view information about the fixed address.

You can also add and modify other information about the fixed addresses you just created. For more information about modifying fixed address information, see [Managing IPv4 DHCP Data](#) on page 841 and [Managing IPv6 DHCP Data](#) on page 869.

Add CNAME Record

A CNAME record maps an alias to a canonical name. You can use CNAME records in both IPv4 forward- and IPv4 reverse-mapping zones to serve two different purposes. (At this time, you cannot use CNAME records with IPv6 reverse-mapping zones.) For more information about CNAME records, see [Managing CNAME Records](#) on page 669.

To add a CNAME record from the Tasks Dashboard:

1. Click **Add CNAME Record** in the IPAM task pack and complete the following in the *Add CNAME Record* wizard:
 - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the CNAME record.
 - **Alias:** Click **Select Zone** to select a DNS zone from the *Zone Selector* dialog box. If you have only one zone, Grid Manager displays the zone name here when you click **Select Zone**. Enter the alias for the canonical name. For an IPv4 reverse-mapping zone, enter the host portion of an IP address. For example, if the full IP address is 10.1.1.1 in a network with a 25-bit netmask, enter **1**. (The 10.1.1.0/25 network contains host addresses from 10.1.1.1 to 10.1.1.126. The network address is 10.1.1.0, and the broadcast address is 10.1.1.127.)
 - **DNS View:** Displays the DNS view of the selected zone.
 - **Canonical Name:** This field displays the domain name of either the current zone or the last selected zone. To add a CNAME record to a forward-mapping zone, enter the complete canonical (or official) name of the host. To add a CNAME record to a reverse-mapping zone, enter *host_ip_addr.prefix.network.in-addr.arpa* (host IP address + 2317 prefix + network IP address + in-addr.arpa). For example, enter 1.0.25.1.1.10.in-addr.arpa. This IP address must match the address defined in the PTR record in the delegated child zone.
 - **Comments:** Enter useful information about this record.
 - **Disable:** Select the check box to disable the record without deleting its configuration. Clear the check box to enable the record.
 2. Save the configuration, or click **Next** to define extensible attributes. For information about extensible attributes, see [Using Extensible Attributes](#) on page 332.
- or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

3. Click **Restart** if it appears at the top of the screen.

You can also add and modify other information about the CNAME record you just created. For more information about modifying the CNAME record, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Add TXT Record

A TXT (text record) record contains supplemental information for a host. For example, if you have a sales server that serves only North America, you can create a text record stating this fact. You can create more than one text record for a domain name. You can add a TXT record from the Tasks Dashboard or the **Data Management** tab. For more information about TXT records, see [Managing TXT Records](#) on page 668.

To add TXT records from the Tasks Dashboard:

1. Click **Add TXT Record** in the IPAM task pack and complete the following in the *Add TXT Record* wizard:
 - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the TXT record.
 - **Name:** If Grid Manager displays a zone name, enter the name to define a TXT record for a host or subdomain. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box. Then, enter the TXT record name. The appliance prefixes the name you enter to the domain name of the selected zone. For example, if you want to create a TXT record for a web server whose host name is `www2.corp100.com` and you define the TXT record in the `corp100.com` zone, enter **www2** in this field. To define a TXT record for a domain whose name matches the selected zone, leave this field empty. The appliance automatically adds the domain name (the same as the zone name) to the TXT record. For example, if you want to create a TXT record for the `corp100.com` domain and you selected the `corp100.com` zone, leave this field empty.
 - **DNS View:** Displays the DNS view of the selected zone.
 - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
 - **Text:** Enter the text that you want to associate with the record. It can contain substrings of up to 255 bytes, up to a total of 512 bytes. Additionally, if you enter leading, trailing, or embedded spaces in the text, add quotes around the text to preserve the spaces. For example: `" v=spf1 include:corp200.com -all "`
 - **Comments:** Enter useful information about this record.
 - **Disable:** Select the check box to disable the record without deleting its configuration. Clear the check box to enable the record.
2. Save the configuration, or click **Next** to define extensible attributes. For information, see [About Extensible Attributes](#) on page 322.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

3. Click **Restart** if it appears at the top of the screen.

Add MX Record

An MX (mail exchanger) record maps a domain name to a mail exchanger. A mail exchanger is a server that either delivers or forwards mail. You can specify one or more mail exchangers for a zone, as well as the preference for using each mail exchanger. A standard MX record applies to a particular domain or subdomain. You can add an MX record from the Tasks Dashboard or the **Data Management** tab. For more information about MX records, see [Managing MX Records](#) on page 665.

To add MX records from the Tasks Dashboard:

1. Click **Add MX Record** in the IPAM task pack and complete the following in the *Add TXT Record* wizard:
 - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the MX record.
 - **Mail Destination:** If Grid Manager displays a zone name, enter the mail destination here. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the mail destination. If you want to define an MX record for a domain whose name matches the zone you selected, leave this field blank. Grid Manager automatically adds the domain name (the same as the zone name) to the MX record. For example, if you want to create an MX record for a mail exchanger serving the corp100.com domain and you selected the corp100.com zone, and leave this field empty.
 If you want to define an MX record for a subdomain, enter the subdomain name. The appliance prefixes the name you enter to the domain name of the selected zone. For example, if you want to create an MX record for a mail exchanger serving site1.corp100.com—a subdomain of corp100.com—and you define the MX record in the corp100.com zone, enter site1 in this field.
 If you want to define an MX record for a domain and all its subdomains, enter an asterisk (*) to create a wildcard MX record.
 - **DNS View:** Displays the DNS view of the selected zone.
 - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
 - **Host Name Policy:** Displays the hostname policy of the selected zone. Ensure that the hostname you enter complies with the hostname restriction policy defined for the zone.
 - **Mail Exchanger:** Enter the fully qualified domain name of the mail exchanger.
 - **Preference:** Select an integer from 10 to 100, or enter a value from 0 to 65535. The preference determines the order in which a client attempts to contact the target mail exchanger.
 - **Comment:** Enter useful information about this record.
 - **Disable:** Select the check box to disable the record without deleting its configuration. Clear the check box to enable the record.
2. Save the configuration, or click **Next** to define extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 or
 Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information about scheduling tasks, see [Scheduling Tasks](#) on page 75.
3. Click **Restart** if it appears at the top of the screen.

CSV Import

You can access **Import Job Manager** and perform CSV imports, manage import jobs, and view import status. You can perform CSV imports from the Task Dashboard and the Toolbar. Selecting **CSV Import** launches the **Import Job Manager**, which allows for importing of data, managing import jobs, and viewing import status. For detailed information about the CSV import feature, see [About CSV Import](#) on page 86.

Enabling the Network Automation Tasks

The Network Automation Tasks pack requires the configuration, licensing and connection of an Infoblox Network Automation appliance to support automation tasks.

Should two superusers be logged in to the NIOS system and one superuser enables the Network Automation Tasks pack on their console, the other superuser will not see the task pack on their console until their next login; the **Disable Network Automation Tasks** from the Configure icon menu shows the correct state.

You install the Network Automation appliance into the managed network, ensuring the appliance is reachable by the NIOS Grid Master. After this is accomplished, you register the Network Automation appliance with NIOS.

1. Click the Configure icon at the top right corner of the **Tasks** page.
2. Choose **Enable Network Automation Tasks**.

The Enable Network Automation Tasks dialog box appears, requesting verification of your action:

Though you can see the change immediately, other users who are currently logged in will not see the change until they log in again.
Are you sure you want to proceed?

3. Click **Yes** to enable the Network Automation Tasks set.
 After a moment, the Automation Tasks panel appears.

Disabling Network Automation Tasks

Should you need to disable the Network Automation Tasks pack, do the following:

1. Click the Configure icon at the top right corner of the **Tasks** page.
2. Choose **Disable Network Automation Tasks**.

The Disable Network Automation Tasks dialog box appears, requesting verification of your action:

Though you can see the change immediately, other users who are currently logged in will not see the change until they log in again.
Are you sure you want to proceed?

3. Click **Yes** to disable the Network Automation Tasks pack.

Registering Network Automation with NIOS

You must register a Network Automation appliance with NIOS to support the Network Automation Tasks. This registration is done directly on the NIOS system. You need the admin account and password for the Network Automation appliance and its hostname or IP address.

Note that when you register Network Automation with a NIOS HA pair, you can register only one interface at a time. Use the IP address of the LAN1 interface, not the VIP address, for registration. When an HA failover occurs, the Network Automation registration is disabled. You can register the Network Automation appliance again after the failover.

1. From the **Dashboards** tab, select the **Tasks** tab.
2. At the top right corner of the Automation Tasks panel, click the Configure icon -> **NetMRI Registration**.
3. In the **NetMRI Registration** dialog, do the following:

- a. Enter the IP address or resolved host name of the Network Automation appliance supporting the Automation task pack.
- b. Enter the **Admin Password**.

This information is specific to the Infoblox Network Automation appliance supporting the Automation tasks in NIOS.

4. Click **Register** to commit settings.

After registration, the **NetMRI Registration** menu item changes to read **NetMRI Deregistration** to support disconnecting from the Network Automation appliance.

You can also start Network Automation from the NIOS Dashboards page.

1. From the **Dashboard** tab, select the **Tasks** tab.
2. In the **Automation Tasks** pane, click the down arrow gadget and select **Launch NetMRI**.

Network Automation launches in a new browser tab.

To check on script executions, go to **Configuration Management** -> **Job Management** side tab -> **Scripts** and check the Last Run column.

The NIOS Task Viewer (see [Using the Task Viewer to View Job Logs and Approve Jobs](#)) also provides the log history of automated jobs.

The Network Automation Task Pack

The Network Automation task pack contains the following tasks:

- [Network Provisioning Task](#)
- [Using the Port Activation Automation Task](#)
- [Specifying a Port Activation Script](#)
- [Assigning a New Script to the VLAN Reassignments Task](#)
- [Rogue DHCP Server Remediation](#)

Depending on your administrative permissions, Grid Manager displays tasks you can access in specific task packs. You can configure your task packs by adding or hiding certain tasks.

To hide tasks in a task pack:

1. Click the Configure icon at the upper right corner of the task pack.
2. In the configuration panel, select the tasks you want to hide from the Active Tasks table. You can use SHIFT+click and CTRL+click to select multiple tasks.
3. Click the left arrow to move the selected tasks to the Available Tasks table.

Click the Configuration icon again to hide the configuration panel after you complete the modification.

Network Automation Task Options

Tasks allow the assignment of job scripts to change and expand task functionality. These scripts reside on the Network Automation appliance and must be readable by the NIOS system to run the automation tasks. You can also select different scripts to execute for automation tasks that provide that feature in NIOS. Three network automation tasks allow for the choosing of non-default scripts for task operation:

- [Network Provisioning Task](#)
- [Using the Port Activation Automation Task](#)
- [Specifying a Port Activation Script](#)

Network Provisioning Task

The Network Provisioning task runs in two modes: a basic mode with a much shorter list of configuration options, and a more complex mode that provides detailed configuration for provisioning a network, including the use of NIOS network views, extensible attributes and network templates.

New networks can be provisioned on routed networks and on switched networks. In the latter case, you can specify the new VLAN number and VLAN name for provisioning, along with the Device Group Device and Interface. the Device Group values are taken from the Device Groups defined on the Network Automation appliance from which NIOS obtains its data.

Network Provisioning supports two types of networks: **IPv4**, in which the new network is IPv4 only, and **IPv4 and IPv6**, in which the new network runs both protocol stacks.

Simple vs. Complex Provisioning

Use of a Network View determines whether you use the simple or detailed views of provisioning a network. A network view is a single routing domain with its own networks and shared networks. In NIOS, all networks must belong to a network view. You can manage networks in one network view independently of other network views. Because network views are mutually exclusive, the networks in each view can have overlapping address spaces with multiple duplicate IP addresses without impacting network integrity.

Also, the same network segment can be present in multiple network views. When you create a new network, you select one view in which to place it, and preserve those values to apply to another view.

You also have the option to provision a single network segment without recourse to NIOS network views. The simple network provisioning option (accessible by simply clicking the IPv4 tool at the top of the Network Provisioning dialog box) allows you to specify as few as three values to configure a network.

The NIOS system also provides a **default** network view, which appears as an option for network provisioning.

If a single network view is configured in NIOS, you will not see a **Network View** option in the Network Provisioning task.

Applying Extensible Attributes

Extensible attributes are associated with a specific network view, and are referenced by the Network Provisioning task. Should you configure a new network using a network view, you may need to consider the application of extensible attributes to the new network (they are not automatically applied, but will appear in the Network Provisioning dialog if those attributes are defined in the chosen Network View). Extensible attributes are generally defined for descriptive and tracking purposes in the network. A network view may have attributes such as Building, Country, Region, Site, State or VLAN, for example. Attributes are defined for network views in NIOS but are not defined by the Network Automation appliance.

If the NIOS system supports only a single network view, no View selections are made for the purposes of network provisioning.

Required settings for provisioning a new network will show a red asterisk (*) by the option name.

To perform an automatic network provisioning task:

1. From the **Dashboards** tab, select the **Tasks** tab -> **Network Provisioning**.
2. Select the network Type for provisioning: **IPv4** or **IPv4 and IPv6**.
3. To configure IPv4 provisioning.
 - a. Enter the **Parent Network** value (or click Select Network to choose the parent network from a list if using a Network View).
 - b. Choose the **Network Template** from the drop-down list if one is provided by the chosen Network View. The Network template is otherwise optional.

The Provision Network task provides subnetting tools.

- c. Drag the **Netmask** slider to the required CIDR mask bit depth (1-32).
 - d. In the **New Network** field, enter the IP prefix for the new network.
 - e. In the **Router Address** field, enter the IP address for the router interface.
 - f. Select any **Extensible Attributes** in the list if they are provided; otherwise, you can create new ones by clicking Add and choosing the **Attribute Name**, **Value** and the **Required** setting.
4. To configure IPv6 provisioning.
 - a. Enter the **Parent Network** value (or click Select Network to choose the parent network from a list if using a Network View).
 - b. Choose the **Network Template** from the drop-down list if one is provided by the chosen Network View. The Network template is otherwise optional.

The Provision Network task provides subnetting tools.

 - c. Drag the **Netmask** slider to the required CIDR mask bit depth (1-32).
 - d. In the **New Network** field, enter the IP prefix for the new network.
 - e. In the **Router Address** field, enter the IP address for the router interface.
 - f. Select any **Extensible Attributes** in the list if they are provided; otherwise, you can create new ones by clicking Add and choosing the **Attribute Name**, **Value** and the **Required** setting.
5. Enter the required name value in the **Interface Hostname** field. (Examples include “eth0” and “serial0.”)
6. Select the DNS **Zone** under which the hostname operates.
7. Choose a device group from the **Device Group** drop-down list.
8. From the **Device** drop-down list, choose the switch or router on which the network will originate.
9. If the selected device is a router, the **VLAN Number** and **VLAN Name** fields will be disabled.

- From the **Interface** list, choose the interface to which the network will be reassigned. The drop-down list contains all the interfaces from the chosen network device, and also shows the ports' respective states (up/down, up/up and so on).

If an interface shows **Routed** or **Switched**, it cannot be selected for provisioning as it is already being used as part of an active network.

- If the chosen device is a switch, enter the new **VLAN Number** on which the new network segment runs.
- If the chosen device is a switch, enter the new **VLAN Name** on which the new network segment runs.
- Click **Provision Network** to commit settings.

The system sends the configuration request to the Network Automation appliance and displays the task configuration sequence.

You can start Network Automation from the registered Network Automation appliance to check job execution.

- In the **Automation Tasks** pane, click the down arrow gadget and select **Launch TAE**.

Network Automation will launch in a new browser tab. To check on script executions, go to **Configuration Management** → **Job Management side** tab → Job History and view details about provisioning jobs and other jobs that execute as a result of NIOS-based automation tasks.

Defining Options for the Network Provisioning Task

The Network Provisioning task provides several configuration options that affect how the task operates.

Hostname provisioning for interfaces is useful for troubleshooting purposes in the network, usually to ensure that an admin knows which router interface they are connecting through to communicate with the device. The hostname value is actually provisioned from within the Network Provisioning task. Enabling the Hostname Required? check box sets the Network Automation appliance to provision the network with hostnames applied to the router interfaces for easier identification.

Network provisioning requires that the system know exactly which IP address the gateway for the network will reside. For provisioning most networks, an Offset value of 1 indicates that the provisioned network gateway IP address ends with the host address of *.*.1, as in 192.168.1.1. An Offset value of 1 will be by far the most common value for provisioning networks. Specifying an offset value other than 1 indicates that the gateway IP is a specified number of host values from the prefix address of the network. For example, setting an IPv4 Gateway Address Offset of 12 indicates that the IP for the gateway ends in *.*.12, as in 10.1.1.12. Offsets work the same way for any size network: for an example such as 10.1.1.64/26, and an offset of 12, the provisioned gateway IP would be 10.1.1.76. Make sure the defined offset value lies within the addressable boundaries of the provisioned network!

The same principles also apply for IPv6 networks, except that the IPv6 value is entered manually in hexadecimal instead of being selected from a drop-down list. Most provisioned IPv6 networks will use a /64 network address.

You can also select a different script from the default for the Network Provisioning task.

To define settings for the Network Provisioning automated task:

- From the **Dashboards** tab, select the **Tasks** tab. Under the **Network Provisioning** task, click the settings icon on the top right.
- If the provisioning process requires a hostname, enable the **Hostname Required?** check box. (The network interface hostname ("eth0," "serial0") and the Zone that it belongs to are defined in the Network Provisioning task.)
- Choose a gateway offset value from the **IPv4 Gateway Address Offset** drop-down list. If no value is selected, the offset value defaults to 1 for the provisioned network address.
- If an IPv6 offset is required for provisioning an IPv6 network or for provisioning a network that supports both IPv4 and IPv6 addressing, enter the **IPv6 Gateway Address Offset** value in hexadecimal. If no value is entered, the offset value defaults to 0000.0000.0000.0001 for the provisioned network address, indicating an offset value of 1 for the gateway IP address.
- In the **Script Name** dropdown, choose the script that you wish to run for the Port Activation task. The scripts are located on the Trinzi Automation 4000 appliance, and referenced for use by NIOS. By default, the bundled **Port Activation** script is selected.
- Click **Save** to commit settings.

- Click Cancel to close the dialog.

The system sends the request to the Network Automation appliance and displays a **Provisioning Network Config updated** notification message.

Using the Port Activation Automation Task

The Port Activation task provides a central console on which the interfaces for any device anywhere in the managed network can be conveniently enabled or disabled. Ports can be taken administratively Up or Down using this task, and all interfaces on a selected device can be activated or deactivated with a series of mouse clicks.

- From the **Dashboards** tab, select the **Tasks** tab -> **Port Activation**.
- Choose the **Device Group** from the drop-down list.
- From the **Device** drop-down list, choose the network device on which port activation will be executed.
The **Interfaces** table lists all interfaces on the current device. The **VLAN** and **VLAN Name** columns list the VLAN assigned to each port (VLAN 1/Default resides on all ports without an explicit VLAN assignment). The **OP Status** column will show the current state of each interface.
- Scroll down the table to locate the interface(s) you want to activate.
- From the **Admin Status column**, select **Up** (or **Down**) from the drop-down list for the chosen interface.
- Set any other interfaces on the current device based on your assigned task.
- Click **Apply** to commit settings.

The system sends the request to the Network Automation appliance and displays the task configuration sequence.

The Port Activation task will also write the full running configuration to memory, making it the saved configuration. If the user made a change to the running configuration, in parallel with the port activation change, and did not save it, those changes will also be saved.

Specifying a Port Activation Script

The Port Activation task provides a central console on which the interfaces for any device in the managed network can be conveniently activated. Ports can be taken administratively Up or Down using this task, and all interfaces on a selected device can be activated or deactivated with a series of mouse clicks.

The Network Automation appliance provides the ability to create new automation scripts for many purposes. You may, for example, wish to create a new Port Activation script and use that as an automation task.

To select a different script from the default choice in the software:

- From the **Dashboards** tab, select the **Tasks** tab. Under the **Port Activation** task, click the settings icon.
- For Port Activation Options, choose a new script from the **Script Name** drop-down list. The scripts are located on the Trinetic Automation 4000 appliance, and automatically referenced for use by NIOS. By default, the bundled **Port Activation** script is selected.
- Click **Save** to commit settings.

The system sends the request to the Network Automation appliance and displays a notification message.

VLAN Reassignment

VLANs can be reassigned to new interfaces on individual L2/L3 switches in the managed network. A VLAN can have a path across several switches; when you make changes on a given switch, make sure that the path is maintained.

To ensure end-to-end connectivity, you may need to change VLAN port assignments on more than one switch in the path. This feature operates with the VLAN Trunking Protocol (VTP). VLAN switching is changed across one port per switch at a time. Should you need to change VLAN assignments across more than one switch in the path, plan accordingly.

VLANs must already be configured on the switch(es) being changed, and be detected by the Network Automation appliance.

- From the **Dashboards** tab, select the **Tasks** tab -> **VLAN Reassignment**.

2. Begin by selecting the **Device Group** from the drop-down list. For **VLAN Reassignments**, you typically choose the Switching device group.
3. From the **Device** drop-down list, choose the switch on which port reassignment will be executed.
4. From the **Port** list, choose the interface to which the VLAN will be reassigned. The **Port** list also shows the Administrative and Operational states of each interface on the current device (Administratively Up/Operationally Down, for example.)

Note: You can reassign a VLAN to a port that is operationally or administratively Down.

The **Current VLAN** value will show the VLAN to which the selected interface is currently assigned.

5. Choose the new VLAN value for port reassignment from the **New VLAN** drop-down list.
6. Click **Move VLAN** to commit settings.

The system sends the configuration request to the Network Automation appliance and displays the task configuration sequence.

The VLAN Reassignment task will also write the full running configuration to memory, making it the saved configuration. If the user made a change to the running configuration, in parallel with the port activation change, and did not save it, those changes will also be saved.

Assigning a New Script to the VLAN Reassignments Task

The Network Automation appliance provides the ability to create new automation scripts for many purposes. You can create and assign a new VLAN Reassignment script and use that for the automation task.

To select a different script from the default choice in the software:

1. From the **Dashboards** tab, select the **Tasks** tab. Under the **VLAN Reassignment** task, click the settings icon.
2. For Port Activation Options, choose a new script from the **Script Name** drop-down list.
3. Click **Save** to commit settings.

The system sends the request to the Network Automation appliance and displays a notification message.

The VLAN Reassignment task will also write the full running configuration to the device's memory, making it the saved configuration. If the user made a change to the running configuration, in parallel with the port activation change, and did not save it, those changes will also be saved.

Provision Bare Metal Device

The **Provision Bare Metal Device** automated task enables automated installation of new switches and routers into the network. The Trinzic Automation task enables cost and convenience savings by detecting the default behavior of new devices on the network, pointing them to customized TFTP servers from which standardized bare-metal configuration files are downloaded and installed onto the new devices.

The **Provision Bare Metal Device** automated task does not provide NIOS-based optional settings; configuration for this task is done in the Trinzic Automation 4000 Network Automation user interface. The automated task is automatically triggered by detection of a network device requiring configuration.

Rogue DHCP Server Remediation

All DHCP servers on the network should be under administrative control. If any device offering DHCP leases to clients on the network is not properly administered, it violates many security guidelines and at the very least may cause configuration problems throughout the network. Some events may be unwitting or innocuous (an office worker installing a wireless access point in their cube to share a resource), or may be an attempt to hijack clients and steal information. To prevent such issues, the Rogue DHCP Server Remediation task enables the detection, location and isolation of such devices.

The **Provision Bare Metal Device** automated task does not provide NIOS-based optional settings; configuration for this task is done in the Trinzic Automation 4000 Network Automation user interface. The automated task is automatically triggered by detection of a network device requiring remediation.

Using the Task Viewer to View Job Logs and Approve Jobs

You can view the logged results from any task run from the Automation Tasks dashboard through a pair of information pages, which are accessed through the Task Viewer window.

A **Job History** page provides a log history of all TAE tasks that have recently run, including all Automation Task types in the dashboard.

A second page, **Issues & Approvals**, provides links to two important items: **Issues**, which displays details about any network issue related to TAE tasks and jobs in an Issue Viewer page from the Network Automation appliance, and **Approvals**, which are jobs that must be approved before the Network Automation appliance can execute the job. For example, the **Isolate Rogue DHCP Server** job must be approved before it will run and attempt to isolate the detected rogue DHCP server in the network.

1. From the **Dashboards** tab, select the **Tasks** tab.
2. In the **Automation Tasks** pane, click the down arrow gadget and select **Task Viewer**. The Task Viewer window appears, displaying a scrollable and sortable Job History table. Important columns include the **Start Time**, the **Job ID** (a numeric value with a clickable link to the TAE Job Details Viewer, which will open in a new browser tab), the **Job Name**, the **User** account that executed the task, the job **Status** and the **# Devices** (the number of devices) against which the task ran.

The Job History page shows the most recent subset of executed TAE jobs. A yellow bar at the top of the table provides a **click here to see more** link, which takes the user to the Network Automation appliance **Job History** page in a new browser tab.

3. If an item appears in the **Issues & Approvals** page, click the link in the **Action** column. You will typically see two different link types: **Issue Details** or **Approve Job**.
 - a. To view an issue in more detail: Clicking an **Issue Details** link displays the Network Automation appliance **Job Details** page in a new browser tab for the selected job.
 - b. To approve a job: Clicking an **Approve Job** link displays the Summary page of the Network Automation **Job Wizard**, with an **Approve Job** button.
4. Click **Close** to close the Task Viewer.

In the Network Automation appliance, you can also check **Configuration Management** → **Job Management** side tab – > **Job History** and view details about any jobs that execute as a result of NIOS-based automation tasks.

ABOUT DASHBOARD TEMPLATES

Superusers can specify the tasks an admin group can perform from the **Tasks Dashboard** tab by creating a dashboard template and assigning it to the admin group. When you create a dashboard template, you define the tasks users in an admin group can perform and specify whether the users can configure their own dashboards when they log in to Grid Manager. When you assign a dashboard template to an admin group, all users in this group can see and perform the tasks you define in the template, provided that the users also have the correct permissions to the objects related to the tasks. For information about administrative permissions, see [Administrative Permissions for Dashboard Tasks](#) on page 214. If the assigned template is unlocked, users can configure tasks on their dashboard. If you lock the dashboard template, users cannot configure task packs on their own dashboards.

Superusers can also restrict limited-access users to access only the **Tasks Dashboard** tab when they log in to Grid Manager. These users cannot manage other core network services through Grid Manager. They can only see the **Tasks Dashboard** tab and access only the tasks defined in the dashboard template, if applicable. This feature is useful when you want to define different levels of admin users and restrict them to specific tasks based on their organizational functions. For information about how to set this restriction, see [Creating Limited-Access Admin Groups](#) on page 156.

To configure and apply dashboard templates, complete the following:

1. Configure dashboard templates, as described in [Adding Dashboard Templates](#) on page 115.
2. Assign dashboard templates to admin groups, as described in [About Admin Groups](#) on page 154.

Adding Dashboard Templates

Only supersuers can configure dashboard templates. Limited-access users may configure task packs depending on the configuration of their assigned dashboard templates.

To add a dashboard template:

1. Log in as a superuser.
2. From the **Dashboards** -> **Tasks** tab, click the Configure icon at the top right corner of a task pack.
3. Select tasks from the Active Tasks table and use the left arrow to move them to the Available Tasks table to hide the tasks, and vice versa. Grid Manager displays the tasks you place in the Active Tasks table. Repeat the steps for all task packs.
4. At the top right corner of the Tasks Dashboard panel, click the Configure icon -> **Configure Template**.
5. In the Dashboard template configuration section, click **Create new template**.
6. In the *Save Dashboard Template* dialog box, complete the following:
 - **Name:** Enter a name for the new dashboard template.
 - **Locked:** When you select this check box and assign this template to an admin group, users in the admin group can only perform the tasks you configure to appear in this template. They cannot configure their dashboards. When you clear this check box, users can still only see the tasks you configure for this template, but they can now configure tasks in the task packs on their dashboards. Note that when you lock a template, it applies to all users in the admin group, including those who have customized dashboards.
7. Click **Save & Close**.
The appliance saves the template and adds it to the **Template** drop-down list.

Resetting Dashboard Templates

Only users with an unlocked dashboard template assigned can reset their dashboards to the template that was originally assigned to them. Users with locked dashboard template cannot configure or reset their dashboards. Also, only superusers can configure dashboard templates.

To reset a dashboard template:

1. Select the **Dashboards** -> **Tasks** tab.
2. For superusers: At the top right corner of the Tasks Dashboard panel, click the Configure icon -> **Reset**. Note that the Configure icon appears only if you are a superuser.
For limited-access users: At the top right corner of the Tasks Dashboard panel, click **Reset**.
The appliance reset your dashboard to the original dashboard template that was assigned to your admin group.

Modifying Dashboard Templates

You can modify an existing dashboard template by locking or unlocking it, and adding or removing tasks from a task pack. However, you cannot change the name of the template. When you change the name of a template, the appliance clones the template and adds the new template to the list. Note that when you modify a locked template that is assigned to an admin group, users in the group automatically adopt the changes you make to the template the next time they log in to Grid Manager.

To modify a dashboard template:

1. From the **Dashboards** -> **Tasks** tab, click the Configure icon at the top right corner of the panel.
2. In the Dashboard template section, select the template you want to modify from the **Template** drop-down list.
Note that Grid Manager displays [L] before the name of a locked template.

3. In the task pack, click the Configure icon at the top right corner.
4. Select tasks from the Active Tasks table and use the left arrow to move them to the Available Tasks table to hide the tasks, and vice versa. Grid Manager displays the tasks you place in the Active Tasks table. Repeat the steps for all task packs.
5. Click **Save**.
6. In the *Save Dashboard Template* dialog box, modify other information, as described in [Adding Dashboard Templates](#).
7. Click **Save & Close**.

Deleting Dashboard Templates

Only superusers can delete dashboard templates. To delete a dashboard template that is currently assigned to an admin group, you must first unassign the template from the admin group. For more information, see [Creating Limited-Access Admin Groups](#) on page 156.

To delete a dashboard template:

1. From the **Dashboards** -> **Tasks** tab, click the Configure icon at the top right corner of the panel.
2. In the Dashboard template section, select the template you want to delete from the **Template** drop-down list.
3. Click **Delete**.
4. In the *Delete Dashboard Template* dialog box, click **Yes**.

Assigning Dashboard Templates

After you create a dashboard template, you can assign it to an admin group. Admin users in this admin group can access the tasks you define in the template.

To assign a dashboard template to an admin group, see [About Admin Groups](#) on page 154.

STATUS DASHBOARDS

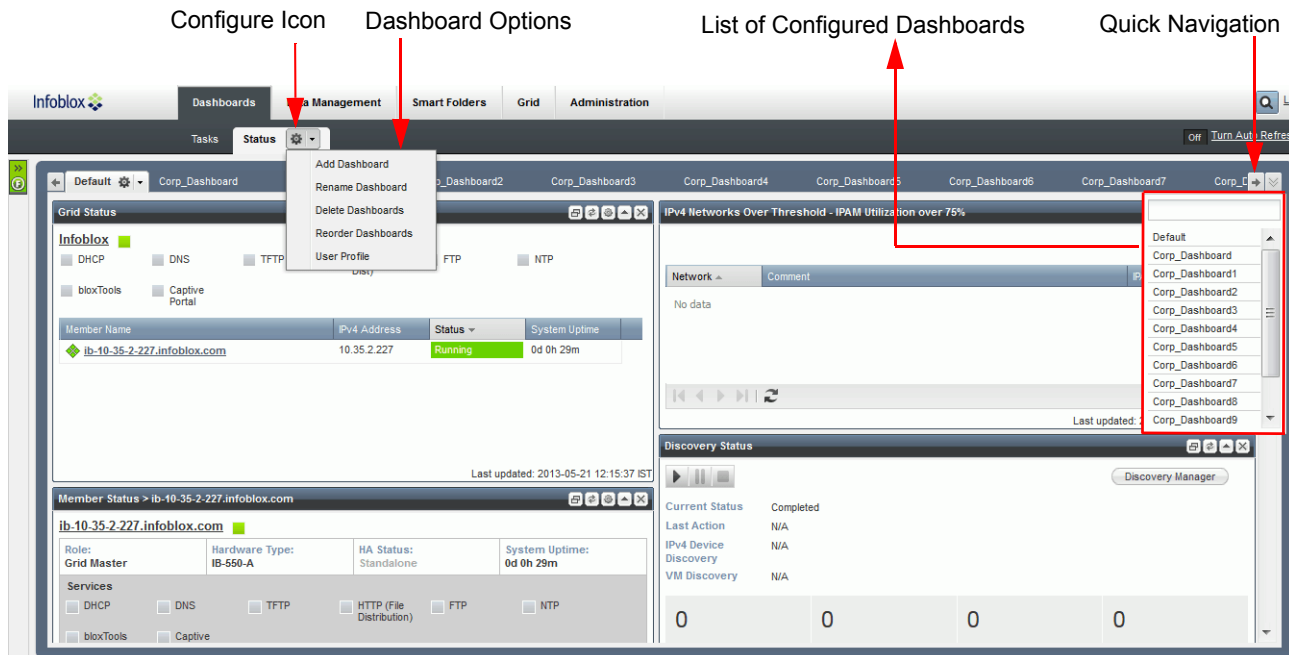
A status dashboard contains widgets from which you can view and manage data. Widgets are the building blocks of status dashboards. For more information about widgets, see [Adding Widgets to Dashboards](#) on page 117. They provide information about different aspects of your Grid and networks. For example, the *Member Status* widget provides general information about a Grid member, and the *Network Statistics* widget provides data for a specified network.

The appliance provides a default status dashboard. Grid Manager displays the default dashboard only when there are more than one widget on the dashboard. You can add and modify widgets in the default dashboard, but you cannot rename or delete it. From a dashboard, you can access your most commonly accessed tasks and monitor appliance status. You can configure your own status dashboards to which you can add widgets that help you manage different data. Configuring multiple status dashboards helps organize widgets in a meaningful way and improves dashboard and widget performance. This is especially useful when you have a Grid serving a large number of Grid members. When you configure a new dashboard, you can use the existing dashboard as a template. You can create up to 100 copies at a time using the **Add Dashboard** option. For information about how to add status dashboards, see [Adding Status Dashboards](#) on page 119.

You can add widgets to different dashboards, however, you can add only one widget at a time on each dashboard. The default number of widgets per dashboard is 10. The maximum number of widgets that you can add on each dashboard is 20 at a time. You can define the number of widgets that can be configured on each dashboard in **User Profile**. This limitation applies only to dashboards that you configure and does not apply to the default dashboard. For information about how to specify the widget limit, see [Configuring Widget Limit per Dashboard](#) on page 119.

If you have configured a lot of status dashboards, you can use the Quick Navigation icon to quickly access each status dashboard. For information, see [Using Quick Navigation](#) on page 120. [Figure 2.1](#) illustrates the typical layout in Grid Manager after you configure multiple status dashboards.

Figure 2.1 Status Dashboard



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

- Add new status dashboards, as described in [Adding Status Dashboards](#) on page 119.
- Rename a dashboard, as described in [Renaming Status Dashboards](#) on page 120.
- Copy or move a widget, as described in [Copying or Moving Widgets](#) on page 120.
- Reorder dashboards, as described in [Reordering Status Dashboards](#) on page 121.
- Delete dashboards, as described in [Deleting Status Dashboards](#) on page 121.
- Configure widget limit, as described in [Configuring Widget Limit per Dashboard](#) on page 119.

Adding Widgets to Dashboards

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

- [Grid Status](#)
- [Grid Upgrade Status](#)
- [Member Status \(System Status\)](#)
- [DNS Statistics](#)
- [Ranges Over Threshold](#)
- [IPv4 Failover Associations Status](#)
- [DHCP Statistics](#)
- [Network Statistics](#)
- [IPv4 Networks Over Threshold](#)
- [Discovery Status](#)
- [Advanced Discovery Status](#)

- [My Commands](#)
- [DDNS Statistics](#)
- [System Activity Monitor](#)
- [File Distribution Statistics](#)
- [Active WebUI Users](#)
- [Microsoft Servers Status Widget](#)
- [Import Job Manager](#)
- [Load Balancer Status](#)
- [Pending Approvals](#)
- [Response Policy Zone \(RPZ\) Statistics](#)
- [Infoblox Community](#)
- [Mobile Devices Status](#)
- [Threat Protection Statistics](#)

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

To add widgets to your dashboard:

1. **Default Status Dashboard:** From the **Dashboards** -> **Status** tab, click the Configure icon -> **Add Content**. This is applicable when you have the default dashboard only.

Configured Status Dashboards: From the **Dashboards** -> **Status** tab, select the configured status dashboard, click the Configure icon -> **Add Content**.

Grid Manager displays thumbnails of the available widgets. Use the scroll bar on the right to scroll through the widgets, as illustrated in the [Figure 2.2](#).

2. Select and drag a widget to the desired location on your dashboard.

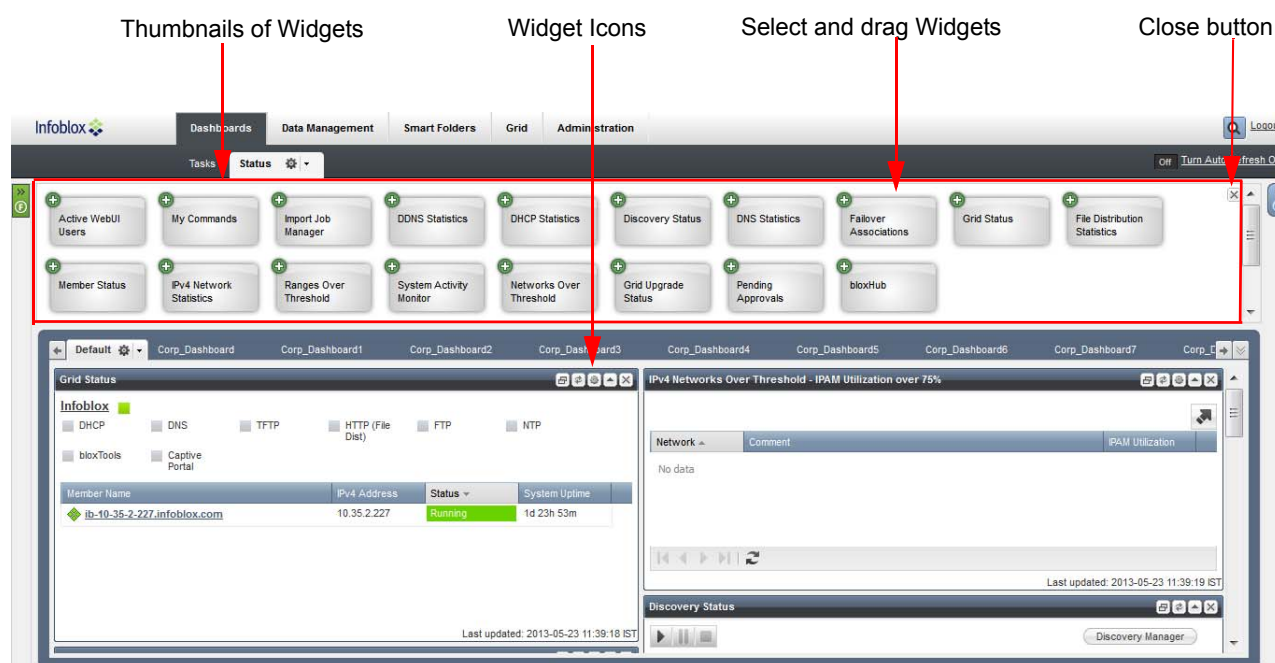
After you add a widget to the dashboard, you can configure it to provide relevant data. You can also copy or move a widget, by selecting and dragging it to its new location on your dashboard. Grid Manager saves your dashboard configuration and displays it the next time you log in.

You can click **Turn Auto Refresh On** at the top of the dashboard to periodically refresh the contents of all widgets. This feature is turned off by default to optimize the performance of Grid Manager.

Widgets have the following icons:

- **Copy/Move:** Click to copy or move the widget from a dashboard to another. For information about how to copy or move, see [Copying or Moving Widgets](#) on page 120.
- **Refresh:** Click to update the content of the widget. Each widget contains a status bar at the bottom that displays the last date and time it was updated.
- **Configure:** Click to hide and show the configuration options of the widget.
- **Toggle:** Click to minimize and restore the widget.
- **Close:** Click to remove the widget from a dashboard.

Figure 2.2 Widgets Panel



Configuring Widget Limit per Dashboard

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

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

Adding Status Dashboards

You can create your own status dashboards and add the widgets that you need. You can configure up to 100 status dashboards at a time. When you create multiple instances of a dashboard, the appliance names each dashboard by adding an incremental suffix to the name of the new dashboard. For example, if you name a new dashboard “Corp_Dashboard” and specify the number of instances as three, then the appliance creates three instances of this new dashboard. In this example, the appliance creates three dashboards: Corp_Dashboard, Corp_Dashboard1, and Corp_Dashboard2. Note that the dashboards you create will not be available to other users. You cannot share dashboards you have created with other users.

To add a new status dashboard:

1. From the **Dashboards** → **Status** tab, click the Configure icon → **Add Dashboard**.
2. In the *Add Dashboard* wizard, complete the following:
 - **Name:** Enter a name for the new dashboard.
 - **Add instances of this new dashboard:** Enter the number of dashboards you want to create. The maximum number of dashboards you can create is 100 at a time.

- **Copy initial content from an existing dashboard:** Select this check box if you want the appliance to copy the contents from an existing status dashboard into the new dashboard. After you select this check box, the appliance displays the list of configured dashboards. Select a dashboard from the list. By default, this check box is not selected.
3. Save the configuration.
The appliance displays all dashboard instances in the **Status** tab.

Using Quick Navigation

You can use the Quick Navigation icon to quickly access a specific dashboard. The appliance provides the Quick Navigation icon at the right corner of the status dashboards, as illustrated in [Figure 2.1](#).

To quickly navigate to a dashboard:

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

Renaming Status Dashboards

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

To rename a dashboard:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Rename Dashboard**.
2. In the *Rename Dashboard* wizard, complete the following:
 - **Select a dashboard:** Select a dashboard from the drop-down list.
 - **Name:** Enter the new name of the dashboard.
3. Do one of the following:
 - Click **Save and Close** to save the new name and close the wizard.
 - Click **Save** to save the new name and continue to rename other dashboards.

To rename a specific dashboard:

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

Copying or Moving Widgets

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

To move or copy a widget:

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

Reordering Status Dashboards

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

To reorder status dashboards:

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

Deleting Status Dashboards

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

To delete multiple dashboards:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Delete Dashboards**.
2. In the *Delete Dashboards* wizard, select the **Dashboard** check box. You can select multiple check boxes for multiple dashboards.
3. Click **Delete**.
4. Click **Yes** in the *confirmation* dialog box.

To delete a specific dashboard:

1. From the **Dashboards** -> **Status** tab -> select the <*Status Dashboard*> tab.
2. Click the Configure icon -> **Delete Dashboard**.
3. In the *Delete Confirmation* dialog box, click **Yes**.




Grid Status

The *Grid Status* widget provides status information about the Grid members and services. Add the *Grid Status* widget to your Dashboard to monitor the Grid status.





You can configure the *Grid Status* widget to display information about all Grid members or only Grid members that have service errors. To modify the *Grid Status* widget, click the Configure icon and select one of the following:

- **Show all Grid members** (this is the default)
- **Only show members with service warnings or errors**

In the upper section of the widget, Grid Manager displays the overall status of the Grid. The Grid status represents the status of the most critical member in the Grid. When all Grid members are running properly, the overall Grid status is green. When one of the members has operational issues, the overall Grid status is red. The status icon can be one of the following:

Icon	Color	Meaning
	Green	All Grid members are operating normally in a “Running” state.
	Yellow	At least one of the Grid members is connecting or synchronizing with its Grid Master.
	Red	At least one of the Grid members does not have a Grid license, is offline, upgrading, downgrading, or shutting down.

This section also displays the overall operational status of the DNS, DHCP, NTP, FTP, TFTP, HTTP (File Distribution), bloxTools, Captive Portal, DNS Accelerator usage, and Reporting services that are currently running on the Grid. The DNS Accelerator usage feature is only available in the IB-4030 appliance. The status icon can be one of the following:

	Green	The enabled service is running properly on one or more Grid members.
	Yellow	At least one of the Grid members is having issues with the enabled service.
	Red	The enabled service is not running properly on at least one of the members. (A red status icon can also appear temporarily when the service is enabled and begins running, but the monitoring mechanism has not yet notified Grid Manager.)
	Gray	The service is not configured or is disabled on at least one Grid member.

The *Grid Status* widget also displays the following information in the member table:

- **Member Name:** The name of the member.
- **IP Address:** The IP address of the member.
- **Status:** The current status of the member.
- **System Uptime:** The duration of time (days, hours, and minutes) that the Grid member has been up and running.

When you select **Only show members with service warnings or errors**, the widget displays only the members that have service errors. The widget does not display any data in the member table if all the services on all members are running properly.

You can click a member link to monitor the detailed status of the selected member. Grid Manager displays the **Grid** tab -> **Member** tab. For information, see [Member Status](#) on page 1004.

Grid Upgrade Status

The *Grid Upgrade Status* widget provides upgrade status of the Grid Master and members. Add the *Grid Upgrade Status* widget to your Dashboard to monitor the upgrade status of the Grid and its members.

The *Grid Upgrade Status* widget displays the following information:

- **Upgrade Status:** The current upgrade status of the Grid. This can be **Running**, **Paused**, **Canceled**, or **Inactive**.
- **Grid Member Upgrade Process Status:** The pie chart shows the number of members that are still processing the upgrade, members that have completed the upgrade, and members that are waiting for the upgrade to happen.
- **Detailed Upgrade Status:** Click this link to access the **Grid** tab -> **Upgrade** tab to see detailed information about the upgrade.

The table on the right shows a summary of the upgrade status of the upgrade groups. It displays the following information:

- **Group:** The name of the upgrade group.
- **Date/Time:** The date and time when the upgrade started on this upgrade group. Note that the time zone is the time zone of the first member in the upgrade group.
- **Completed:** Indicates whether the upgrade is complete or not.

Member Status (System Status)

The *Member Status* widget provides status information about the system resources and services of a Grid member, including the reporting server. The *System Status* widget provides the operational status about an independent appliance. Add a *Member Status* widget to your Dashboard for each Grid member that you want to monitor. The widget always displays the services that a Grid member is running. You can then configure it to display additional information and specify how the information is displayed.





You can modify the *Member Status* or the *System Status* widget by clicking the Configure icon. If you have an independent appliance, you can only configure some of the following:

- For *Member Status* widget only: Click **Select Member** to select a Grid member for display. When you select the reporting server, the widget displays reporting related information, such as reporting usage and reporting warning count.
- Select the information you want to display:
 - **Show Role:** For *Member Status* widget only. Click to display whether the appliance is a Grid Master, Grid Master candidate, or Grid member. An independent appliance does not have a Grid license installed.
 - **Show Hardware Type:** Click to display the appliance hardware model.
 - **Show HA Status:** Click to display whether the appliance is part of an HA pair. It displays one of the following:
 - **Standalone:** The Grid member is an independent appliance.
 - **HA OK:** The Grid member is part of an HA pair that is functioning properly.
 - **HA Broken:** The appliance is part of an HA pair that is not operating properly. You can check the logs to determine the problem.
 - **Show System Uptime:** Click to display the duration of time (days, hours, and minutes) that the Grid member has been up and running.
- **Statistics:** Select the data that you want to display and its format:
 - **CPU:** Click to display the percentage of CPU that is in use. Select either **Dial** or **Bar** for the display format.
 - **Memory:** Click to display the current percentage of memory that is in use. Select either **Dial** or **Bar** for the display format.
 - **Database:** Click to display the percentage of the database that is in use. Select either **Pie** or **Bar** for the display format.
 - **Disk:** Click to display the percentage of the data partition on the hard disk drive in use. Select either **Pie** or **Bar** for the display format.
 - **System Temperature:** Click to display the system temperature. Depending on the hardware model, the system temperature may not be available. Select to display the temperature in either **Celsius** or **Fahrenheit**.
 - **CPU Temperature:** Click to display the CPU temperature. Depending on the hardware model, the CPU temperature may not be available. Select to display the temperature in either **Celsius** or **Fahrenheit**.
 - **DNS Accelerator Usage:** Click to display the percentage of DNS Cache Acceleration usage, if available. The Member Status is yellow, if the DNS cache acceleration utilization exceeds a predefined threshold. This feature is only available in the IB-4030 appliance.

Click the Configuration icon again to hide the configuration panel after you complete the modification.

Grid Manager displays the hostname of the appliance at the top of the widget. You can click the name link to view detailed information about the appliance. The widget also displays the upgrade status if the member is currently in the process of an upgrade. If the member is scheduled for an upgrade, the **Scheduled for upgrade** link appears. You can click this link to access the **Grid** tab -> **Upgrade** tab to view more details about the date and time of the scheduled upgrade.

The widget also displays the service status of the following: FTP, TFTP, HTTP (File Distribution), DNS, DHCP, NTP, bloxTools, Captive Portal, DNS Accelerator, IF-MAP (for IF-MAP server only), and Reporting in the Services section. The service status can be one of the following:

Icon	Color	Meaning
	Green	The service is enabled and running properly.
	Yellow	The service is enabled, but there may be some issues that require attention.
	Red	The service is enabled, but it is not running properly or is out of synchronization. (A red status icon can also appear temporarily when a service is enabled and begins running, but the monitoring mechanism has not yet notified the GUI engine.)
	Gray	The service is not configured or is disabled.

The widget also displays the statistics you specified, such as CPU usage, memory and database usage, in the format you selected.

When you select the reporting server, you can also see the following information:

- **Reporting Usage:** Displays the daily consumption rate for the reporting service.
- **Reporting Warning Count:** When the data usage on the reporting server approaches or reaches the daily maximum limit, the appliance sends an SNMP trap and email notification, if configured. When you receive five (5) violation notifications in a rolling period of 30 days, you cannot view reports or configure reporting related functions. You must then contact Infoblox Technical Support to resolve the issue.

For more information about reporting, see [Infoblox Reporting Solution](#) on page 1113.

DNS Statistics

The *DNS Statistics* widget provides statistics for a member or for a zone. The zone statistics are cumulative, collected from all the members that are authoritative servers for zones or are hosting stub zones. The widget displays the totals for each type of DNS response as well as a line graph that tracks the responses per second.

You can add a *DNS Statistics* widget to your Dashboard for each zone or member DNS server on the Grid.

To configure the *DNS Statistics* widget, click the Configure icon and do the following:

- Click **Select Member**. In the *Member Selector* dialog box, choose a Grid member to display statistics for all its stub zones and authoritative zones.

or

- Click **Select Zone**. In the *Zone Selector* dialog box, choose a DNS zone to display statistics for that zone only.

The widget displays only the option that you selected on your subsequent logins. For example, if you clicked **Select Member**, the widget displays the **Select Member** option only, and not the **Select Zone** option, when you log in again.

- **Graph Configuration:** Select which DNS messages you want to track in the **Responses per Second** graph.
 - **Success:** The number of successful queries.
 - **NXDOMAIN:** The number of queries for domain names that did not exist in the database.
 - **Referral:** The number of queries that became referrals.
 - **NXRRSET:** The number of queries for domain names that did not have the requested records.
 - **Failure:** The number of queries that failed due to reasons other than nonexistent domain names or records in a domain.
 - **Recursion:** The number of recursive queries for which the name server sent queries to other name servers.

The widget displays the following information:

- **DNS Responses** tab: Displays a pie chart and the total number of each type of message. It also displays the total number of full and incremental zone transfers that the Grid member performed.
- **Responses per Second** tab: Displays a line graph that tracks the DNS responses received per second, within an hour. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

Ranges Over Threshold

The *Ranges Over Threshold* widget enables you to monitor IPv4 DHCP range usage from your Dashboard. It lists the IPv4 ranges that are allocated above a specified threshold and thus may warrant your attention. The default threshold is 75%. For information, see [Configuring Thresholds for DHCP Ranges](#) on page 807. Note that the appliance highlights disabled IPv4 ranges in gray.

The widget displays the IPv4 ranges with utilization percentages that surpass the threshold.

To configure the *Ranges Over Threshold* widget, click the Configure icon and do the following:

- **Network View:** Select a network view in which you want to monitor the IPv4 ranges. This field is displayed only when you have more than one network view.
- **Threshold:** Enter a new threshold value. The default is 75%.

In addition, you can do the following:

- Click the Export button to export the list of IPv4 ranges that surpass the threshold to a file in CSV format.
- Click the Refresh button to refresh the data in the list.

IPv4 Failover Associations Status

The *IPv4 Failover Associations Status* widget enables you to monitor the status of the failover associations from your Dashboard. It lists all the failover associations in the Grid and displays their names and status. The widget also displays the primary and secondary servers in the association. When you click a failover association link or a status link, Grid Manager displays the Failover Association section where you can get detailed information about the failover association. For information, see [Monitoring Failover Associations](#) on page 888.

In addition, you can do the following:

- Click the Export button to export the list of failover associations to a file in CSV format.
- Click the Refresh button to refresh the data in the list.

DHCP Statistics

The *DHCP Statistics* widget displays statistics about the different types of DHCP messages that a Grid member sends and receives. The widget displays the totals for each type of DHCP message as well as a line graph that tracks the messages per second.

You can add a *DHCP Statistics* widget to your Dashboard for each member DHCP server in the Grid. If the DHCP service is not enabled or is offline, the widget displays a message indicating that the DHCP statistic are not available.

To configure the *DHCP Statistics* widget, click the Configure icon and do the following:

- **Protocol:** Select either **IPv4** or **IPv6**.
- Click **Select Member**. In the *Member Selector* dialog box, select a Grid member from the list.
- **Graph Configuration:** This section lists IPv4 or IPv6 messages, depending on the protocol you selected.
- Select which IPv4 messages you want to track in the **Messages per Second** graph.
 - **Discovers:** The number of DHCPDISCOVER messages that the Grid member received from DHCP clients. A DHCP client broadcasts a DHCPDISCOVER message to obtain an IP address.
 - **Offers:** The number of DHCPOFFER messages that the Grid member sent to DHCP clients. If the Grid member has an IP address that it can allocate to the DHCP client that sent the DHCPDISCOVER message, the Grid member responds with a DHCPOFFER message that includes the IP address and configuration information.
 - **Requests:** The number of DHCPREQUEST messages that the Grid member received from DHCP clients. A DHCP client sends DHCPREQUEST messages when it selects a lease, connects to the network, and if it renews the lease.
 - **Acks:** The number of DHCPACK messages that the Grid member sent to DHCP clients. When the Grid member receives a DHCPREQUEST message, it responds with a DHCPACK message to confirm the IP address selected by the DHCP client.
 - **Nacks:** The number of DHCPNACK messages that the Grid member sent to DHCP clients. The Grid member sends a DHCPNACK message when a DHCP client requests an IP address that is not valid for the network.
 - **Declines:** The number of DHCPDECLINE messages that the Grid member received. A DHCP client sends a DHCPDECLINE message to a DHCP server when it discovers that the IP address offered by a DHCP server is already in use.
 - **Inform:** The number of DHCPINFORM messages that the Grid member received. A client that did not receive its IP address from the DHCP server can send it a DHCPINFORM message to retrieve configuration parameters, such as the IP addresses of DNS servers in the network.
 - **Releases:** The number of DHCPRELEASE messages that the Grid member received. A DHCP client sends a DHCPRELEASE message when it terminates its lease and releases its IP address.

Select which IPv6 messages you want to track in the **Messages per Second** graph.

- **Declines:** The number of Decline messages that the Grid member received. A DHCP client sends a Decline message to a DHCP server when it discovers that the IP address offered by a DHCP server is already in use.
- **Renews:** The number of Renew messages that the Grid member received. A DHCP client sends a Renew message to a DHCP server to extend the lifetimes on the leases granted by the DHCP server and to update other properties.
- **Information Requests:** The number of Information-Request messages that the Grid member received. A client sends an Information-Request message to retrieve configuration parameters, such as the IP addresses of DNS servers in the network.
- **Solicits:** The number of Solicit messages that the Grid member received, including Solicit messages embedded in Relay-Forward messages. A DHCP client sends a Solicit message to locate DHCP servers.
- **Requests:** The number of Request messages that the Grid member received. A DHCP client sends a Request message to request one or more IP addresses and configuration parameters from a DHCP server.
- **Rebinds:** The number of Rebind messages that the Grid member received. A DHCP client sends a Rebind message to extend the lifetime of its lease and to update configuration parameters.
- **Releases:** The number of Release messages that the Grid member received. A DHCP client sends a Release message when it terminates its lease and releases its IP address.

- **Advertises:** The number of Advertise messages that the Grid member sent. When a DHCP server receives a Solicit message, it can respond with an Advertise message to indicate that the server is available for DHCP service.
- **Replies:** The number of Reply messages that the Grid member sent. A DHCP server sends a Reply message that includes IP addresses and configuration parameters when it responds to Solicit, Request, Renew or Rebind message. It sends a Reply message with configuration parameters only when it responds to an Information-Request message.

The widget displays the following information:

- **DHCP Messages** tab: Displays a pie chart and the totals for each type of DHCP message. It also displays the number of Deferred Updates, which are DDNS update requests which are deferred because the DNS primary was not reachable when the update was first attempted.
- **Messages per Second** tab: Displays a line graph that tracks the DHCP messages that were sent and received per second, within an hour. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

Network Statistics

The *Network Statistics* widget provides information about IP address usage in an IPv4 network. You can monitor several networks simultaneously to view the distribution of address resources. Such information can indicate if there is a sufficient number of available addresses in each network. It can also provide information about the distribution of address resources, indicating if there are too many unused addresses in one network while all the addresses in another are in use.

Add a *Network Statistics* widget to your Dashboard for each network that you want to monitor. You can monitor IPv4 networks only.

To configure the *Network Statistics* widget, click the Configure icon and do the following:

- Select one of the following chart types:
 - **Pie**
 - **Bar**
- Click **Select Network**. In the *Network Selector* dialog box, choose a network from the list and click **Select**.
Note that if multiple network views were previously configured, Grid Manager displays the default network view. You can choose another network view from the drop-down list, and then select a network.

The *Network Statistics* widget displays the following information about the selected network:

- **IPAM Utilization:** When you define a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network.

When you define a network container that contains subnets, this is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. For example, when you define a /16 network and then 64 /24 networks underneath it, the /16 network container is considered 25% utilized even when none of the IP addresses in the /24 networks is in use.

You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.

- **Unmanaged:** The number of discovered IP addresses that do not have corresponding records on the appliance, such as A records, PTR records, fixed address records, host records, or leases. To obtain this data, you must run a discovery process on the network first.

- **Conflicts:** The number of IP addresses that have either a MAC address conflict or a DHCP range conflict. To obtain this data, you must run a discovery process on the network first. A discovered host has a MAC address conflict when its MAC address is different from that specified in its fixed address, DHCP lease, or host record. A discovered host has a DHCP range conflict when it is part of a DHCP range, but it does not have a matching fixed address or DHCP lease, and it is not part of an exclusion range.

IPv4 Networks Over Threshold

The *IPv4 Networks Over Threshold* widget enables you to monitor IPv4 network and IP address usage from your Dashboard. It lists the IPv4 networks that are allocated above a specified threshold and thus might warrant your attention. The default threshold is 75%.

For network containers, the threshold is the percentage of IP address space that has been allocated. For subnets, it is the percentage of used addresses, except the broadcast and network addresses. The widget displays the network containers and subnets with utilization percentages that surpass the threshold.

To configure the *Networks Over Threshold* widget, click the Configure icon, and then complete the following:

- **Network View:** This field appears only if you have more than one network view. Select the network view in which you want to monitor the threshold.
- **Threshold:** Enter a new threshold value. The default is 75%.
- **Type:** Select **IPAM Utilization** or **IPv4 DHCP Utilization**. For information, see [Managing IPv4 DHCP Data](#) on page 841.

In addition, you can do the following:

- Click the Export button to export the list of networks that surpass the threshold to a file in CSV format.
- Click the Refresh button to refresh the data in the list.

Discovery Status

The appliance can run an IP discovery to detect and obtain information about active hosts in specified networks. It can also run a VM discovery to detect virtual entities on VMware vSphere servers. For information about the discovery process, see [Chapter 13, Network Discovery](#), on page 493.

You can add the *Discovery Status* widget to your Dashboard. From this widget, you can access Discovery Manager and configure parameters for a discovery. You can do the following from the widget:

- Start a discovery immediately. For more information about immediate discovery, see [Starting a Discovery Immediately](#) on page 505.
- Schedule a discovery for a later date and time. For more information about discovery, see [Scheduling a Discovery](#) on page 506.
- Configure a recurring discovery. For more information about recurring discovery, see [Configuring a Recurring Discovery](#) on page 506.
- Click the Start button to start a discovery process.
- Click the Pause button to temporarily pause the process.
- Click the Stop button to stop the process.

This widget displays the status of discovery tasks. If there are no active discovery tasks, the widget displays the discovery results of the previous tasks. For information about starting and scheduling a discovery task, see [Guidelines for Starting and Scheduling a Discovery](#) on page 505.

After you start a discovery, the *Discovery Status* widget displays a status bar that indicates the discovery is in progress. It also tracks the number of networks in an IP discovery and the number of vSphere servers and virtual machines in a VM discovery. You can click the Refresh icon to update the discovery status.

The widget displays the following information about the discovery process:

- **Current Status:** If a discovery is in progress, this field displays its current status. Otherwise, it displays the date and time of the last discovery.
- **Last Action:** Displays the last operation and the admin who initiated it.

- **IPv4 Device Discovery:** Displays the total number of IPv4 networks and the IPv4 network and IP address range on which the IP discovery is currently running. You can click **Refresh** to update this information.
- **VM Discovery:** Displays the total number of vSphere servers, the server on which the VM discovery is currently running, and the server IP address or FQDN. You can click **Refresh** to update this information.

The *Discovery Status* widget also displays the following information about the last discovery:

- **Discovered:** The total number of active hosts in the network.
- **Managed:** The number of discovered IP addresses that are managed by the NIOS appliance. These IP addresses have an A record, PTR record, fixed address record, host record, lease, or are within a configured DHCP range.
- **Unmanaged:** The number of discovered IP addresses that do not have corresponding records on the appliance, such as A records, PTR records, fixed address records, host records, or leases.
- **Conflicts:** The number of discovered hosts that have a MAC address conflict or are part of a configured DHCP range, but do not have a fixed address or lease record and are not part of an exclusion range.

Advanced Discovery Status

With the correct licensing, dedicated NIOS appliances operating as Grid members can perform infrastructure device discovery. NIOS appliances with the Discovery license operate primarily for discovery tasks and do not perform core DNS or DHCP network functions. Discovery appliances, called Probes, collect all network device data and compile it into a database. A separate NIOS appliance, called a Consolidator, aggregates the collected device information from the Probes and synchronizes with the Infoblox Grid Master.

For more information about discovery and its features and requirements, see the chapter *Network Insight* on page 517 and its associated sections.

The Advanced Discovery Status widget provides several basic counts describing the general state of device discovery within the Grid, and for networks outside the Grid being inventoried by the NIOS appliances designated for discovery. The widget divides counters into two categories: **Networks** and **Assets**. Network counters refer to counts of managed and unmanaged networks discovered by Probe appliances. Asset counters refer to counts of specific types of network devices, termed Assets, which are comprised of end hosts, enterprise servers, enterprise printers, and any other enterprise asset that exists in an end-user network segment. The widget counters include:

In the **Networks** category:

- **Discovered:** The total number of networks discovered by Probe appliances.
- **Managed:** The number of discovered networks that are currently managed by the NIOS Grid. These IP networks have been converted from Unmanaged status to Managed status.
- **Unmanaged:** The number of discovered networks that are counted as Unmanaged by the NIOS Grid Master. After a network is discovered and catalogued by a Probe appliance, its default state as a network is Unmanaged.

In the **Assets** category:

- **Discovered:** The total number of Assets discovered by Probe appliances.
- **Managed:** The number of discovered assets that are currently managed by the NIOS Grid. These devices have been converted from Unmanaged status to Managed status.
- **Unmanaged:** The number of IPs with discovered data that are counted as Unmanaged by the NIOS Grid Master, and have not been converted into a Host or a Fixed IP Address. After an Asset is discovered and catalogued by a Probe appliance, its default state is Unmanaged.
- **Conflicts:** The number of discovered assets that have a MAC address conflict or are part of a configured DHCP range, but do not have a fixed address or lease record and are not part of an exclusion range.

My Commands

The *My Commands* widget provides easy access to commands that you frequently use, so you can perform your tasks without leaving the Dashboard. You can add one *My Commands* widget to your Dashboard.

To configure the *My Commands* widget, click the Configure icon and do the following:

- Select a command from the **Available** list and click the › arrow to move it to the **Selected** list. You can always toggle the commands between the two lists. Select multiple commands by using SHIFT-click and CTRL-click.

DDNS Statistics

The *DDNS Statistics* widget provides information about the dynamic DNS (DDNS) updates that occur on the DNS service of a selected Grid member. The widget displays the total number of DDNS updates that succeeded, failed, and that were rejected. It also displays a line graph that tracks the status of the DDNS updates per second.

You can add a *DDNS Statistics* widget to your Dashboard for each DNS server on the Grid that accepts dynamic DNS updates.

To configure the *DDNS Statistics* widget, click the Configure icon and do the following:

- Click **Select Member**. In the *Member Selector* dialog box, select a Grid member from the list.
- **Graph Configuration**: Select which updates you want to track in the **Updates per Second** graph:
 - **Success**: The number of DDNS update requests that succeeded.
 - **Prerequisite Reject**: The number of DDNS update requests that were rejected because the prerequisite conditions specified in the request were not met.
 - **Reject**: The number of DDNS update requests that were rejected by the DNS service.
 - **Failure**: The number of DDNS update requests that failed.

The widget displays the following information:

- **DDNS Updates** tab: Displays totals for each type of update.
- **Updates per Second** tab: Displays a line graph that tracks the status of the DDNS updates. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

System Activity Monitor

The *System Activity Monitor* widget provides information about the following resources on the selected Grid member: CPU, system memory, NIC usage, and information about VLAN interfaces. By default, the widget displays the system activity of the Grid Master. You can add a *System Activity Monitor* widget to your Dashboard for each Grid member whose resources you want to monitor.

To configure the *System Activity Monitor* widget, click the Configure icon and select a Grid member and the resources that you want to track:

- Click **Select Member**. In the *Member Selector* dialog box and select a Grid member from the list.
- **CPU**: Select which type of CPU usage you want to track:
 - **User**: The CPU usage of user applications, such as programs and libraries.
 - **System**: The CPU usage of the kernel and drivers.
 - **Idle**: The percentage of CPU that is not in use.
- **System Memory**: Select which portion of the system memory you want to track:
 - **Real Memory Used**: The physical RAM usage.
 - **Swap Used**: The swap area usage. The swap area is the disk area that temporarily holds a process memory image.
- **NIC Usage**: Select how you want to measure network traffic:
 - **Bytes**: Reports the number of bytes.
 - **Packets**: Reports the number of packets.
- **NIC Settings**: Select the port on which you want to measure network traffic. If you have configured VLANs, Grid Manager displays them in the format LAN1 nnnn or LAN2 nnnn, where nnnn represents the associated VLAN ID. For example, a VLAN configured on LAN1 can be displayed as LAN1 297 and a LAN2 VLAN can be LAN2 21. For more information about VLANs, see [About Virtual LANs](#) on page 346.

The *System Activity Monitor* widget displays a tab for each resource: **CPU**, **System Memory**, and **NIC Usage**.

Each tab contains a line graph that tracks the resource utilization per second. The graph in the **CPU** tab tracks the percentage of CPU usage. The graph in the **System Memory** tab tracks the memory utilization percentage. The graph in the **NIC Usage** tab tracks either bytes or packets per second.

The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

File Distribution Statistics

The *File Distribution Statistics* widget enables you to monitor the status of file distributions services from the Dashboard. The widget provides an overall status of file distribution on all members in the Grid. It also displays the file system utilization for the file distribution subsystem.

The service status displays one of the following:

- **OK:** All file distribution services are running properly.
- **Stopped:** All file distribution services are stopped.
- **Warning:** The file distribution services are not running properly.
- **Error:** The file distribution services encounter an error.

You can click the link to view detailed information about the file distribution services. Grid Manager displays the Members tab in the File Distribution tab.

To configure the *File Distribution Statistics* widget, click the Configure icon and select one of the following chart types:

- **Pie**
- **Bar**

The *File Distribution Statistics* widget displays the following information:

- **File System Utilization:** The percentage of utilization of the overall allocated file distribution subsystem space on all members. You can use this information to verify if there is sufficient space for file distribution in the Grid.

Active WebUI Users

The *Active WebUI Users* widget provides information about the users who are logged in to Grid Manager or System Manager. It does not include users who are using the Infoblox API or are logged in to the serial console.

You can add only one *Active WebUI Users* widget to the Dashboard. You must have a superuser account to add this widget to the Dashboard.

It displays the following information about each user:

- **User ID:** The user name.
- **Source Address:** The IP address of the management station the user used to connect to Grid Manager.
- **Logged In Since:** The date and time the user logged in.
- **Idle Time:** The number of minutes the user has not had any activity on Grid Manager.
- **User Agent:** The system used to access Grid Manager, such as the browser version and platform information.

You can sort the columns and hide or display each one. You can also export the list to a .csv file.

Microsoft Servers Status Widget





The *Microsoft Servers Status* widget displays the operational status of each Microsoft server managed by the Grid. Grid Manager displays this widget only when at least one member in the Grid has a Microsoft management license. You can configure this widget to display the status of all Microsoft servers or only those with warnings and errors. To modify the *Microsoft Servers Status* widget, click the Configure icon and select one of the following:

- **Show all Microsoft servers**
- **Only show servers with service warnings or errors**

The *Microsoft Servers Status* widget displays the following information about each Microsoft server:





- **Server Name:** The hostname of the Microsoft server.
- **IP Address:** The IP address of the Microsoft server.
- **Status:** The connection status of the Microsoft server.
 - **OK:** The Grid member is connected to the Microsoft server.
 - **Connecting:** The Grid member is connecting to the Microsoft server.
 - **Error:** The Grid member tried to connect to the Microsoft server, but failed. You can check the syslog for any messages.
 - **Not Available:** The Microsoft server is disabled. The Grid member does not try to connect to disabled servers.
- **DNS:** The status of the DNS service on the Microsoft server.

The DNS service status can be one of the following:

Icon	Color	Meaning
	Green	The DNS service is functioning properly.
	Red	The DNS service is stopped.
	Yellow	The DNS service is starting or stopping.
	Gray	Management of the Microsoft DNS server is disabled.

- **DHCP:** The status of the DHCP service on the Microsoft server.

The DHCP service status can be one of the following:

Icon	Color	Meaning
	Green	The DHCP service is functioning properly.
	Red	The DHCP service is stopped.
	Yellow	The DHCP service is starting or stopping.
	Gray	Management of the Microsoft DHCP server is disabled.

Import Job Manager

The **Import Job Manager** on the Status Dashboard displays the status of CSV import jobs you have submitted. You can start a file import from **Import Job Manager** and control and monitor it from this widget. You can also launch **Import Job Manager** from the Task Dashboard or the Toolbar. For more information, see [About CSV Import](#) on page 86. You can click the Refresh icon or configure auto refresh to update the status.

The widget displays the following information about the import jobs that were submitted in the past 14 days:

- **User Name:** The admin user who submitted the CSV import. Only superusers can view this column.
- **Status:** The current status of the import job. The status can be one of the following:
 - **RUNNING:** The job is being executed.
 - **PENDING:** The job is in queue for execution.
 - **COMPLETED:** The import is completed. Check the **Message** field for information about the import.

- **UPLOADED:** The data file has been uploaded, but import is not started.
- **STOPPED:** The job has been stopped. You can select the job and restart the import.

Note: After a service restart, all **RUNNING** jobs go into **STOPPED** state; all **PENDING** jobs continue to be queued for execution.

- **Submitted:** The timestamp when the job was submitted.
- **Completed:** The timestamp when the job was completed. This field is blank if the job has not been completed yet.
- **File Name:** The CSV data file name.
- **Message:** This field displays the number of rows of data that has been processed and the number of rows of data the import has detected errors. Depending on the import options, Grid Manager displays the row number at which it stops the import when it encounters an error, or the total number of rows it has processed by skipping over the erroneous data. For example, if there are 100 rows of data and you select “On error: Stop importing,” and there is an error in row 90, the appliance displays **90 of 100 completed, 1 error**. If you select “On error: Skip to the next row and continue,” the appliance displays **100 of 100 completed, 1 error**.

Note: Superusers can view all CSV import jobs and limited-access users can view only the jobs they submitted.

Load Balancer Status

The *Load Balancer Status* widget displays the operational status of GLBs (Global Load Balancers) managed by the Grid. Grid Manager displays this widget only when at least one member in the Grid has a Load Balancer license. You can configure this widget to display the status of all GLBs or only those with warnings and errors. To modify the *Load Balancer Status* widget, click the Configure icon and select one of the following:

- **Show all Load Balancers**
- **Only show servers with service warnings or errors**

The *Load Balancer Status* widget displays the following information about each load balancer:

- **Name:** The name of the load balancer.
- **IP Address or FQDN:** The IP address or FQDN of the load balancer.
- **Version:** The TMOS version that is running on the load balancer.
- **Status:** The connection status of the GLB.
 - **OK:** The Grid member is connected to the GLB.
 - **Unknown:** The Grid member is unable to contact the GLB and cannot retrieve any status details. This can be caused by incorrect IP address, FQDN, username, or password.
 - **Error:** The GLB has a connection error. Click the Detailed Status icon to view detailed information or check the syslog for any error messages.
 - **Warning:** Certain issues, such as Grid member failures or licensing issues, have occurred. Click the Detailed Status icon to view detailed information or check the syslog for messages to determine the reason for the warning.
 - **Disabled:** The load balancer is disabled. The Grid member does not try to connect to disabled GLB.

Pending Approvals

The *Pending Approvals* widget provides information about tasks that are pending your approvals. Add the *Pending Approvals* widget to your Dashboard to monitor tasks that require your approvals.

You can select a task and perform the following:

- Click the Approve icon to approve the task.
- Click the Reject icon to disapprove the task.

You can also click **Task Manager** to access the **Administration** tab -> **Workflow** tab -> **Task Manager** tab.

The *Pending Approvals* widget displays the following information about each task that requires your approval:

- **Task ID:** The ID associated with the task. The appliance assigns an ID to a task in chronological order.
- **Submitter:** The username of the admin who scheduled or submitted the task.
- **Ticket Number:** The reference number entered by the submitter to identify the task. You can enter up to 20 alphanumeric characters.
- **Scheduled Time:** The date, time, and time zone when the task was scheduled for execution.
- **Affected Object:** The name or value of the object that is associated with the task. For example, if the task involves an A record, this field displays the domain name of the record. If it is a fixed address, it displays the IP address of the fixed address.
- **Object Type:** The object type. For example, the appliance can display A Record or Fixed Address.
- **Action:** The operation the appliance performs in this task. The can be: **Add, Modify, Delete, or Network Discovery.**
- **Submitted Time:** The date, time, and time zone when the task was submitted. You can select this for display. It is not displayed by default.

Response Policy Zone (RPZ) Statistics

The *Response Policy Zone (RPZ) Statistics* widget provides statistical information about RPZ hits. This widget contains the following tabs: *RPZ Recent Hits*, *Trend* and *Health*. Note that you must select a member to view data in the corresponding tabs. In this widget, you can select RPZ rules for which you want to view details. You can do the following in this widget:

- Click **Select Member**. In the *Member Selector* dialog box, choose a Grid member to view the RPZ hits, or statistics, or RPZ zones and their last updated date and time.
- Select a graph configuration, **Client Hits**, **Passthru Hits**, **Blocked Hits**, or **Substituted Hits**, to view details of a specific RPZ rule. You can select either one or all the available graph configurations. Note that **Client Hits** is displayed only when the graph type is **Line Diagram**.
- Select a graph type, **Stacked Diagram** or **Line Diagram**, to display data in the required diagrammatic format. This option is enabled only when you click the *Trend* tab and disabled when you click the *RPZ Recent Hits* or *Health* tabs. For more information, see *Trend* on page 135.
- Click **View Syslog** to view the last 20 RPZ events that are logged in the syslog. For more information, see *Previewing the Syslog* on page 136.
- Click the *RPZ Recent Hits* tab to view information about the latest five RPZ hits with unique client addresses. For more information, see *RPZ Recent Hits* on page 134.
- Click the *Trend* tab to view RPZ hit statistics on the selected member. For more information, see *Trend* on page 135.
- Click the *Health* tab to view information about RPZ zones on the selected member and their last updated times. For more information, see *Health* on page 136.

Note that you must install the RPZ license and enable **RPZ logging** to access this widget. For more information about installing licenses and enabling RPZ logging, see *License Requirements and Admin Permissions* on page 1235 and *Setting DNS Logging Categories* on page 1015.

RPZ Recent Hits

The *RPZ Recent Hits* tab displays the data that is collected from the most recent hits of five unique clients, identified by their IP addresses, during the last 24 hours. NIOS retrieves this data from the syslog. This tab does not display any data when there are no syslog messages or if RPZ logging is disabled. NIOS displays an error message if RPZ logging is disabled. For more information about enabling RPZ logging, see *Setting DNS Logging Categories* on page 1015.

Grid Manager retrieves recent hits from the selected member. If a member has an RPZ license installed, then NIOS will parse the syslog every 60 seconds to collect the data. NIOS parses the generated data to identify the five most recent hits. It searches for these fields in the syslog message: CEF: data string(RPZ syslog) and src fields.

The NIOS appliance remembers the start and end time of previously searched operations to optimize the recent hits data collection, so that the same data is not searched again. Note that when the same client makes repeated queries in the last 24 hours, then there might be less than five unique client hits. You cannot sort or filter values in this tab.

This tab displays the following information:

- **Client IP Address:** IP address of the client that made the recent hits.
- **Requested FQDN:** The domain name or IP address that triggered the RPZ rule. For example, consider an RPZ rule test.com.rpz.com, which queries for test.com. In this example, test.com is the requested FQDN.
- **RPZ Entry:** The RPZ rule that queried a domain name or an IP address. In the above example, test.com.rpz.com is the RPZ rule.
- **Timestamp:** The date and time when the hit occurred.

Consider an example in which you query an RPZ zone and the NIOS appliance logs the following message in the syslog:

```
CEF:0|Infoblox|NIOS|6.9.0-219291|RPZ-QNAME|NODATA|4|app=DNS dst=10.35.101.14
src=10.36.0.251 spt=44460 view=_default qtype=A msg="rpz QNAME NODATA rewrite w18.vg [A]
via w18.vg.fireeye.com"
```

This tab displays information in the corresponding fields as follows:

Fields	Description
Client IP Address	Data is retrieved from the src field. Example: 10.36.0.251
Requested FQDN	It is retrieved from the data between the rewrite and [A] via fields. Example: w18.vg.
RPZ Entry	It is retrieved from the data after the via in msg field. Example: w18.vg.fireeye.com
Timestamp	This is listed in the syslog.

You can export data displayed in this tab by clicking the *Export* icon. For more information, see [Exporting Displayed Data](#) on page 91.

Trend

The *Trend* tab displays statistics of RPZ hits on a member during the last 60 minutes. You can use a stacked graph or a line graph to view the hits. DNS service generates RPZ statistics for the selected member. Each of the RPZ policy is represented with a different color. This tab displays the following information:

- **Client Hits:** Total number of queries that triggered an RPZ policy. Note that this option is not displayed when you choose **Stacked Diagram**, but displayed only when you choose **Line Diagram**.
- **Passthru Hits:** Total number of queries that triggered a **Passthru** RPZ rule. For more information about passthru rules, see [Managing Passthru Rules](#) on page 1239.
- **Blocked Hits:** Total number of queries that triggered a **Block (No Data)** or **Block (No Such Domain)** RPZ rule. For more information, see [Managing Block \(No Data\) Rules](#) on page 1241 or [Managing Block \(No Such Domain\) Rules](#) on page 1240 respectively.
- **Substituted Hits:** Total number of queries that triggered a **Substitute (Domain Name)** or **Substitute (Record)** RPZ rule. For more information, see [Managing Substitute \(Domain Name\) Rules](#) on page 1242 and [Managing Substitute \(Record\) Rules](#) on page 1243.
- **Timestamp:** The graph displays a 24 hours time window.

Note the following about this tab:

- The statistical data in DNS service will be reset when you stop and restart the DNS service or if you force an active DNS service to restart regardless of its state. This results in loss of prior data.
- Using this graph, you can view the timestamp of statistics collection.

Health

The *Health* tab displays information of RPZ zones on a member and their last updated date and time. This data is retrieved directly from the database. Note that you cannot sort or filter values in this tab. You can export the data displayed in this tab by clicking the *Export* icon. For more information, see [Exporting Displayed Data](#) on page 91.

Previewing the Syslog

You can view the RPZ events that are logged in the syslog for a selected Grid member. Note that the preview displays only the last 20 RPZ events from the syslog. This wizard displays the following information:





- **Timestamp:** The date and time when the hit occurred.
- **Facility:** The location on the syslog server sorting the log message.
- **Level:** The severity of the message. This can be **ALERT**, **CRITICAL**, **DEBUG**, **EMERGENCY**, **ERROR**, **INFO**, **NOTICE**, or **WARNING**.
- **Server:** The name of the server that logged this message, plus the process ID.
- **Message:** Detailed information about the RPZ query.

You can click the **Go to Syslog Viewer** link to view the RPZ events that are logged in the syslog. NIOS displays all the RPZ events that are logged in the syslog for the selected member and the **Quick Filter** is set to **RPZ Incident Logs** by default. For more information, see [Viewing RPZ in the Syslog](#) on page 1266.

Infoblox Community

The *Infoblox Community* widget displays the latest news from Infoblox. It provides links to video clips that show you how to perform certain tasks, such as how to prepare for IPAM Express and how to add a network. You can click available links in the widget to get more information about Infoblox products and solutions.

Note that content in the Infoblox Community widget may not be displayed in certain versions of Mozilla FireFox, Google Chrome, and Microsoft Internet Explorer due to restrictions these browsers use to block certain secure data. Follow these steps to unblock the *Infoblox Community* widget and view data in your respective browser:

- **Mozilla FireFox:** Click the *Shield* icon  in the address bar and choose **Disable Protection on This Page** from the drop-down list. The icon in the address bar changes to a warning triangle  and content is displayed in the *Infoblox Community* widget. For more details, refer to information at <https://blog.mozilla.org/tanvi/2013/04/10/mixed-content-blocking-enabled-in-firefox-23/>.
- **Google Chrome:** Click the *Shield* icon  in the address bar and click **Load unsafe script** in the pop-up box. Chrome automatically refreshes the webpage and loads the content in the *Infoblox Community* widget. For more details, refer to information at <https://support.google.com/chrome/answer/1342714?hl=en>.
- **Internet Explorer:** Click the *Compatibility View* icon  adjacent to the address bar. The browser refreshes and the *Security Warning* dialog box is displayed. Click **no** in the dialog box. The **Only Secure content is displayed** pop-up blocker is displayed at the bottom of the browser. Click the **Show all content** button in this pop-up blocker to view the content. For more details, refer to the information at <http://windows.microsoft.com/en-in/internet-explorer/use-compatibility-view#ie=ie-8>.

Mobile Devices Status

The *Mobile Devices* widget provides information about the number of active leases of the DHCP fingerprint devices managed by the Grid. The widget displays a pie chart indicating the number of active leases in percentile for each of the device category. For information about device category, device class, and device type, see [Table 2.2](#). You can click the Refresh icon or configure auto refresh to update the status.

Note: The *Mobile Devices* widget updates its data every 15 minutes. A device might not be displayed in this widget if its lease expires within 15 minutes.

To configure the *Mobile Devices* widget, click the Configure icon and do the following:

- Click **Select Network View**. In the *Network View Selector* dialog box, select a network view from the list and click **OK**.

Note that if multiple network views were previously configured, Grid Manager displays the default network view. You can select another network view from the *Network View Selector* dialog box.

The widget displays the number of active leases for the following device classes:

- Mac OS** - Displays all devices that were detected to be running Mac OS.
- Windows** - Displays all devices that were detected to be running Windows.
- Android Mobile** - Displays Smartphones/PDAs/Tablets that were detected to be running Android.
- Apple Mobile** - Displays Smartphones/PDAs/Tablets that were detected to have Apple in the DHCP fingerprint information.
- No Match** - Displays all devices whose fingerprint information does not match with any of the standard/custom DHCP fingerprint data stored in the appliance. For information about Standard and Custom DHCP Fingerprints, see [Standard and Custom DHCP Fingerprints](#) on page 1033.
- Other** - Displays all devices that belong to a device class other than those listed above.

Table 2.2 List of device types and classes

Category	Device Class	Device Type
Windows	Windows	Microsoft Windows 2000
		Microsoft Windows 2003
		Microsoft Windows 8
		Microsoft Windows Vista/7 or Server 2008
		Microsoft Windows XP
Mac OS	Macintosh	Apple Mac OS 9
		Apple Mac OS X, TV (HD)
Apple Mobile	Smartphones/PDAs/Tables	Apple iPod
		Apple iPod, iPhone, iPad or TV (SD)
Android Mobile	Smartphones/PDAs/Tables	Android Phone/Tablet (Generic)
		Android Phone/Tablet (HTC, older devices)
		Android Phone/Tablet (Motorola, older devices)
		Android Phone/Tablet (Sony Ericsson, older devices)
		Android Phone/Tablet (Unknown devices)
		Android Phone/Tablet (Vizio tablet, Others)
		Android Phone/Tablet (newer devices)
		Android tablets (Samsung, Others)
		ZTE N9120 Android

Threat Protection Statistics

The *Threat Protection Statistics* widget displays statistics about Advanced DNS Protection related events by severity. The event statistics are cumulative, collected from all members that supports Advanced DNS Protection. For information about the threat protection feature, see [About Infoblox Advanced DNS Protection](#) on page 1214.

This widget displays the total numbers of events by each severity level as well as a line graph that tracks events by the second.

To configure the *Threat Protection Statistics* widget, click the Configure icon and do the following:

- Click **Select Member**. In the *Member Selector* dialog box, select a Grid member from the list.
- **Graph Configuration**: Select the severity level for the event you want to track in the **Event Count by Severity** graph.
 - **Critical**: The number of critical events.
 - **Major**: The number of major events.
 - **Warning**: The number of warning events.
 - **Informational**: The number of informational events.

The widget displays the following information:

- **All Events per Severity Level** tab: Displays a pie chart and the totals for each type of event severity.
- **Event Count by Severity** tab: Displays a line graph that tracks the security events captured per second. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to view the coordinates of any point in the graph.



Chapter 3 Smart Folders

This chapter explains how to create and use smart folders to organize your core network services data. It includes the following sections:

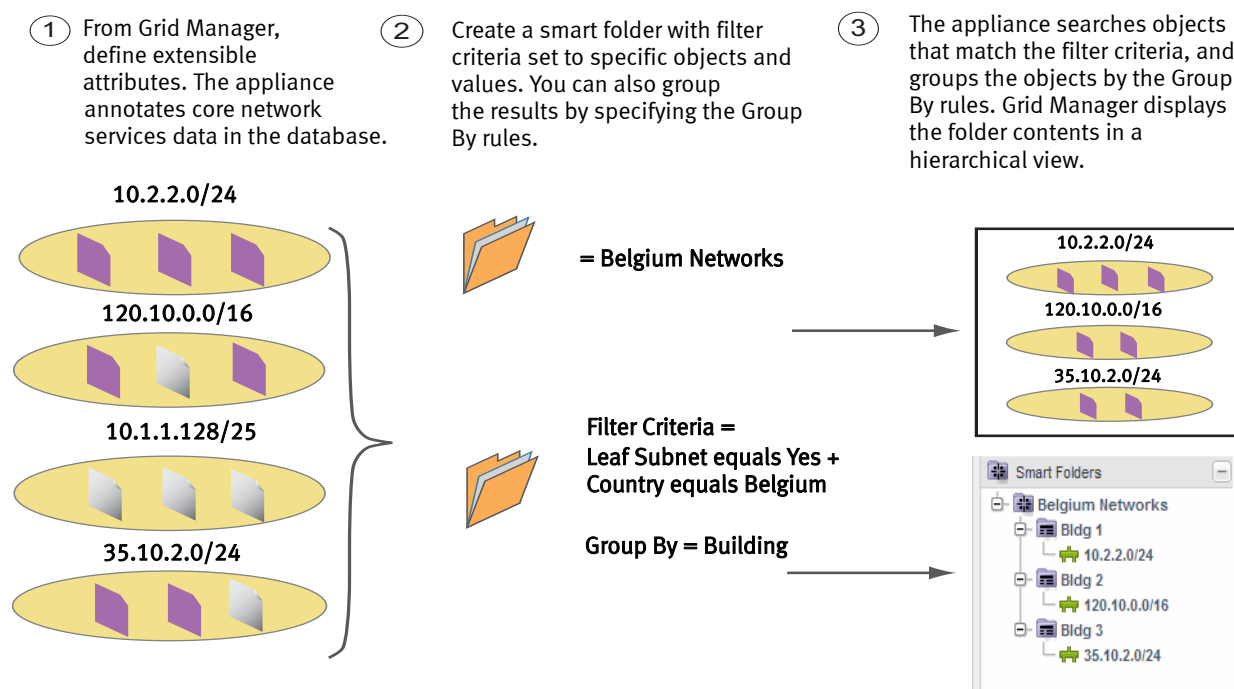
- [*About Smart Folders*](#) on page 140
 - [*Global Smart Folders*](#) on page 141
 - [*My Smart Folders*](#) on page 141
 - [*Predefined Smart Folders*](#) on page 142
- [*Creating Smart Folders*](#) on page 142
- [*Viewing and Modifying Data in Smart Folders*](#) on page 143
- [*Modifying Smart Folders*](#) on page 144
- [*Deleting Smart Folders*](#) on page 144
- [*Saving a Copy of a Smart Folder*](#) on page 145
- [*Printing and Exporting Data in Smart Folders*](#) on page 145

ABOUT SMART FOLDERS

Use smart folders to organize your core network services data. Depending on your administrative roles and business needs, you can filter your data by object types, names, extensible attributes, and discovered data such as conflicts, unmanaged data, or the virtual entity data, and then place the filtered results in a smart folder. You can also group the filtered results by defining up to 10 extensible attributes as the Group By rules. For example, you can create a smart folder that contains all the networks you manage in Belgium, and then group the networks by building number, as illustrated in [Figure 3.1](#).

Once you set up a smart folder, the appliance displays up-to-date information based on your filter and grouping criteria each time you access the folder. You can also view and modify object information in the folder. For information, see [Viewing and Modifying Data in Smart Folders](#) on page 143. With smart folders, you can organize your data in a meaningful way and quickly obtain the information you need to perform specific tasks without searching the entire database.

Figure 3.1 Creating Smart Folders



Before you set up your smart folders, decide how you want to organize your data. You can specify search and Group By criteria to help you group information in a meaningful way. You can also decide whether you want to include objects that do not contain attribute values when you use the Group By criteria to group filtered data by extensible attributes. For information, see [Creating Smart Folders](#) on page 142. Note that a smart folder becomes invalid when you delete an extensible attribute that the folder uses as a filter or Group By criterion. You must redefine the extensible attribute and reconfigure the folder criteria to validate the smart folder.

In Grid Manager, you can create smart folders in both the Global Smart Folders and My Smart Folders panels. In Global Smart Folders, you can create smart folders to which other administrators can create links. Only administrators with superuser accounts can create, edit, and delete global smart folders. For information, see [Global Smart Folders](#) on page 141. You can create personal folders as well as links to global smart folders in My Smart Folders. For information, see [My Smart Folders](#) on page 141.

Each smart folder you create can contain up to 2,000 objects. When the number of objects exceeds 2,000, Grid Manager sorts and displays the first 2,000 objects only. It also displays a warning message at the top of the panel. In this case, you may want to redefine your filter criteria to further refine the filtered data in your smart folders.

To create smart folders, follow these procedures:

1. Determine how you want to organize your core network services data.
2. Identify the fields that you want to use to group networks or define extensible attributes for the data that you want to track. For information about extensible attributes, see [About Extensible Attributes](#) on page 322.

Note: Infoblox strongly recommends that you use **Type** as one of the filter criteria to improve system performance.

3. Create smart folders in either the My Smart Folders or Global Smart Folders panel. For information, see [Creating Smart Folders](#) on page 142.

Global Smart Folders

You can create global smart folders to share among administrators. You must log in as a superuser account to create, edit, and delete global smart folders. All other users have read-only access to global smart folders. You can create as many folders as you need in Global Smart Folders. You can also save a local copy of an existing folder, depending on your administrative permissions. For information, see [Saving a Copy of a Smart Folder](#) on page 145.

Grid Manager displays a list of global smart folders in the list panel.

When you log in as a superuser and mouse over a global smart folder, the following icons appear:

- **Information:** Displays information about the selected smart folder. Information includes comments and filter criteria for the folder. It also displays the Group By rules.
- **Edit:** Click this icon to edit the definition and filter criteria for the smart folder.
- **Create link:** Click this icon to create a link to the smart folder. The link to this folder is placed in My Smart Folders.
- **Delete:** Click this icon to delete the smart folder. This operation does not affect the objects that are in the folder. Only the smart folder is deleted.

My Smart Folders

In My Smart Folders, you can create personal smart folders and links to global smart folders. You can create up to 200 smart folders, including links to global smart folders. When you create links to global smart folders, you can only view information in the folders. However, you can create a local copy of the global smart folder in its current state for editing purposes. Note that when the original global smart folder is updated, information in your local copy is not updated. For information, see [Saving a Copy of a Smart Folder](#) on page 145. When you delete a link to a global smart folder in this tab, only the link is deleted. There is no impact on the information in the original global smart folder.

Grid Manager displays a list of smart folders in the list panel. The same list of smart folders is also displayed in the *Finder* panel. For information, see [Finder Panel](#) on page 59.

When you mouse over a smart folder in the list panel, the following icons appear:

- **Information:** Displays information about the selected smart folder. Information includes comments and filter criteria of the folder. It also displays how you grouped the filtered data.
- **Edit:** Click this icon to edit the definition and filter criteria for the smart folder.
- **Delete:** Click this icon to delete the smart folder. This operation does not affect the objects or networks that are in the folder. Only the smart folder is deleted.

Predefined Smart Folders

The appliance can detect remote clients through their DHCP fingerprints, or through device information discovered through SNMP and other device and network detection protocols. You can use predefined smart folders to view lease history, IP addresses, network infrastructure devices, and other related information for remote clients that contain DHCP fingerprint information related to the following device groups:

- **Smartphone, PDA, Tablet Devices:** Includes all devices that were detected as smartphones, PDAs, and tablets.
- **Microsoft Windows Devices:** Includes all devices that were detected to be running Windows OS.
- **Apple MAC OS Devices:** Includes devices that were detected to be running MAC OS.
- **Conflicts:** Includes hosts detected in the network that have a MAC Address conflict. A discovered host has a MAC address conflict when its MAC address is different from that specified in its fixed address, DHCP lease, or host record.
- **Discovered Routers/Switches:** Includes core network infrastructure devices of the specific Router, Switch, or Switch-Router types discovered by NIOS using the discovery feature set. Clicking on a device name opens the device page under **Data Management → Devices** and shows the **Interfaces** page for the chosen device.
- **Gaming Console Devices:** Includes devices that were detected as gaming consoles.
- **Router and Wireless Access Point Devices:** Includes devices there were detected as routers or wireless access point devices.
- **Unmanaged:** Shows all unmanaged devices.

Note: For information about DHCP fingerprints, see [About DHCP Fingerprints](#) on page 1033. For information about discovery, see [Network Insight](#) beginning on page 517.

CREATING SMART FOLDERS

You can create personal smart folders in My Smart Folders. You can also create global folders to share among administrators in Global Smart Folders when you log in as a superuser account. Each time you access a smart folder, you obtain up-to-date information about the core network services data that match the filter criteria you set for the folder. You can also set Group By rules to group the filtered data by extensible attributes. Grid Manager displays a hierarchical view of the data using the Group By rules you define.

To create a smart folder:

1. Click the **Smart Folders** tab.
2. Click the **My Smart Folders** tab to create a personal smart folder.
or
If you logged in with a superuser account, click the **Global Smart Folders** tab to create a global smart folder.
3. Click **Create**.
4. In the *Smart Folder* data panel, complete the following:
 - **Name:** Enter the name of the smart folder.
 - **Comment:** Optionally, enter additional information about the smart folder.
 - In the first drop-down list, select a field as the filter. You can select a network view or a record type as the filter. Grid Manager highlights extensible attributes in gray. You can also group the default data by adding a Group By rule without adding a filter. The default filter is “Type equals Network/Zone/Range/Member”.

Note: Infoblox strongly recommends that you use **Type** as the first filter criterion to improve system performance.

- In the second drop-down list, select an operator for the filter.

- Enter or select a value for the selected field and operator. Depending on the field and operator that you select, the field can be a text or an integer field. It can also be a drop-down list or a calendar widget. The default is **Network/Zone/Range/Member** if you select **Type** in the first field. Grid Manager displays all the networks, zones, DHCP ranges, and members in the results table. The results table may display the name in its native characters if the name was originally entered as an IDN (Internationalized Domain Name). For example, the **Name** column in the results table displays 网络, a zone name in Chinese.

Note: When you select “IF-MAP publishing equals Yes” as the filter, only networks and ranges that have IF-MAP publishing explicitly enabled appear in the results table. When you select “IF-MAP publishing equals No,” only networks and ranges that have IF-MAP publishing explicitly disabled appear in the results table.

- Optionally, click **+** to add another filter. You can also click **Apply** to view the filtered data in the results table.
- Optionally, select the **Group Results** check box to organize the filtered data. You can also disable a Group By filter by deselecting the check box.
- From the **Group by** drop-down list, select an extensible attribute by which you want to group the filtered data. For example, if you want to group the filtered data by building number, you can select **Building** from the drop-down list. To add additional Group By rules, click the **+** icon, and then select a field from the drop-down list. You can apply up to 10 Group By rules. You can also delete a rule by selecting the rule and clicking the **-** icon.
- After you add all filter criteria and Group By rules, click **Apply**. Grid Manager displays the filtered data in the results table. Note that in the **Name** field, the appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address. If you select to include objects with no attribute values in the Group By rules, the appliance may take longer to process the results.

5. Click **Save** to save the smart folder.

VIEWING AND MODIFYING DATA IN SMART FOLDERS

After you set up a smart folder, the appliance searches for matching objects based on the filter criteria you specified for the folder. Grid Manager also groups the objects by the Group By rules you specify. If you select to include objects with no attribute values, the appliance may take longer to process the results. Each smart folder you create can contain up to 2,000 objects. When the number of objects exceeds 2,000, Grid Manager sort and displays the first 2,000 objects and a message at the top of the panel. In this case, you may want to redefine your filter criteria to further refine the filtered data in your smart folders.

Grid Manager displays smart folders hierarchically in a tree view based on your Group By rules in the following:

- Smart Folder section in the Finder panel
- Selectors from which you can select a smart folder

In the smart folder list panel however, Grid Manager displays all the smart folders in a flat list. You can modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.

In the smart folder data panel, Grid Manager displays the first hierarchical level of the smart folder based on your Group By rules. If you do not configure any Group By rule, Grid Manager displays all the objects in the results table. If you select to include objects with no attribute values, Grid Manager also includes these objects in the hierarchical view. Depending on your Group By rules, you can view detailed information about the objects by clicking the object link and drilling down to the lowest hierarchical level, and then opening an object. To go back to a previous hierarchical view, click the link of the corresponding level in the breadcrumb.

To view detailed information about an object:

1. In the Smart Folder data panel, click the object link until you drill down to the last hierarchical level of the folder.
2. Grid Manager displays the following information:

- **Name:** The name or IP address of the object. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **Comment:** Information about the object.
- **Type:** The object type.
- **Site:** The site to which the object belongs. This is one of the predefined extensible attributes.

You can also select other available extensible attributes for display, and sort the data in ascending or descending order by column.

3. Select an object check box, and then do one of the following:
 - Click the Open icon to display the data in the network list or IP address list.
 - Click the Edit icon to modify or schedule the modification of the object configuration. Grid Manager displays the corresponding editor depending on the object you select.
 - Click the Delete icon to delete the object or click the Schedule Deletion icon to schedule the deletion of the object.

You can also print or export the data in this panel. For information, see [Infoblox Grid Manager](#) on page 43.

MODIFYING SMART FOLDERS

After you create a smart folder, you can modify its filter and grouping criteria.

To modify a smart folder:

1. Go to **Smart Folders**.
2. Click **My Smart Folders** to modify personal smart folders.
or
Click **Global Smart Folders** to modify global smart folders if you logged in with a superuser account.
3. Mouse over to the smart folder that you want to modify.
4. Click the Edit icon. You can also click the Edit icon next to the name of the smart folder in the data panel.
5. Make the appropriate changes in the Smart Folder data panel as described in [Creating Smart Folders](#) on page 142.

DELETING SMART FOLDERS

You can delete personal smart folders in My Smart Folders. However, you must log in as a superuser account to delete global smart folders.

To delete a smart folder:

1. Click the **Smart Folders** tab.
2. Click the **My Smart Folders** tab to delete personal smart folders.
or
Click the **Global Smart Folders** tab to delete global smart folders.
3. Mouse over to the smart folder that you want to delete.
4. Click the Delete icon. In the *Delete Smart Folder* dialog box, click **Yes**.

SAVING A COPY OF A SMART FOLDER

You can make a copy of an existing smart folder, add or change filter criteria, and then rename the folder accordingly. You can also create a local copy of the global smart folder in its current state for editing purposes. In My Smart Folders, you can save a folder copy only in My Smart Folders. In Global Smart Folders however, you can save a folder copy in either My Smart Folders or Global Smart Folders. You must have superuser permissions to save a global smart folder copy in Global Smart Folders. Note that when the original global smart folder is updated, information in your local copy is not updated.

To save a copy of a smart folder:

1. Click **My Smart Folders** to save a folder copy in this tab.

or

Click **Global Smart Folders** to save a folder copy in either this tab or My Smart Folders. To save a smart folder copy in Global Smart Folders, log in as a superuser account.

2. Select the smart folder that you want to save as a copy.

3. Click **Save copy as**.

4. Grid Manager saves the folder copy in My Smart Folders when you save the folder copy in this tab.

or

The *Save Smart Folder As* dialog box appears when you perform this function in Global Smart Folders. Select one of the following:

- **My Smart Folders:** Saves the copy in My Smart Folders.
- **Global Smart Folders:** Saves the copy in Global Smart Folders.

Click **OK**.

Note: For the folder copy, the appliance appends the word *Copy* to the original name of the smart folder. You can change the name of the folder at anytime by editing the folder.

PRINTING AND EXPORTING DATA IN SMART FOLDERS

You can print the list of networks that are on the current smart folder page, or you can export all the data in CSV (comma separated value) format. For information, see [About Tasks](#) on page 72 and [Exporting Data to Files](#) on page 89.



PART 2 APPLIANCE ADMINISTRATION

This section provides information about configuring admin groups, roles, and accounts, and defining the appropriate permissions. It explains how to configure and manage a Grid or an independent appliance, and set operational parameters. It also describes the file distribution services (TFTP, FTP, and HTTP) and the bloxTools environment. It includes the following chapters:

- [Chapter 4, *Managing Administrators*](#), on page 149
- [Chapter 5, *Deploying a Grid*](#), on page 221
- [Chapter 6, *Deploying Independent Appliances*](#), on page 273
- [Chapter 7, *Managing Appliance Operations*](#), on page 303
- [Chapter 8, *File Distribution Services*](#), on page 389
- [Chapter 9, *Managing NIOS Software and Configuration Files*](#), on page 405
- [Chapter 10, *bloxTools Environment*](#), on page 431
- [Chapter 11, *RIR Registration Updates*](#), on page 437



Chapter 4 Managing Administrators

This chapter describes the various tasks associated with setting up admin groups, admin roles, admin accounts, and permissions. It contains the following sections:

- [About Admin Accounts](#) on page 152
- [About Admin Groups](#) on page 154
 - [Creating Superuser Admin Groups](#) on page 155
 - [Creating Limited-Access Admin Groups](#) on page 156
- [About Admin Roles](#) on page 157
 - [Creating Admin Roles](#) on page 157
- [Managing Admin Groups and Admin Roles](#) on page 158
 - [Modifying Admin Groups and Roles](#) on page 158
 - [Deleting Admin Groups and Roles](#) on page 159
 - [Viewing Admin Groups](#) on page 159
 - [Viewing Admin Roles](#) on page 159
 - [Viewing Admin Group Assignments](#) on page 160
- [About Administrative Permissions](#) on page 160
 - [Defining Global Permissions](#) on page 161
 - [Defining Object Permissions](#) on page 162
 - [Defining DNS and DHCP Permissions for Grid Members](#) on page 164
 - [Applying Permissions and Managing Overlaps](#) on page 166
 - [Managing Permissions](#) on page 167
- [Authenticating Administrators](#) on page 169
- [Creating Local Admins](#) on page 169
 - [Managing Passwords](#) on page 170
 - [Modifying and Deleting Admin Accounts](#) on page 171
- [About Remote Admins](#) on page 171
- [Authenticating Admins Using RADIUS](#) on page 173
 - [Authentication Protocols](#) on page 174
 - [Accounting Activities Using RADIUS](#) on page 174
 - [Configuring Remote RADIUS Servers](#) on page 174
 - [Configuring RADIUS Authentication](#) on page 175
 - [Configuring a RADIUS Authentication Server Group](#) on page 175

- [*Authenticating Admins Using Active Directory*](#) on page 177
 - [*Configuring an Active Directory Authentication Service Group*](#) on page 178
- [*Authenticating Admin Accounts Using TACACS+*](#) on page 179
 - [*TACACS+ Accounting*](#) on page 180
 - [*Configuring TACACS+*](#) on page 180
 - [*Configuring a TACACS+ Authentication Server Group*](#) on page 180
- [*Authenticating Admins Using LDAP*](#) on page 182
 - [*Authentication Protocols*](#) on page 183
 - [*Configuring LDAP*](#) on page 183
 - [*Configuring an LDAP Server Group*](#) on page 183
- [*Defining the Authentication Policy*](#) on page 185
 - [*Configuring a List of Authentication Server Groups*](#) on page 185
 - [*Configuring a List of Remote Admin Groups*](#) on page 186
- [*Authenticating Admins Using Two-Factor Authentication*](#) on page 187
 - [*Best Practices for Configuring Two-Factor Authentication*](#) on page 189
 - [*Configuring the OSCP Authentication Server Group*](#) on page 189
 - [*Viewing the OSCP Authentication Server Group*](#) on page 191
- [*Changing Password Length Requirements*](#) on page 191
- [*Notifying Administrators*](#) on page 191
- [*Administrative Permissions for Common Tasks*](#) on page 193
- [*Administrative Permission for the Grid*](#) on page 195
 - [*Administrative Permissions for Grid Members*](#) on page 195
 - [*Administrative Permissions for Network Discovery*](#) on page 196
 - [*Administrative Permissions for Scheduling Tasks*](#) on page 196
 - [*Administrative Permissions for Microsoft Servers*](#) on page 197
- [*Administrative Permissions for IPAM Resources*](#) on page 198
 - [*Administrative Permissions for IPv4 and IPv6 Networks*](#) on page 198
 - [*Administrative Permissions for Hosts*](#) on page 199
- [*Administrative Permissions for DNS Resources*](#) on page 199
 - [*Administrative Permissions for DNS Views*](#) on page 200
 - [*Administrative Permissions for Zones*](#) on page 201
 - [*Administrative Permissions for Resource Records*](#) on page 202
 - [*Administrative Permissions for Shared Record Groups*](#) on page 203
 - [*Administrative Permissions for DNS64 Synthesis Groups*](#) on page 204
- [*Administrative Permissions for DHCP Resources*](#) on page 205
 - [*Administrative Permissions for Network Views*](#) on page 206
 - [*Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks*](#) on page 207
 - [*Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations*](#) on page 208
 - [*Administrative Permissions for IPv4 or IPv6 DHCP Enabled Host Addresses*](#) on page 209
 - [*Administrative Permissions for IPv4 and IPv6 DHCP Ranges*](#) on page 210
 - [*Administrative Permissions for IPv4 or IPv6 DHCP Templates*](#) on page 211
 - [*Administrative Permissions for Roaming Hosts*](#) on page 212
 - [*Administrative Permissions for MAC Address Filters*](#) on page 212
 - [*Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories*](#) on page 213

-
- [*Administrative Permissions for File Distribution Services*](#) on page 214
 - [*Administrative Permissions for Dashboard Tasks*](#) on page 214
 - [*Administrative Permissions for OCSP Server Groups and CA Certificates*](#) on page 215
 - [*Administrative Permissions for Load Balancers*](#) on page 216
 - [*Administrative Permissions for Named ACLs*](#) on page 217
 - [*Administrative Permissions for DNS Threat Protection*](#) on page 217

ABOUT ADMIN ACCOUNTS

A user must have an admin account to log in to the NIOS appliance. Each admin account belongs to an admin group, which contains roles and permissions that determine the tasks a user can perform. For information, see [About Admin Groups](#) on page 154.

When an admin connects to the appliance and logs in with a username and password, the appliance starts a two-step process that includes both authentication and authorization. First, the appliance tries to authenticate the admin using the username and password. Second, it determines the authorized privileges of the admin by identifying the group to which the admin belongs. It grants access to the admin only when it successfully completes this process.

The NIOS appliance can authenticate users that are stored on its local database as well as users stored remotely on an Active Directory domain controller, a RADIUS server, a TACACS+ server or an LDAP server. The group from which the admin receives privileges and properties is stored locally.

The appliance can also authenticate users with smart cards that contain X.509 client certificates. The status of these certificates is stored remotely on OSCP (Online Certificate Status Protocol) responders. NIOS uses two-factor authentication to validate these users. For more information about two-factor authentication and how to configure it, see [Authenticating Admins Using Two-Factor Authentication](#) on page 187.

The tasks involved in configuring administrator accounts locally and remotely are listed in [Table 4.1](#).

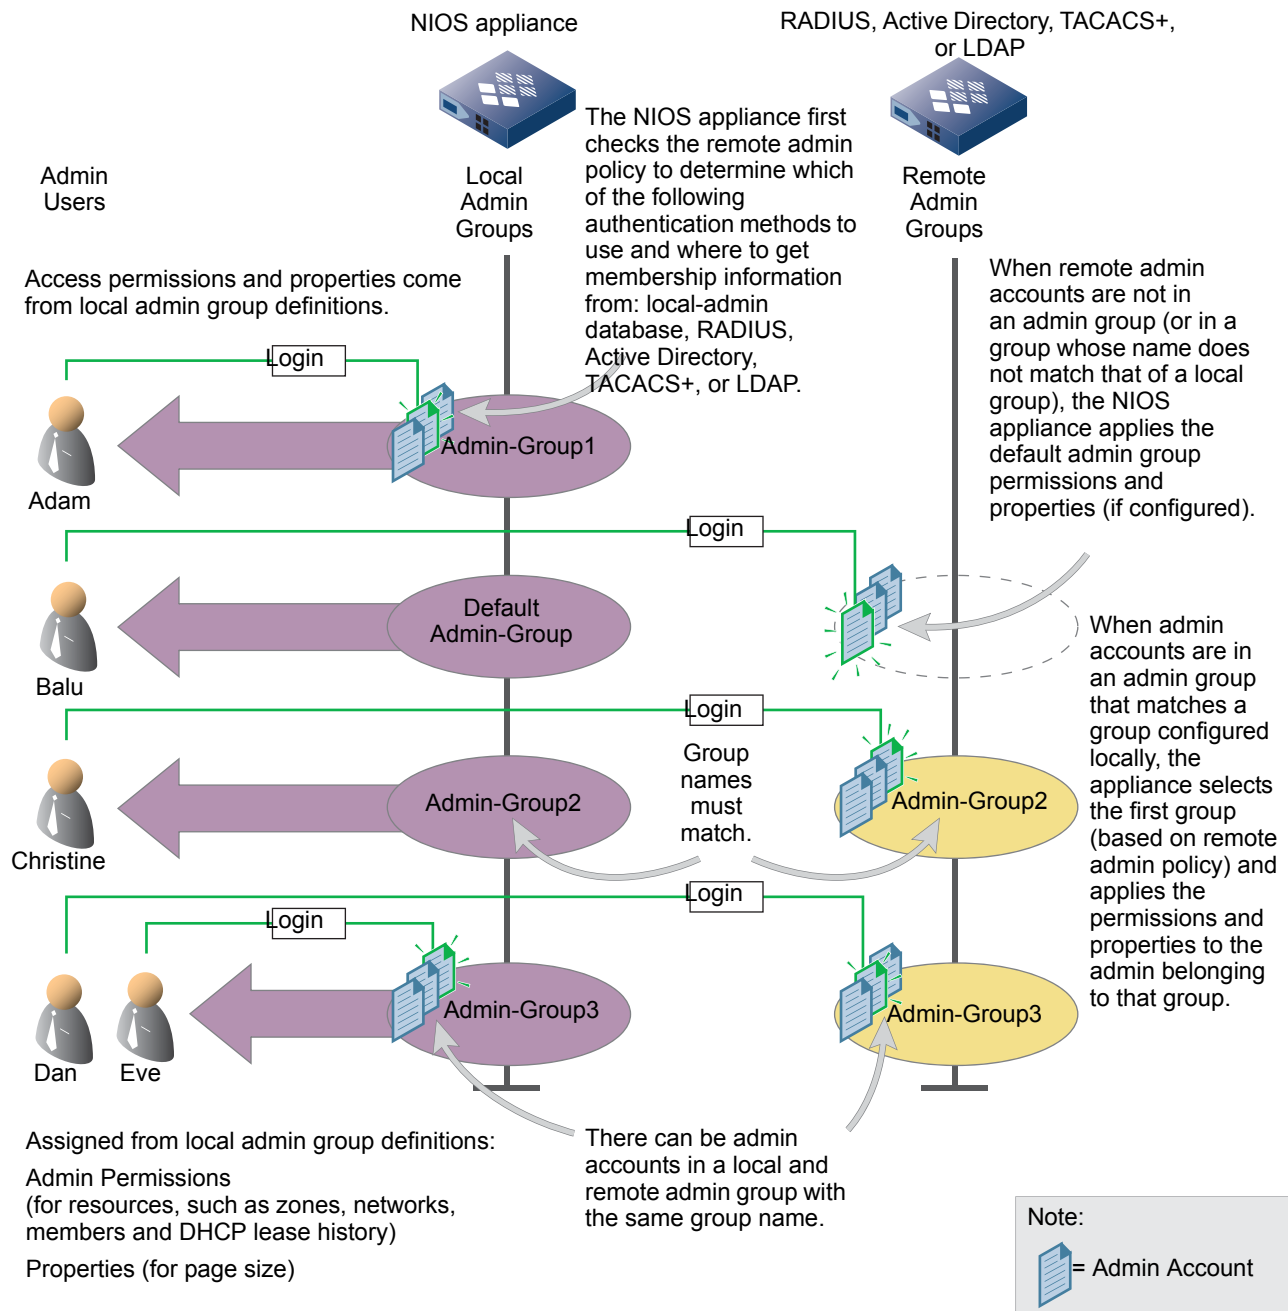
Table 4.1 Storing Admin Accounts Locally and Remotely

	NIOS Appliance	RADIUS server/AD Domain Controller/TACACS+ server/LDAP server
To store admin accounts locally	<ul style="list-style-type: none"> Use the default admin group (“admin-group”) or define a new group Set the privileges and properties for the group Add admin accounts to the group 	
To store admin accounts remotely	<ul style="list-style-type: none"> Configure communication settings with a RADIUS server, an Active Directory domain controller, TACACS+ server, or LDAP server <p>If you use admin groups on the RADIUS server, Active Directory domain controller, TACACS+ server, or LDAP server:</p> <ul style="list-style-type: none"> Configure admin groups that match the remote admin groups Set the privileges and properties for the groups <p>If you do not use admin groups on the RADIUS server, Active Directory domain controller, TACACS+ server, or LDAP server:</p> <ul style="list-style-type: none"> Assign an admin group as the default 	<ul style="list-style-type: none"> Configure communication settings with the NIOS appliance <p>If you use admin groups:</p> <ul style="list-style-type: none"> Import Infoblox VSAs (vendor-specific attributes) (if RADIUS) Define an admin group with the same name as that on the NIOS appliance Define admin accounts and link them to an admin group <p>If you do not use admin groups:</p> <ul style="list-style-type: none"> Define admin accounts

The admin policy defines how the appliance authenticates the admin: with the local database, RADIUS, Active Directory, TACACS+, or LDAP. You must add RADIUS, Active Directory, TACACS+, or LDAP as one of the authentication methods in the admin policy to enable that authentication method for admins. See [Defining the Authentication Policy](#) on page 185 for more information about configuring the admin policy.

[Figure 4.1](#) illustrates the relationship of local and remote admin accounts, admin policy, admin groups, and permissions and properties.

Figure 4.1 Privileges and Properties Applied to Local and Remote Admin Accounts



Complete the following tasks to create an admin account:

1. Use the default admin group or create an admin group. See [About Admin Groups](#) on page 154.
2. Define the administrative permissions of the admin group. See [About Administrative Permissions](#) on page 160.
3. Create the admin account and assign it to the admin group.
 - To add the admin account to the local database, see [Creating Local Admins](#) on page 169.
 - To configure the appliance to authenticate the admin account stored remotely, see [About Remote Admins](#) on page 171.

ABOUT ADMIN GROUPS

All administrators must belong to an admin group. The permissions and properties that you set for a group apply to all administrators assigned to that group. You can assign a dashboard template to an admin group. A dashboard template specifies the tasks an admin group can access through the **Tasks Dashboard** tab when they log in to Grid Manager. For information about dashboard templates, see [About Dashboard Templates](#) on page 114. You can also restrict certain user groups to manage specific tasks in the **Tasks Dashboard** tab only. These users cannot manage other core network services through Grid Manager. For information about how to apply this restriction, see [Creating Limited-Access Admin Groups](#) on page 156.

To define admins who can perform specific core network service tasks, you can set up admin groups and assign them permissions for those tasks. To control when and whether certain tasks should be performed, you can add an admin group to an approval workflow and define the admins as submitters or approvers. A submitter is an admin whose tasks require approvals before execution, and an approver is an admin who can approve the submitted tasks. When you add submitter and approver groups to an approval workflow, you have control over who can perform which mission critical tasks and whether and when the tasks should be executed. For more information about how to create and configure approval workflows, see [Configuring Approval Workflows](#) on page 80.

There are three types of admin groups:

- **Superuser** – Superuser admin groups provide their members with unlimited access and control of all the operations that a NIOS appliance performs. There is a default superuser admin group, called **admin-group**, with one superuser administrator, **admin**. You can add users to this default admin group and create additional admin groups with superuser privileges. Superusers can access the appliance through its console, GUI, and API. In addition, only superusers can create admin groups.
- **Limited-Access** – Limited-access admin groups provide their members with read-only or read/write access to specific resources. These admin groups can access the appliance through the GUI, API, or both. They cannot access the appliance through the console.
- **Default** – When upgrading from previous NIOS releases, the appliance converts the ALL USERS group to the Default Group when the ALL USERS Group contains admin accounts. The appliance does not create the Default Group if there is no permission in the ALL USERS group. The permissions associated with the ALL USERS group are moved to a newly created role called Default Role. Supported in previous NIOS releases, the ALL USERS group was a default group in which you defined global permissions for all limited-access users. This group implicitly included all limited-access users configured on the appliance.

All limited-access admin groups require either read-only or read/write permission to access certain resources, such as Grid members, and DNS and DHCP resources, to perform certain tasks. Therefore, when you create an admin group, you must specify which resources the group is authorized to access and their level of access.

Only superusers can create admin groups and define their administrative permissions. There are two ways to define the permissions of an admin group. You can create an admin group and assign permissions directly to the group, or you can create roles that contain permissions and assign the roles to an admin group.

Complete the following tasks to assign permissions directly to an admin group:

1. Create an admin group, as described in [Creating Limited-Access Admin Groups](#) on page 156.
2. Assign permissions to the admin group, as described in [About Administrative Permissions](#) on page 160.

Complete these tasks to assign admin roles to an admin group:

1. Create an admin role, as described in [About Admin Roles](#) on page 157.
2. Define permissions for the newly created admin role, as described in [Creating Admin Roles](#) on page 157.
3. Create an admin group and assign the role to the group, as described in [Creating Limited-Access Admin Groups](#) on page 156.

After you have created admin groups and defined their administrative permissions, you can assign administrators to the group.

- For local admins, see [Creating Local Admins](#) on page 169.
- For remote admins, see [About Remote Admins](#) on page 171.

Creating Superuser Admin Groups

Superusers have unlimited access to the NIOS appliance. They can perform all operations that the appliance supports. There are some operations, such as creating admin groups and roles, that only superusers can perform.

Note that there must always be one superuser admin account, called “admin”, stored in the local database to ensure that at least one administrator can log in to the appliance in case the NIOS appliance loses connectivity to the remote admin databases such as RADIUS servers, AD domain controllers, TACACS+ servers, LDAP servers, or OCSP responders.

NIOS comes with a default superuser admin group (**admin-group**). It also automatically creates a new admin group, **fireeye-group**, when you add the first FireEye RPZ (Response Policy Zone). Infoblox recommends that you do not add another admin group with the same name as the default or FireEye admin group. Note that the FireEye admin group is read-only and you cannot assign permissions to it. For more information about FireEye RPZs, see [About FireEye Integrated RPZs](#) on page 1254. You can create additional superuser admin groups, as follows:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab, and then click the Add icon.
2. In the *Add Admin Group* wizard, complete the following:
 - **Name:** Enter a name for the admin group.
 - **Comment:** Enter useful information about the group, such as location or department. For **fireeye-group**, NIOS displays **Group used to receive FireEye alerts** in this field.
 - **Disable:** Select this to retain an inactivated profile for this admin group in the configuration. For example, you may want to define a profile for recently hired administrators who have not yet started work. Then when they do start, you simply need to clear this check box to activate the profile.
3. Click **Next** and complete the following:
 - **Superusers:** Select this to grant the admin accounts that you assign to this group full authority to view and configure all types of data and perform all tasks.
4. Click **Next** and complete the following to define the dashboard template:
 - **Dashboard Template:** From the drop-down list, select the dashboard template you want to assign to this superuser group. When you apply a dashboard template to an admin group, the template applies to all users in the group. The default is **None**, which means that users in this group can access all licensed tasks in the **Tasks Dashboard** tab if they have the correct permissions to the task-related objects. Note that if you want to delete a template, you must first unassign the template from an admin group, or select **None**, before you can delete it. For more information about dashboard templates, see [About Dashboard Templates](#) on page 114.
5. Click **Next** to add admin email addresses if you want the appliance to send approval workflow notifications to a list of email addresses for the admin group. Complete the following in the Email Address table:

Click the Add icon and Grid Manager adds a row to the table. Enter the email address of the admin who should receive workflow notifications. You can click the Add icon again to add more email addresses. You can also select an email address and click the Delete icon to delete it. To modify an email address, click the **Email Address** column and modify the existing address.

Note: When you configure an approval workflow and select **Group Email Address(es)** as the approver notification addresses, the appliance sends workflow notifications to all email addresses you have added to this table. For information, see [Configuring Approval Workflows](#) on page 80.

6. Optionally, click **Next** to add extensible attributes to the admin group. For information, see [About Extensible Attributes](#) on page 322.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

You can do one of the following after you create a superuser admin group:

- Add local admins to the superuser group. For information, see [Creating Local Admins](#) on page 169.
- Assign the superuser group to remote admins. For information, see [About Remote Admins](#) on page 171.

Creating Limited-Access Admin Groups

When you create a limited-access admin group, you can assign roles to it. The group then inherits the permissions of its assigned roles. In addition, you can assign permissions directly to the group. Only superusers can create admin groups.

To create a limited-access admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab, and then click the Add icon.
2. In the *Add Admin Group* wizard, complete the following:
 - **Name:** Enter a name for the admin group.
 - **Comment:** Enter useful information about the group, such as location or department.
 - **Disable:** Select this to retain an inactivated profile for this admin group in the configuration. For example, you may want to define a profile for recently hired administrators who have not yet started work. Then when they do start, you simply need to clear this check box to activate the profile.

3. Click **Next** and complete the following:

- **Superusers:** Clear this check box to create a limited-access admin group.
- **Roles:** Optionally, click the Add icon to add an admin role to the admin group. In the *Role Selector* dialog box, select the roles you want to assign to the admin group, and then click the Select icon. Use Shift+click and Ctrl+click to select multiple admin roles. You can assign up to 21 roles to an admin group. The appliance displays the selected roles in the list box.

When an admin group is assigned multiple roles, the appliance applies the permissions to the group in the order the roles are listed. Therefore if there are overlapped permissions among the roles, the appliance uses the permission from the role that is listed first and ignores the others. You can reorder the list by selecting a role and clicking the arrow keys to move the role up and down the list. To delete a role, select it and click the Delete icon.

- **Allowed Interfaces:** Specify whether the admin group can use the Grid Manager GUI and the API (application programming interface) to configure the appliance.
 - **GUI:** Select this to allow the admin group to use the GUI.
 - **API:** Select this to allow the admin group to use the API.

4. Click **Next** and complete the following to define the dashboard template:

- **Dashboard Template:** From the drop-down list, select the dashboard template you want to assign to this superuser group. When you assign a dashboard template to an admin group, the template applies to all users in the group. The default is **None**, which means that users in this group can perform all licensed tasks in the **Tasks Dashboard** tab if they have the correct permissions to the task-related objects. Note that if you want to delete a template, you must first unassign the template from an admin group, or select **None**, before you can delete it. For more information about dashboard templates, see [About Dashboard Templates](#) on page 114.
- **Display Taskflow Dashboards Only:** Select this check box if you want to restrict this admin group to access only the Tasks Dashboard in Grid Manager. Note that when you select this check box, users in this admin group have access to the tasks you specified in the selected dashboard template, if applicable. They cannot perform any other tasks or manage any core network services in Grid Manager the next time they log in to the system.

5. Click **Next** to add admin email addresses if you want the appliance to send approval workflow notifications to a list of email addresses for the admin group. Complete the following in the Email Address table:

Click the Add icon and Grid Manager adds a row to the table. Enter the email address of the admin who should receive workflow notifications. You can click the Add icon again to add more email addresses. You can also select an email address and click the Delete icon to delete it. To modify an email address, click the **Email Address** column and modify the existing address.

Note: When you configure an approval workflow and select **Group Email Address(es)** as the approver notification addresses, the appliance sends workflow notifications to all email addresses you have added to this table. For information, see [Creating Approval Workflows](#) on page 81.

6. Optionally, click **Next** to add or delete extensible attributes for this admin group. For information, see [About Extensible Attributes](#) on page 322.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

ABOUT ADMIN ROLES

An admin role is a group of permissions that you can apply to one or more admin groups. Roles allow you to quickly and easily apply a suite of permissions to an admin group. You can define roles once and apply them to multiple admin groups. The appliance contains the following system-defined admin roles:

- **DHCP Admin:** Provides read/write access to all network views, all DHCP MAC filters, all Grid members, and all Microsoft servers that are managed by the Grid. It also provides read-only access to all DHCP templates and DHCP lease history.
- **DNS Admin:** Provides read/write access to all Grid members, all Microsoft servers that are managed by the Grid, all shared record groups, and all DNS views.
- **File Distribution Admin:** Provides read/write access to Grid file distribution properties.
- **Grid Admin:** Provides read/write access to all DNS views, all shared record groups, all members, all Microsoft servers that are managed by the Grid, all network views, all DHCP MAC filters, all DHCP templates, DHCP lease history, Grid File distribution properties, network discovery, task scheduling, and all Dashboard tasks.
- **Load Balancer Admin:** Provides read/write access to all load balancer resources.
- **PKI Admin:** Provides read/write access to all HSM groups, all OCSP groups, and all CA certificates.
- **DHCP Fingerprint:** Provides read/write access to all DHCP fingerprint related objects.

You can assign these system-defined roles to admin groups and create additional roles based on the job functions in your organization. If you are creating a role that has similar permissions to an existing role, you can copy the role and then make the necessary modifications to the new role. Thus you do not have to create each new role from scratch.

You can assign up to 21 roles to an admin group, and you can assign a role to more than one admin group. When you make a change to a role, the appliance automatically applies the change to that role in all admin groups to which the role is assigned.

Creating Admin Roles

There are two ways to create an admin role. You can create a new role and define its permissions, or you can copy an existing role and redefine the configuration for the new role.

To create a new role from scratch:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab, and then click the Add icon.
2. In the *Add Role* wizard, complete the following:
 - **Name:** Enter a name for the role.
 - **Comment:** Enter useful information about the role. For example, if you are creating a role for IT personnel, you can put the information here.
 - **Disable:** Select this to retain an inactivated profile for this admin role in the configuration.
3. Optionally, click **Next** to add extensible attributes to this role. For information, see [About Extensible Attributes](#) on page 322.

4. Click **Next** and select one of the following:
 - **Save & Add Permissions:** Save the entry and add permissions to the role. Grid Manager displays the **Permissions** tab with the newly created role selected. You can then add permissions to this role. For information, see [About Administrative Permissions](#) on page 160.
 - **Save & Close:** Save the entry and close the wizard.
 - **Save & Edit:** Save the entry and continue to edit.
 - **Save & New:** Save the entry and open a new wizard.

To copy an existing role:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab -> *admin_role* check box, and then click **Clone** from the Toolbar.
2. The *Copy Role* editor provides the following tabs from which you can modify data for the new role:
 - *General:* Enter the name and information about the new role. You can also disable the role in this tab.
 - *Admin Groups:* Displays a list of admin groups that are currently using this role. You cannot modify the list.
 - *Extensible Attributes:* Add and delete extensible attributes that are associated with the admin role. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
3. Save the configuration and click **Restart** if it appears at the top of the screen.
The appliance displays the new role in the **Roles** tab.

After you create roles, you can do the following:

- Define their permissions. For information and guidelines on defining permissions, see [About Administrative Permissions](#) on page 160.
- Assign roles to admin groups, as described in [Creating Limited-Access Admin Groups](#) on page 156.

MANAGING ADMIN GROUPS AND ADMIN ROLES

After you create an admin group or an admin role, you can view, modify, and delete it.

Modifying Admin Groups and Roles

To modify an admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin_group* check box, and then click the Edit icon.
2. The *Admin Group* editor provides the following tabs from which you can modify data:
 - **General:** You can modify the following data.
 - **Name:** Modify the name of the admin group.
 - **Comment:** Enter useful information about the group, such as location or department.
 - **Disable:** Select this to retain an inactivated profile for this admin group in the configuration. For example, you may want to define a profile for recently hired administrators who have not yet started work. Then when they do start, you simply need to clear this check box to activate the profile.
 - **Roles:** Modify the data as described in [Creating Limited-Access Admin Groups](#) on page 156.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the admin group. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Deleting Admin Groups and Roles

You can remove any default or custom admin group as long as it is not your own admin group or the last admin group. You can also delete any default or custom admin role. The appliance puts the deleted roles in the Recycle Bin, if enabled.

To delete an admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin_group* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

To delete an admin role:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab -> *admin_role* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

Viewing Admin Groups

You can view the list of admin groups that are currently in the Grid. To view admin groups, from the **Administration** tab, select the **Administrators** tab -> **Groups** tab.

Grid Manager displays the following information:

- **Name:** The name of the admin group.
- **Superuser:** Indicates whether the admin accounts that you assign to this group have full authority to view and configure all types of data. The value can be **Yes** or **No**.
- **Comment:** The information about the admin group.

You can select the additional fields, **Disabled** and **Site**, for display.

You can also do the following:

- Sort the data in ascending or descending order by column.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Print or export the data.

Viewing Admin Roles

You can view the list of admin roles that are currently in the Grid. To view admin roles, from the **Administration** tab, select the **Administrators** tab -> **Roles** tab.

Grid Manager displays the following information:

- **Name:** The name of the admin role.
- **System:** Indicates whether the admin role is system defined or not. The value can be **Yes** or **No**.
- **Comment:** The information about the admin role.

You can select the additional fields, **Disabled** and **Site**, for display. You can also do the following:

- Sort the data in ascending or descending order by column.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Print or export the data.

Viewing Admin Group Assignments

After you define permissions for an admin role, you can assign it to multiple admin groups. You can view the list of admin groups to which an admin role is assigned, as follows:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab -> *admin_group* check box, and then click the Edit icon.
2. In the *Role* editor, select the **Admin Groups** tab.
Grid Manager displays the list of admin groups to which the role is assigned.

ABOUT ADMINISTRATIVE PERMISSIONS

Admin permissions define actions that an admin role or admin group can perform for specific resources. For example, when you set a read/write (RW) permission for a DNS zone and then assign the permission to an admin group, all admin users in the group can view, add, modify, and delete DNS objects in the zone. On the other hand, if the assigned permission is read-only (RO), admins can only view and search for objects in the zone.

Following are permissions you can set for supported resources:

- Read/Write (RW): Allows admins to add, modify, delete, view, and search for a resource.
- Read-Only (RO): Allows admins to view and search for a resource. Admins cannot add, modify, or delete the resource.
- Deny: Prevents admins from adding, modifying, deleting, and viewing a resource. This is the default permission for all resources.

You can assign permissions to an admin role that you then reassign to an admin group, or you can assign permissions directly to an admin group. By default, the superuser group (**admin-group**) has full access to all resources on the appliance. Superusers can create limited-access admin groups and grant them permissions for resources at the global and object levels. Limited-access admin groups must have either read-only or read/write permissions assigned in order to view information or perform tasks for any supported objects.

When you assign permissions at the global level, the permissions apply to all objects that belong to the specified resource. For example, when you define a read/write permission to all DHCP networks, the permission applies to all DHCP ranges and fixed addresses in the networks. For information about global permissions, see [Defining Global Permissions](#) on page 161.

You can also define permissions at a more granular level, such as for a specific Grid member, DNS zone, RPZ (Response Policy Zone), network, and even an individual database object, such as a resource record or fixed address. When you define a permission at the object level, admins with this permission can only manage the specified object and its associated objects. For information about object permissions, see [Defining Object Permissions](#) on page 162.

You can use global and object permissions to restrict admins to specific DNS and DHCP resources on specific Grid members by assigning the appropriate permissions. You can use this feature to separate DNS and DHCP administration on selected Grid members. For more information, see [Defining DNS and DHCP Permissions for Grid Members](#) on page 164.

You can configure global permissions, object permissions, and member DNS and DHCP permissions for default and custom admin groups and roles. However, you cannot define permissions for the factory default roles, such as DHCP Admin.

The appliance supports the following resource groups:

- Grid resources: Includes Grid DNS properties, Grid DHCP properties, all Grid members, Microsoft servers that are managed by the Grid, network discovery, task scheduling, CSV imports, and all dashboard tasks.

- IPAM resources: Includes network views, IPv4 and IPv6 networks, and host records.
- DHCP resources: Includes Grid DHCP properties, network views, IPv4 networks, host records, DHCP ranges, DHCP fixed addresses/reservations, DHCP enabled host addresses, Mac filters, shared networks, DHCP templates, lease history, and roaming hosts.
- DNS resources: Includes Grid DNS properties, DNS views, DNS zones, Response Policy Zones, host records, bulk hosts, all DNS resource records, and all shared records.
- File distribution resources: Includes Grid-level file distribution properties.
- Reporting resources: Includes Grid-level reporting properties.
- Administration resources: Includes all OSCP server groups and CA certificates.
- GLB (Global Load Balancer) resources: Includes all NIOS managed GLB objects.
- DHCP fingerprint resources: Includes all DHCP fingerprint related objects.
- Named ACL resources: Includes all named ACLs (access control lists).

The appliance applies permissions hierarchically in a parent-child structure. When you define a permission for a resource, the permission applies to all the resources and objects contained within that resource. For example, if you grant an admin group read/write permission for a network, it automatically has read/write permission for all objects in the network. However, you can override the network-level permission by setting a different permission, read-only or deny, for a fixed address or a DHCP enabled host address. Permissions for a Grid member apply to all zones and resource records served by that Grid member, and permissions for a network view apply to all DHCP resources within that view. To override permissions set at a higher level, you must define permissions at a more specific level. To define permissions for a more specific level, see the following:

- Permissions for common tasks, as described in [Administrative Permissions for Common Tasks](#) on page 193.
- Permissions for the Grid and Grid members, as described in [Administrative Permission for the Grid](#) on page 195.
- Permissions for IPAM resources, such as IPv6 networks, as described in [Administrative Permissions for IPAM Resources](#) on page 198.
- Permissions for DNS resources, such as DNS views and A records, as described in [Administrative Permissions for DNS Resources](#) on page 199.
- Permissions for DHCP resources, such as network views and fixed addresses, as described in [Administrative Permissions for DHCP Resources](#) on page 205.
- Permissions for file distribution services, as described in [Administrative Permissions for File Distribution Services](#) on page 214.
- Permissions for OSCP server groups and CA certificates, as described in [Administrative Permissions for OSCP Server Groups and CA Certificates](#) on page 215.
- Permissions for GLB and GLB objects, as described in [Administrative Permissions for Load Balancers](#) on page 216.

When you set permissions that overlap with existing permissions, Grid Manager displays a warning about the overlaps. You can view detailed information and find out which permissions the appliance uses and which ones it ignores. For information, see [Applying Permissions and Managing Overlaps](#) on page 166.

Defining Global Permissions

You can define permissions at a global level for an admin group or admin role.

To define global permissions:

1. For an admin group: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin_group* in the Groups table, and then click the Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.

or

For an admin role: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin_role* in the Roles table, and then click Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.

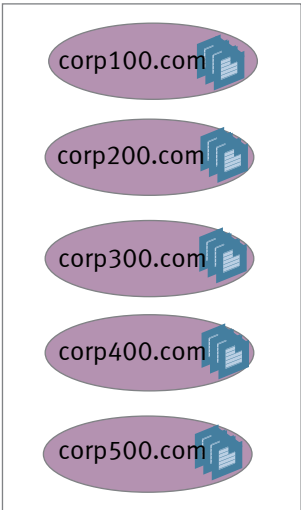
- Grid Manager displays the *Manage Global Permissions* editor. For an admin group, the appliance displays the selected admin group in the **Group Permission** field. For an admin role, the appliance displays the selected admin role in the **Role Permission** field. You can also select a different group or role from the drop-down list.
- Select the resources that you want to configure from the **Permission Type** drop-down list. Depending on your selection, Grid Manager displays the corresponding resources for the selected permission type in the table.
- Select **Read/Write**, **Read-Only**, or **Deny** for the resources you want to configure. By default, the appliance denies access to resources if you do not specifically configure them.
- Optionally, select additional resources from the **Permission Type** drop-down list. Grid Manager appends the new resources to the ones that you have already configured. Define the permissions for the resources you select.
- Save the configuration and click **Restart** if it appears at the top of the screen.

Defining Object Permissions


You can add permissions to specific objects for selected admin groups or roles. When you add permissions to objects, you can select multiple objects with the same or different object types. When you select multiple objects with the same object type, you can apply permissions to the selected objects as well as the sub object types that are contained in the selected objects. As described in [Figure 4.2](#), when you select five DNS forward-mapping authoritative zones, the appliance displays the object type “AuthZone” for all the zones. Since all five DNS zones are of the same object type, you can also apply permissions to all the resource records in these zones. The appliance displays the resources in the resource section of the *Create Object Permissions* editor. You can choose one or more of the resources to which you want to apply permissions.

Figure 4.2 Selecting Multiple Objects with the Same Object Type

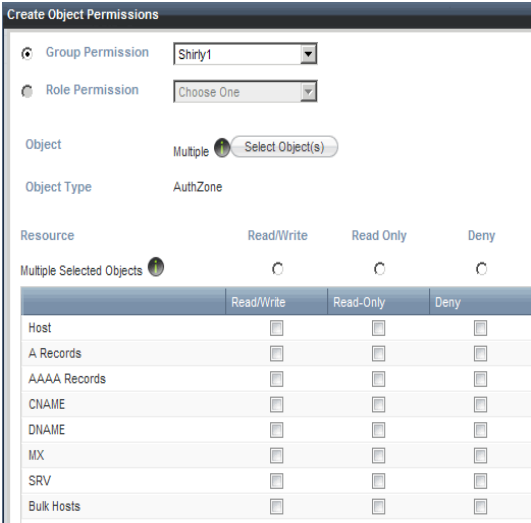
① You select five forward-mapping authoritative DNS zones that have resource records such as A records, Hosts, and CNAME records.



② Since all DNS zones have the same object type, you can apply object permissions to all the DNS zones as well as to all the resource records in the DNS zones.




③ The appliance displays the common resources in the **Resources in Selected Objects** column.



When you select multiple objects with more than one object type, you can add permissions to the selected objects as well as to the sub object types that are common among the selected objects. For example, when you select three DNS forward-mapping authoritative zones and two DNS IPv4 reverse-mapping authoritative zones as illustrated in [Figure 4.3](#), you can apply permissions to all the five DNS zones as well as to the CNAME, DNAME, and host records in these zones because CNAME, DNAME, and host records are the common sub object types in these zones.

Figure 4.3 Multiple Objects with Common Sub Object Types

When you select three DNS forward-mapping authoritative zones and two IPv4 reverse-mapping authoritative zones, you can apply object permissions to all the DNS zones as well as the CNAME, DNAME and Host records in these DNS zones.



To define object permissions for an admin group or role:

1. For an admin group: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin_group* in the Groups table, and then click the Add icon -> **Object Permissions** from the Create New Permission area or select **Add** -> **Object Permissions** from the Toolbar.
or
For an admin role: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin_role* in the Roles table, and then click Add icon -> **Object Permissions** from the Create New Permission area or select **Add** -> **Object Permissions** from the Toolbar.
2. Grid Manager displays the *Create Object Permissions* wizard. For an admin group, the appliance displays the selected group in the **Group Permission** field. For an admin role, the appliance displays the selected admin role in the **Role Permission** field. You can also select a different group or role from the drop-down list.
3. Click **Select Object(s)**. Grid Manager displays the *Object Selector* dialog box.
4. In the *Object Selector* dialog box, complete the following:
 - Enter a value or partial value of an object in the first field. This field is not case-sensitive. For example, if the object to which you want to define permissions contains “Infoblox”, enter Infoblox here.
 - Select the object type for which you are searching in the **Type** drop-down list. By default, the appliance searches all object types.
 - In the operator drop-down list, select an operator for the filter criteria. Depending on what you select in the first filter field, this list displays the relevant operators for the selection.
 - In the value field, enter or select the attribute value for the first filter field. Depending on what you select for the first two filter fields, you can either enter a value or select a value from a drop-down list.
5. Click **Search**. The appliance lists all matching objects in the table. You can select multiple object types by clicking the Add icon to add more filter criteria. You can also click **Reset** to clear all entries.
6. Select the check boxes of the objects to which you are defining permissions, and then click the Select icon.
7. In the *Create Object Permissions* wizard, do the following:
 - **Object**: Displays the name of the selected object. When you select multiple objects, the appliance displays **Multiple** here. Mouse over to the information icon to view the list of objects to which you are defining permissions.
 - **Object Type**: Displays the object type of the selected object. When you select more than one object type, the appliance displays **Multiple** here.

- **Resource:** Displays the selected objects. When you select more than one object type, the appliance displays **Multiple Selected Objects** here. Mouse over to the information icon to view the list of objects to which you are defining permissions. Grant the resources an appropriate permission: **Read/Write**, **Read Only**, or **Deny**.

8. Save the configuration and click **Restart** if it appears at the top of the screen.

Grid Manager displays a warning message when the permissions you define here overlap with other permissions in the system. Click **See Conflicts** to view the overlapping permissions in the *Permissions Conflict* dialog box. For information, see [Applying Permissions and Managing Overlaps](#) on page 166.

You can also set permissions for specific objects from the objects themselves. For example, to define permissions for a particular Grid member, navigate to that Grid member and define its permissions.

To define the permissions of a specific object:

1. Navigate to the object. For example, to define permissions for a particular network, from the **Data Management** tab, select the **IPAM** tab → *network* check box, and then click the Edit icon.
2. In the editor, select the **Permissions** tab, and then do one of the following:
 - Click the Add icon to add permission to the object. In the **Admin Group/Role Selector** dialog box, select an admin group or role to which you want to assign the permission, and then click the Select icon.
 - Modify the permission and resource type of a selected admin group or role.
 - Select an admin group or role and click the Delete icon to delete it.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Defining DNS and DHCP Permissions for Grid Members

You can restrict certain admin groups or roles to perform specific DNS and DHCP tasks on specific Grid members by assigning the correct global and object permissions. You can use this feature to separate the DNS and DHCP administration on different Grid members. For example, you can create an admin group or role that can only create, modify, and delete DHCP ranges in a specific network on a specific member in the Grid. This admin group or role is restricted to the specified tasks on the selected Grid member. It cannot perform other DNS or DHCP tasks on this member, and it cannot perform the specified tasks on other Grid members.

For example, you can define permissions that allow admins to create, modify, and delete DHCP ranges in network 10.0.0.0/8 on Grid member “sales.infoblox.com” by granting read/write object permissions to all DHCP ranges, network 10.0.0./8, and member DHCP on sales.infoblox.com. Admins with these permissions can only add, modify, and delete DHCP ranges in network 10.0.0.0/8 on Grid member sales.infoblox.com. They cannot perform other DHCP or DNS tasks on the member, and they cannot perform these tasks on other Grid members.

For information about required permissions for specific DNS and DHCP tasks, see [Administrative Permissions for Common Tasks](#) on page 193.

You can define the following DNS and DHCP permissions for an admin group or role:

- **Grid DNS or Grid DHCP:** Admins with read/write permissions can manage any DNS or DHCP resources on any Grid members. They can also modify Grid DNS or Grid DHCP properties and any member DNS and member DHCP properties. Admins with read-only permissions can only view DNS or DHCP resources. They cannot modify any DNS or DHCP resources or restart related services.
- **Member DNS or Member DHCP:** Admins with read/write permissions can perform the defined DNS or DHCP tasks only on the specified Grid member, not any other members. They can also modify DNS or DHCP properties on the specified member. Admins with read-only permission cannot assign the Grid member to any DNS or DHCP resources.
- **Restart DNS or Restart DHCP on member:** Admins with read/write permissions can restart the DNS or DHCP service on the specified Grid member, not any other members. However, they cannot modify DNS or DHCP properties on the member. They can assign the specified Grid member to any DNS or DHCP resources, but they cannot assign any other Grid members to DNS or DHCP resources.

To specify member DNS and DHCP permissions, define DNS or DHCP permissions at the global or object level for an admin group or admin role, as described in [Defining Global Permissions](#) on page 161 and [Defining Object Permissions](#) on page 162. Ensure that you include the Grid member object to which you want to restrict DNS or DHCP administration.

You can also control whether the admins can modify DNS or DHCP properties on a member, as described in [Modifying Permissions on a Grid Member](#) on page 165.

Modifying Permissions on a Grid Member

Admins can perform different tasks on a Grid member based on the permissions they have. [Table 4.2](#) outlines the permissions and the tasks admins can perform on a Grid member:

Table 4.2 Member Permissions and Tasks

	Grid Member	Member DNS or DHCP Properties	Restart DNS or DHCP on Grid Member
Read/Write	<ul style="list-style-type: none"> • Modify member properties • Restart, reboot, and shutdown member • Modify member DNS and DHCP properties • Restart member DNS and DHCP services • Assign and un-assign member to DNS and DHCP objects 	<ul style="list-style-type: none"> • Modify member DNS or DHCP properties • Restart member DNS or DHCP service • Assign and un-assign member to DNS or DHCP objects 	<ul style="list-style-type: none"> • Restart member DNS or DHCP service • Assign and un-assign member to DNS or DHCP objects
Read-only	<ul style="list-style-type: none"> • View member DNS and DHCP properties 	<ul style="list-style-type: none"> • View member DNS or DHCP properties 	<ul style="list-style-type: none"> • N/A (You cannot define a read-only permission)
Deny	<ul style="list-style-type: none"> • Cannot modify member, DNS, and DHCP properties • Cannot restart related services • Cannot assign member to DNS and DHCP objects 	<ul style="list-style-type: none"> • Cannot modify member, DNS, and DHCP properties • Cannot restart related services • Cannot assign member to DNS and DHCP objects 	<ul style="list-style-type: none"> • Cannot modify member, DNS, and DHCP properties • Cannot restart related services • Cannot assign member to DNS and DHCP objects

After you add permissions to an admin group or role for a specific Grid member, you can modify the member permissions and resources. Note that when you modify the member permissions and resources, the appliance updates the permissions of the admin group or role accordingly.

To modify Grid member permissions:

1. From the **Data Management** tab, select the **DHCP** or **DNS** tab → **Members** tab → *Grid_member*, and then click the Edit icon.
2. In the *Member DHCP Properties* or *Member DNS Properties* editor, select the **Permissions** tab.
3. Click a permission in the Permissions table, select a different permission from the **Permissions** drop-down list or select a different resource from the **Resources** drop-down list. Note that when you select **Restart DNS** or **Restart DHCP**, the admins with this permission can only restart the DNS or DHCP service on the selected member. They cannot modify DNS or DHCP properties of this member.
4. Save the configuration. Note that the appliance automatically updates the permissions of the corresponding admin group or role in the **Administration** tab.

Applying Permissions and Managing Overlaps

When an admin tries to access an object, the appliance checks the permissions of the group to which the admin belongs. Because permissions at more specific levels override those set at a higher level, the appliance checks object permissions hierarchically—from the most to the least specific. In addition, if the admin group has permissions assigned directly to it and permissions inherited from its assigned roles, the appliance checks the permissions in the following order:

1. Permissions assigned directly to the admin group.
2. Permissions inherited from admin roles in the order they are listed in the **Roles** tab of the *Admin Group* editor.

For example, an admin from the DNS1 admin group tries to access the a1.test.com A record in the test.com zone in the Infoblox default view. The appliance first checks if the DNS1 admin group has a permission defined for the a1.test.com A record. If there is none, then the appliance checks the roles assigned to DNS1. If there is no permission defined for the a1.test.com A record, the appliance continues checking for permissions in the order listed in [Table 4.3](#). The appliance uses the first permission it finds.

Table 4.3 Permission Checking

The appliance checks object permissions from the most to the least specific, as listed.	For each object, the appliance checks permissions in the order listed.
1. a1.test.com A record	a. DNS1 admin group
2. A records in test.com	b. Role 1, Role 2, Role 3...
3. test.com	
4. All zones in the default view	
5. Default view	
6. All A records	
7. All zones	
8. All DNS views	

An admin group that is assigned multiple roles and permissions can have overlaps among the different permissions. As stated earlier, the appliance uses the first permission it finds and ignores the others. For example, as shown in [Table 4.4](#), if an admin group has read/write permission to all A records in the test.com zone and a role assigned to it is denied permission to test.com, the appliance provides read/write access to A records in the test.com zone, but denies access to the test.com zone and all its other resource records.

Table 4.4 Directly-Assigned Permissions and Roles

Permission assigned to the admin group	Read/Write to all A records in the test.com zone
Permission inherited from an admin role	Deny to the test.com zone
Effective permissions	Deny to the test.com zone Read/Write to all A records in test.com zone Deny to all other resource records in test.com zone

If the group has multiple roles, the appliance applies the permissions in the order the roles are listed. If there are overlaps in the permissions among the roles, the appliance uses the permission from the role that is listed first. For example, as shown in [Table 4.5](#), the first role assigned to the admin group has read-only permission to all A records in the test.com zone and the second role has read/write permission to the same records. The appliance applies the permission from the first admin role.

Table 4.5 Multiple Roles

Role 1 permission	Read-only to all A records in the test.com zone
Role 2 permission	Read/Write to all A records in test.com zone Read/Write to all MX records in test.com zone
Effective permissions	Deny to the test.com zone Read-only to all A records in the test.com zone Read/Write to all MX records in test.com zone

You can check for overlapped permissions when you add permissions to roles and to admin groups, and when you assign roles to an admin group. When you create a permission that overlaps with existing permissions, Grid Manager displays a warning message and the **See Conflicts** link on which you click to view the overlapped permissions. For information, see [Viewing Overlapping Permissions](#) on page 167. You can also use the quick filter **Overlaps** to filter overlapped permissions, the appliance lists permissions that overlap with other permissions. If you want to change the permission the appliance uses, you must change the order in which the roles are listed or change the permissions that are directly assigned to the admin group. For information, see [Creating Limited-Access Admin Groups](#) on page 156.

Viewing Overlapping Permissions

When you click **See Conflicts** to view overlapping permissions, Grid Manager displays the following information in the *Permission Overlap* dialog box:

- **Resource:** The name of the object or resource.
- **Type:** The object type.
- **Permission:** The permission granted. This can be Read/Write, Read-Only, or Deny.
- **Inherited From:** Indicates the source from which the permission is inherited.
- **Conflict Status:** Indicates whether the permission is being used or ignored. In a permission overlap, the group permission always overrides the role permission if both permissions are set at the same level (global or object). However, if the permissions are set at different levels, the permission at a more specific level overrides that set at a higher level.
- **Role/Group Name:** The name of the admin group or admin role.

You can click the arrow key next to the resource to view the permission that is being ignored in the overlap.

Managing Permissions

After you define permissions for an admin group and role, you can do the following:

- View the permissions, as described in [Viewing Permissions](#) on page 168.
- Modify the permissions, as described in [Modifying Permissions](#) on page 168.
- Delete the permission, as described in [Deleting Permissions](#) on page 168.

Viewing Permissions

Only superusers can view the permissions of all admin groups.

To view the permissions of an admin group or role:

1. From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab.
2. For an admin group: Select an admin group in the Groups table.
or
For an admin role: Select an admin role in the Roles table.
3. Grid Manager displays the following information in the Permissions table:
 - **Group/Role:** The name of the admin group or role.
 - **Permission Type:** The type of permissions. This can be Administrative Permissions, Named ACL Permissions, DHCP Permissions, DNS Permissions, File Distribution Permissions, Grid Permissions, IPAM Permissions, or Security Permissions.
 - **Resource:** The name of the object. For example, this field displays **All Hosts** if you have defined permissions for all the hosts in the Grid.
 - **Resource Type:** The object type. For example, this can be Host, PTR record, or Shared Network.
 - **Permission:** The defined permission for the resource.

When you click **Show All** for Admins, Groups, and Roles, Grid Manager displays all the admin accounts, admin groups, and admin roles in their respective tables.

Filtering the List of Permissions

You can filter the permissions you want to view by selecting one of the following from the quick filter menu:

- **Effective Permissions:** Select to view only the permissions that the appliance is using for this group. The permissions that were ignored due to overlaps are not listed in this view.
- **Overlaps:** Select to view only the overlapped permissions.
- **All Configured Permissions:** Select to view all permissions.

Modifying Permissions

You can modify the permissions of user-defined admin roles and admin groups. You cannot modify the permissions of system-defined admin roles. When you change the permissions of a role that has been assigned to multiple admin groups, the appliance automatically applies the change to the role in all admin groups to which it is assigned.

To modify the existing permissions of a role or an admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab.
2. For an admin group: Select an admin group in the Groups table.
or
For an admin role: Select an admin role in the Roles table.
3. In the Permissions table, select the resource that you want to modify, and then click the Edit icon.
4. In the *Mange Global Permissions* or *Create Object permissions* editor, select the new permission: **Read/Write**, **Read-Only** or **Deny** for the resource.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Deleting Permissions

You can remove permissions from user-defined admin roles and admin groups. You cannot remove permissions from system-defined admin roles. When you remove permissions from a role, they are removed from the role in all admin groups to which the role is assigned. You can remove a permission from a group as long as it is not inherited from a role. You cannot remove permissions that are inherited from a role.

To delete a permission:

1. From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab.
2. For an admin group: Select an admin group in the Groups table.
or
For an admin role: Select an admin role in the Roles table.
3. In the Permissions table, select the resource that you want to modify, and then click the Delete icon.
4. In the *Delete Permission Confirmation* dialog box, click **Yes**.

AUTHENTICATING ADMINISTRATORS

The NIOS appliance supports the following authentication methods: local database, RADIUS, Active Directory, LDAP, and TACACS+. The appliance can use any combination of these authentication methods. It authenticates admins against its local database by default. Therefore, if you want to use local authentication only, you must configure the admin groups and add the local admin accounts, as described in [Creating Local Admins](#) on page 169.

To authenticate admins using RADIUS, Active Directory, TACACS+, or LDAP in addition to local authentication, you must define those services on the appliance and define the admin authentication policy. For information, see [About Remote Admins](#) on page 171.

The appliance also supports two-factor authentication that validates smart card users, such as the U.S. Department of Defense (CAC) Common Access Card users. In two-factor authentication, NIOS first authenticates admins through the admin authentication policy, and then validates the admin client certificates through the OCSP service. For more information about two-factor authentication and how to configure it, see [Defining the Authentication Policy](#) on page 185.

Note: If you are using remote authentication, you must always have at least one local admin in a local admin group to ensure connectivity to the NIOS appliance in case the remote servers become unreachable.

CREATING LOCAL ADMINS

When you create an admin account, you must specify the name, password, and admin group of the administrator. You can also control in which time zone the appliance displays the time in the audit log and the DHCP and IPAM tabs of Grid Manager, such as the *DHCP Lease History* and *DHCP Leases* panels. The appliance can use the time zone that it automatically detects from the management system that the admin uses to log in. Alternatively, you can override the time zone auto-detection feature and specify the time zone.

To create an admin account and add it to an admin group:

1. Log in as a superuser.
2. From the **Administration** tab, select the **Administrators** tab -> **Admins** tab, and then click the Add icon.
or
From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin_group*, and then click the Add icon.
3. In the *Add Administrator Basic* wizard, complete the following:
 - **Login:** Enter a name for the administrator. This is the username that the administrator uses to log in.
 - **Password:** Enter a password for the administrator to use when logging in.
 - **Confirm Password:** Enter the same password.
 - **Email Address:** Enter the email address for this administrator. The appliance uses this email address to send scheduling notifications.

- **Admin Group:** Click **Select** to specify an admin group. If there are multiple admin groups, Grid Manager displays the *Admin Group Selector* dialog box from which you can select one. An admin can belong to only one admin group at a time.
NIOS appliance creates a new group, **fireeye-group**, when you add the first FireEye zone. The FireEye admin group is read-only and you cannot assign permissions to it. Select **fireeye-group** for the admin group and add users to this group. For more information, see [About FireEye Integrated RPZs](#) on page 1254.
 - **Comment:** Enter useful information about the administrator.
 - **Disable:** Select this check box to retain an inactive profile for this administrator in the configuration. For example, you might want to define a profile for a recently hired administrator who has not yet started work. Then when he or she does start, you simply need to clear this check box to activate the profile.
4. Optionally, click **Next** to add extensible attributes to the admin account. For information, see [About Extensible Attributes](#) on page 322.
 5. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing Passwords

Superusers can define requirements for the passwords of local admins according to your organization's policies. In addition to specifying the minimum password length, you can define rules that specify the character types that are allowed in the password. You can also specify whether passwords expire, their duration, and when reminders are sent to the users. Additionally, you can require admins to change their passwords when they first log in or after their passwords are reset.

You set the requirements at the Grid level, so they apply to all local admins who log in to the Grid. The requirements that you define appear in the User Profile of all local admins and when users are required to change their password.

To define the password requirements for local admins:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties**.
3. In the *Grid Properties* editor, select the **Password** tab and complete the following:
 - **Minimum Password Length:** Specify the minimum number of characters that are required in a password.
 - **Password Complexity:** Specify the following password requirements:
 - the minimum number of lowercase characters
 - the minimum number of uppercase characters
 - the minimum number of numeric characters
 - the minimum number of symbol characters. The allowed characters are: ! @ # \$ % ^ & * ()
 You must also specify how many characters an admin must change when revising a password.
 - **Password must expire:** Select this check box to enable passwords to expire after a specified period. Specify the duration of each password and the number of days before the expiration that the appliance sends a reminder.
 - **Force password change at next login:** Select this check box to force all new users to change their passwords when they first log in and to force existing users whose passwords were just reset to change their passwords.

Note: The “force password change at next login” feature does not apply to admin users in the **fireeye-group**. These users will not be prompted to change their passwords at the next login. Their original passwords continue to work. For information about FireEye integrated RPZs, see [About FireEye Integrated RPZs](#) on page 1254.

4. Click **Save & Close**.

Modifying and Deleting Admin Accounts

You can modify and delete admin accounts that you create, but you can only partially modify the default superuser account “admin”—and only when you are logged in as a superuser account. Furthermore, because there must always be a superuser account on the appliance, you can only remove the default “admin” account after you create another superuser account.

To modify an admin account:

1. From the **Administration** tab, select the **Administrators** tab -> **Admins** tab -> *admin_account* check box, and then click the Edit icon.
or
From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin_group* -> *admin_account* check box, and then click the Edit icon.
2. The *Administrator* editor provides the following tabs from which you can modify data:
 - **General:** In the **General Basic** tab, modify data of the admin account as described in [Creating Local Admins](#) on page 169.
In the **General Advanced** tab, complete the following:
 - **Time Zone:** Select a time zone from the drop-down list if you want to specify the time zone for the administrator. By default, the appliance automatically detects the time zone from the management system that the administrator uses to connect to the appliance. The appliance uses this time zone when it displays the timestamps for relevant data.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the admin account. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

To delete an admin account:

1. From the **Administration** tab, select the **Administrators** tab -> **Admins** tab -> *admin_account* check box, and then click the Delete icon.
or
From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin_group* -> *admin_account* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

ABOUT REMOTE ADMINS

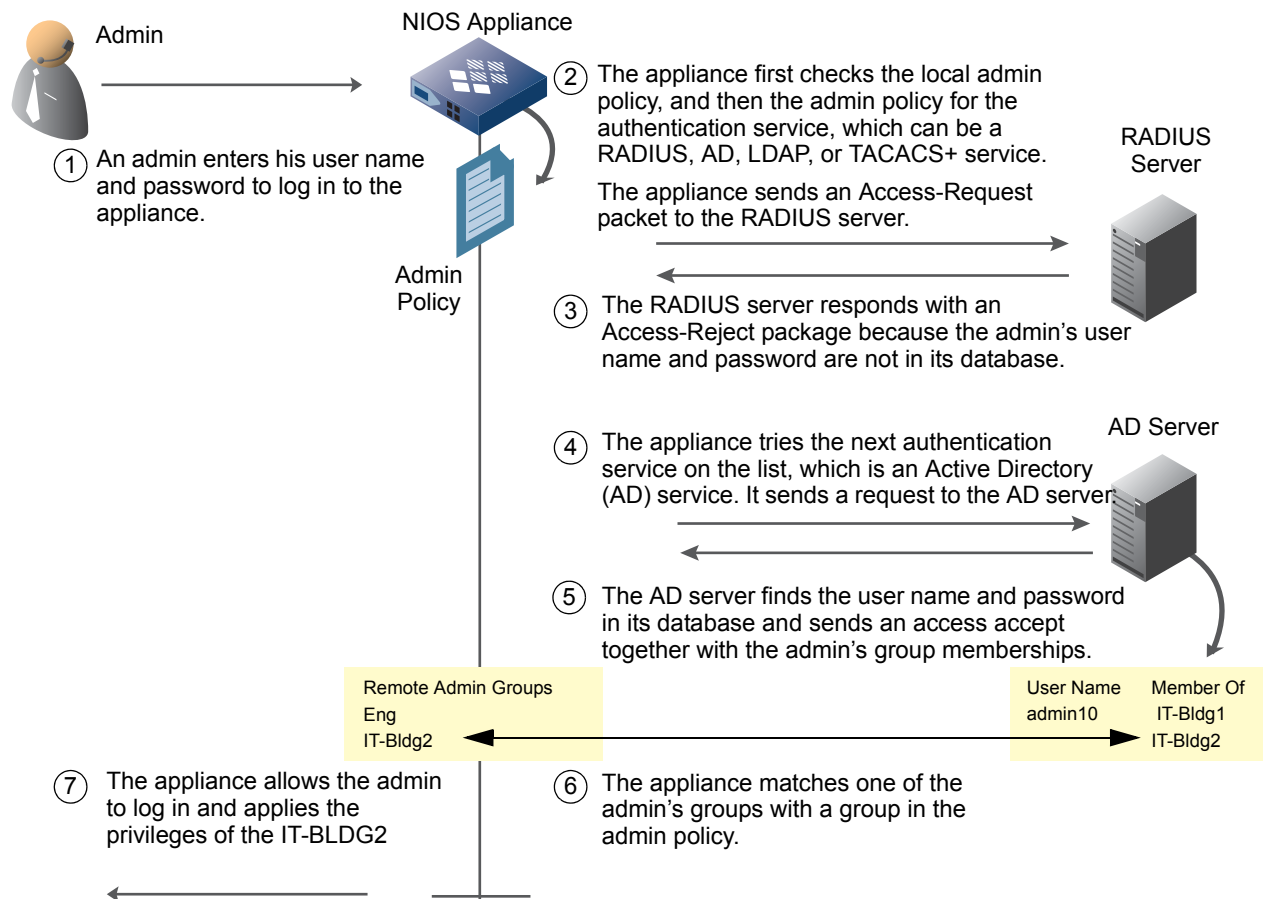
NIOS can authenticate admins whose user credentials are stored remotely on RADIUS servers, AD domain controllers, LDAP servers, or TACACS+ servers. You can configure authentication server groups for each type of server to which NIOS sends authentication requests. For example, you can create a server group for RADIUS servers and another server group for AD domain controllers. Then in the admin authentication policy, you can list which authentication server groups to use and in what order.

In addition, if admin groups are configured on the remote authentication server, you can configure admin groups with the same names on the NIOS appliance and list them in the authentication policy as well. Then if the remote authentication server provides the admin group name while authenticating an admin, NIOS can automatically assign the admin to the matching admin group specified in the authentication policy. You can also create a default admin group for all admins that are authenticated through a remote authentication service.

[Figure 4.4](#) illustrates the authentication and authorization process for remote admins. In the example, two authentication server groups are configured—a RADIUS server group and an AD server group. When an admin logs in with a user name and password, the appliance uses the first service listed in the admin policy to authenticate the admin. If authentication fails, the appliance tries the next service listed, and so on. It tries each service on the list until it is successful or all services fail. If all services fail, then the appliance denies access and generates an error message in syslog.

If authentication succeeds, the NIOS appliance tries to match the admin group names in the admin policy to any groups received from the remote server. If it finds a match, the NIOS appliance applies the privileges of that group to the admin and allows access. If the appliance does not find a match, then it applies the privileges of the default group. If no default group is defined, then the appliance denies access.

Figure 4.4 Authenticating Remote Admins



Only superusers can perform the following tasks to configure NIOS to authenticate admins using remote authentication servers:

- Configure the authentication server groups. You can create multiple RADIUS, LDAP, and AD server groups, but only one TACACS+ server group.
 - For information about RADIUS authentication, see [Authenticating Admins Using RADIUS](#) on page 173.
 - For information about AD authentication, see [Authenticating Admins Using Active Directory](#) on page 177.
 - For information about TACACS+ authentication, see [Authenticating Admin Accounts Using TACACS+](#) on page 179.
 - For information about LDAP authentication, see [Authenticating Admins Using LDAP](#) on page 182.
- Configure admin groups with names that match those on the remote server. For information about admin groups, see [About Admin Groups](#) on page 154.

- Configure the admin policy, as described in [Defining the Authentication Policy](#) on page 185.

Note: Infoblox strongly recommends that even if you are using remote authentication, you always have at least one local admin in a local admin group to ensure connectivity to the appliance in case the remote servers become unreachable. Also, when you delete an authentication server group, the appliance removes it from the system. Deleted authentication server groups are not moved to the Recycle Bin. Once deleted, the authentication server groups no longer exist in the system.

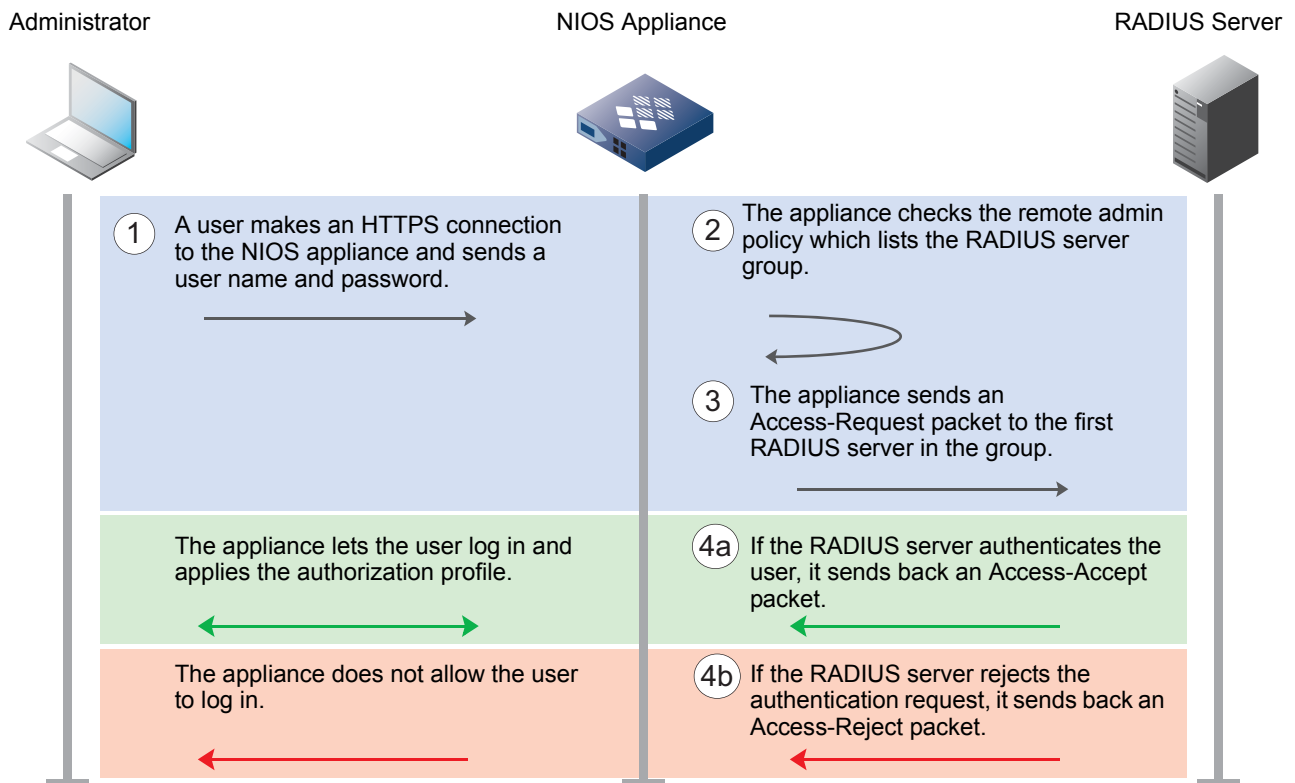
You can also authenticate users with smart cards that contain X.509 client certificates. The status of these certificates is stored remotely on OCSP responders. You can configure NIOS to authenticate these admins through the two-factor authentication method. For more information about two-factor authentication and how to configure it, see [Defining the Authentication Policy](#) on page 185.

AUTHENTICATING ADMINS USING RADIUS

RADIUS provides authentication, accounting, and authorization functions. The NIOS appliance supports authentication using the following RADIUS servers: FreeRADIUS, Microsoft, Cisco, and Funk.

When NIOS authenticates administrators against RADIUS servers, NIOS acts similarly to a network access server (NAS), which is a RADIUS client that sends authentication and accounting requests to a RADIUS server. [Figure 4.5](#) illustrates the RADIUS authentication process.

Figure 4.5 Authentication using a RADIUS server



Authentication Protocols

When you configure the NIOS appliance to authenticate admins against a RADIUS server group, you must specify the authentication protocol of each RADIUS server, which can be either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol).

PAP tries to establish the identity of a host using a two-way handshake. The client sends the user name and password in clear text to the NIOS appliance. The appliance uses a shared secret to encrypt the password and sends it to the RADIUS server in an Access-Request packet. The RADIUS server uses the shared secret to decrypt the password. If the decrypted password matches a password in its database, the user is successfully authenticated and allowed to log in.

With CHAP, when the client tries to log in, it sends its user name and password to the NIOS appliance. The appliance then creates an MD5 hash of the password together with a random number that the appliance generates. It then sends the random number, user name, and hash to the RADIUS server in an Access-Request package. The RADIUS server takes the password that matches the user name from its database and creates its own MD5 hash of the password and random number that it received. If the hash that the RADIUS server generates matches the hash that it received from the appliance, then the user is successfully authenticated and allowed to log in.

You can configure one of the following modes to send the authentication request to the RADIUS server:

- **Ordered:** In this mode, the authentication request is sent to the first server in the list. The authentication request is sent to the next server only when the first server is out of service or unavailable.
- **Round Robin:** In this mode, the first authentication request is sent to a server chosen randomly in a group. If there is no response from the server, continued attempts are performed sequentially until it selects the last server in the list. Then it starts with the first server in the list and continues the selection process until all the servers have been attempted.

Accounting Activities Using RADIUS

You can enable the accounting feature on the RADIUS server to track whether an administrator has initiated a session. After an administrator successfully logs in, the appliance sends an Accounting-Start packet to the RADIUS server.

Configuring Remote RADIUS Servers

For NIOS to communicate with a RADIUS server, you must also set up the remote RADIUS server to communicate with the NIOS appliance.

Note: If you have two Infoblox appliances in an HA pair, enter both the members of the HA pair as separate access appliances and use the LAN or MGMT IP address of both appliances (not the VIP address), if configured.

Depending on your particular RADIUS server, you can configure the following RADIUS server options to enable communication with the NIOS appliance:

- Authentication Port
- Accounting Port
- Domain Name/IP Address of the NIOS appliance
- Shared Secret Password
- Vendor Types

Configuring Admin Groups on the Remote RADIUS Server

Infoblox supports admin accounts on one or more RADIUS servers.

On the remote RADIUS server, do the following to set up admins and associate them with an admin group:

- Import Infoblox VSAs (vendor-specific attributes) to the dictionary file on the RADIUS server
- For third-party RADIUS servers, import the Infoblox vendor file (the Infoblox vendor ID is 7779)
- Define the admin group

- Associate one or more remote admin accounts with the admin group
- Add and activate a policy for the admin accounts, but do not associate the policy with a policy group that contains an infoblox-group-info attribute.

Refer to the documentation for your RADIUS server for more information.

Configuring RADIUS Authentication

To configure NIOS to use one or more RADIUS server groups to authenticate administrators, you must do the following:

- Configure at least one RADIUS authentication server group. For more information, see [Configuring a RADIUS Authentication Server Group](#) on page 175.
- Define admin groups for the admins that are authenticated by the RADIUS servers and specify their privileges and settings. The group names in NIOS must match the admin group names on the RADIUS server. See [About Admin Groups](#) on page 154 for information about defining admin groups.
- In the authentication policy, add the RADIUS server groups and the admin groups that match those on the RADIUS server. You can also designate an admin group as the default group for remote admins. NIOS assigns admins to this group when it does not find a matching group for a remote admin. See [Defining the Authentication Policy](#) on page 185 for more information about configuring the policy.

Configuring a RADIUS Authentication Server Group

You can add multiple RADIUS servers to the group for redundancy. When you do, the appliance tries to connect to the first RADIUS server on the list and if the server does not respond within the maximum retransmission limit, then it tries the next RADIUS server on the list. NIOS tries to connect to each RADIUS server in the order the servers are listed. If it does not receive a response within the configured timeout period and has tried to connect the specified retry value, then it tries the next RADIUS server on the list. It logs an error to syslog when it fails to connect to any of the servers in the group.

After you add a RADIUS server to the NIOS appliance, you can validate the configuration. The appliance uses a pre-defined username and password when it tests the connection to the RADIUS server. The pre-defined user name is “Infoblox_test_user” and the password is “Infoblox_test_password”. Do not use these as your administrator username and password.

To configure a RADIUS authentication server group:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the Add icon in the **RADIUS Services** subtab.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
 - **Name:** Enter the name of the server group.
 - **RADIUS Servers:** Click the Add icon and enter the following:
 - **Server Name or IP Address:** Enter the FQDN or the IP address of the RADIUS server that is used for authentication.
 - **Comment:** Enter additional information about the RADIUS server.
 - **Authentication Port:** The destination port on the RADIUS server. The default is 1812. This field is required only if you do not enable accounting on the RADIUS server. This field is not required if you enable accounting to configure an accounting-only RADIUS server.
 - **Authentication Type:** Select either PAP or CHAP from the drop-down list. The default is PAP.
 - **Shared Secret:** Enter the shared secret that the NIOS appliance and the RADIUS server use to encrypt and decrypt their messages. This shared secret is a value that is known only to the NIOS appliance and the RADIUS server.
 - **Enable Accounting:** Select this to enable RADIUS accounting for the server so you can track an administrator’s activities during a session. When you enable accounting, you must enter a valid port number in the **Accounting Port** field.

- **Accounting Port:** The destination port on the RADIUS server. The default is 1813.
- **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for administrator authentication communications with just this RADIUS server.
- **Disable server:** Select this to disable the RADIUS server if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server.
- Click **Test** to test the configuration. If the NIOS appliance connects to the RADIUS server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the RADIUS server, the appliance displays a message indicating an error in the configuration.
- Click **Add** to add the server to the list.

When you add multiple RADIUS servers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact a RADIUS server. You can move a server up or down the list by selecting it and clicking the up or down arrow.

You can also delete a RADIUS server by selecting it and clicking the Delete icon.

- **Authentication:** Optionally, modify the authentication settings. These settings apply to all RADIUS servers that you configure on the NIOS appliance.
 - **Timeout(s):** Specify the number of seconds that the appliance waits for a response from the RADIUS server.
 - **Retries:** Specify how many times the appliance attempts to contact an authentication RADIUS server. The default is 5.

If you have configured multiple RADIUS servers for authentication and the NIOS appliance fails to contact the first server in the list, it tries to contact the next server, and so on.

- **Accounting:** Optionally, modify the Accounting settings.
 - **Timeout(s):** Specify the number of seconds that the appliance waits for a response from the RADIUS server.
 - **Retries:** Specify how many times the appliance attempts to contact an accounting RADIUS server. The default is 1000.
- **Mode:** Specifies how the appliance contacts the RADIUS servers. The default is Ordered List.
 - **Ordered List:** The Grid member always selects the first RADIUS server in the list when it sends an authentication request. It queries the next server only when the first server is considered down.
 - **Round Robin:** The Grid member sends the first authentication request to a server chosen randomly in a group. If there is no response from the server, the Grid member selects the next server in the group. Continued attempts are performed sequentially until it selects the last server in the group. Then it starts with the first server in the group and continues the selection process until all the servers have been attempted.
- **Comment:** Enter useful information about the RADIUS service.
- **Disable:** Select this to disable RADIUS authentication for the servers listed in the table.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

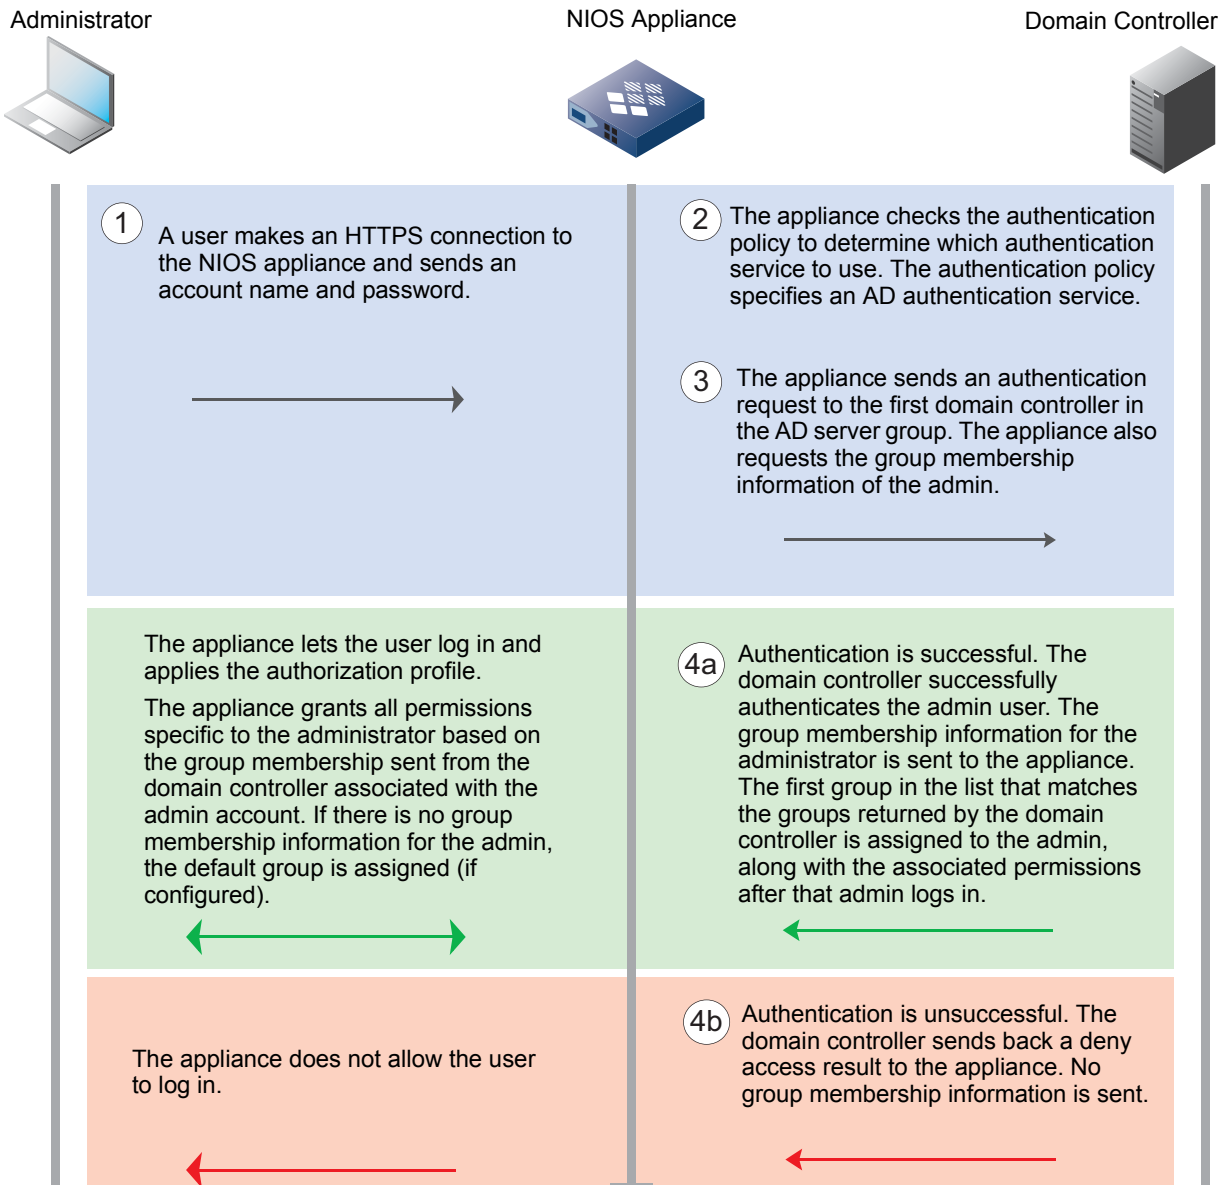
Note that the following fields in the wizard do not apply to this feature: **Enable NAC Filter**, **Cache Time to Live**, and **Recovery Interval**. They are used with the NAC Integration feature described in [Chapter 30, Authenticated DHCP](#), on page 915.

AUTHENTICATING ADMINS USING ACTIVE DIRECTORY

Active Directory™ (AD) is a distributed directory service that is a repository for user information. The NIOS appliance can authenticate admin accounts by verifying user names and passwords against Active Directory. In addition, the NIOS appliance queries the AD domain controller for the group membership information of the admin. The appliance matches the group names from the domain controller with the admin groups on its local database. It then authorizes services and grants the admin privileges, based upon the matching admin group on the appliance.

[Figure 4.6](#) illustrates the Active Directory authentication process.

Figure 4.6 Authentication Using a Domain Controller



To configure NIOS to authenticate administrators using Active Directory domain controller groups, you must first configure user accounts on the domain controller. Then, on the NIOS appliance, do the following:

- Configure one or more AD authentication server group on the appliance and add AD domain controllers to the group. For information about configuring an AD authentication service group for admins, see [Configuring an Active Directory Authentication Service Group](#) on page 178.

- If you configured admin groups on the AD controller, you must create those same groups on the NIOS appliance and specify their privileges and settings. Note that the admin group names must match those on the AD domain controller. You can specify a default group as well. The NIOS appliance assigns admins to the default group if none of the admin groups on the NIOS appliance match the admin groups on the AD domain controller or if there are no other admin groups configured. For information about configuring group permissions and privileges, see [About Admin Groups](#) on page 154.
- Add the newly configured Active Directory service to the list of authentication services in the admin policy, and add the admin group names as well. See [Defining the Authentication Policy](#) on page 185 for more information about configuring an admin policy.

Configuring an Active Directory Authentication Service Group

You can add multiple domain controllers to an AD authentication server group for redundancy. The NIOS appliance tries to connect with the first domain controller on the list. If it is unable to connect, it tries the next domain controller on the list, and so on.

To configure an Active Directory authentication server group on the NIOS appliance:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Active Directory Services** subtab and click the Add icon.
3. In the *Add Active Directory Authentication Service* wizard, complete the following:
 - **Name:** Enter a name for the service.
 - **Active Directory Domain:** Enter the AD domain name.
 - **Domain Controllers:** Click the Add icon and complete the following to add an AD domain controller:
 - **Server Name or IP Address:** Enter the FQDN or the IP address of the AD server that is used for authentication.
 - **Comment:** Enter additional information about the AD server.
 - **Authentication Port:** Enter the port number on the domain controller to which the appliance sends authentication requests. The default is 389.
 - **Encryption:** Select **SSL** from the drop-down list to transmit through an SSL (Secure Sockets Layer) tunnel. When you select SSL, the appliance automatically updates the authentication port to 636. Infoblox strongly recommends that you select this option to ensure the security of all communications between the NIOS appliance and the AD server. If you select this option, you must upload a CA certificate from the AD server. Click **CA Certificates** to upload the certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
 - **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for administrator authentication communications with just this AD server.
 - **Disable server:** Select this to disable an AD server if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server.
 - Click **Test** to test the configuration. If the NIOS appliance connects to the domain controller using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the server, the appliance displays a message indicating an error in the configuration.
 - Click **Add** to add the domain controller to the group.

When you add multiple domain controllers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact a domain controller. You can move a server up or down the list by selecting it and clicking the up or down arrow.

You can also delete a domain controller by selecting it and clicking the Delete icon.

- **Timeout(s):** The number of seconds that the NIOS appliance waits for a response from the specified authentication server. The default is 5.
 - **Comment:** Enter additional information about the service.
 - **Disable:** Select this to retain an inactive AD authentication service profile.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

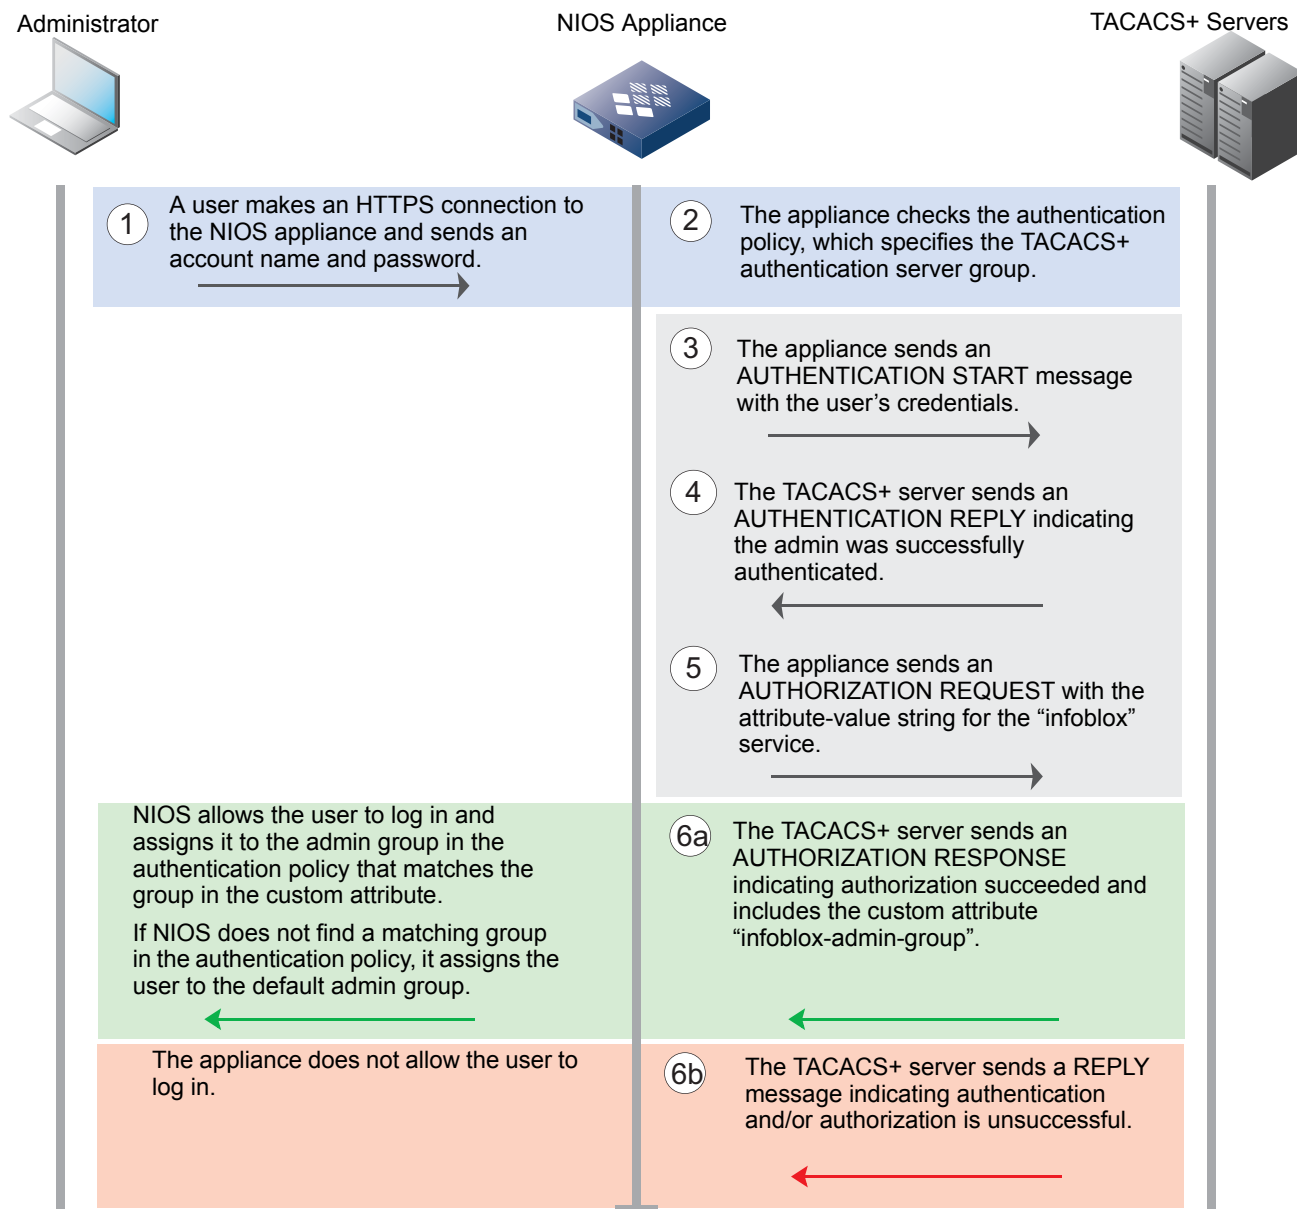
AUTHENTICATING ADMIN ACCOUNTS USING TACACS+

You can configure NIOS to authenticate admins against TACACS+ (Terminal Access Controller Access-Control System Plus) servers. TACACS+ provides separate authentication, authorization, and accounting services. To ensure reliable delivery, it uses TCP as its transport protocol, and to ensure confidentiality, all protocol exchanges between the TACACS+ server and its clients are encrypted. For detailed information about TACACS+, refer to the Internet draft <http://tools.ietf.org/html/draft-grant-tacacs-02>.

In addition, you can configure a custom service, infoblox, on the TACACS+ server, and then define a user group and specify the group name in the custom attribute infoblox-admin-group. Ensure that you apply the user group to the custom service infoblox. On NIOS, you define a group with the same name and add it to the authentication policy. Then when the TACACS+ server responds to an authentication and authorization request and includes the infoblox-admin-group attribute, NIOS can match the group name with the group in the authentication policy and automatically assign the admin to that group.

Figure 4.7 illustrates the TACACS+ authentication and authorization process when PAP/CHAP authentication is used.

Figure 4.7 TACACS+ Authentication



TACACS+ Accounting

When you enable TACACS+ accounting, NIOS sends the TACACS+ accounting server a TACACS+ accounting event with the same information that it sends to the Audit Log for any user command/event. NIOS sends an accounting start packet when a user first logs in successfully using TACACS+ authentication, and it sends an accounting STOP packet when a user logs out of the GUI or CLI or when a GUI or CLI session times out. If a product restarts or software failure occurs, NIOS drops any outstanding accounting packets. Note that audit log entries that are greater than 3,600 characters are truncated in accounting events sent to TACAS+ servers.

Configuring TACACS+

Complete the following tasks to enable NIOS and the TACACS+ servers to communicate.

On each TACACS+ server that you are adding to the authentication server group:

- For Windows TACACS+ servers, add the NIOS appliance as an AAA client. This step is not required for LINUX TACACS+ servers.
- Determine which user group on the TACACS+ server is used to match the admin group in NIOS, and then configure the following settings for the user group:
 - Add “infoblox” as a custom service.
 - Define the custom attribute for the group, in the format: **infoblox-admin-group=group_name**. For example, **infoblox-admin-group=remoteadmins1**. The group name can have a maximum of 64 characters.

On the NIOS appliance:

- Create a TACACS+ authentication server group. You can create only one TACACS+ server group. For more information, see [Configuring a TACACS+ Authentication Server Group](#) on page 180.
- Create the local admin group in NIOS that matches the user group on the TACACS+ server. Note that the NIOS admin group name must match the group name specified in the TACACS+ server and in the custom attribute. For example, if the custom attribute is **infoblox-admin-group=remoteadmins1**, then the admin group name must be **remoteadmins1**. In addition, you can designate a default admin group for remote admins. For information about configuring group permissions and privileges, see [About Admin Groups](#) on page 154.
- In the authentication policy, add the newly configured TACACS+ server group and the TACACS+ admin group name. See [Defining the Authentication Policy](#) on page 185 for more information about configuring an admin policy.

Configuring a TACACS+ Authentication Server Group

You can add multiple TACACS+ servers to the TACACS+ authentication server group. NIOS sends authentication requests to the TACACS+ servers in the order they are listed. NIOS sends authentication requests to the first server on the list. If that server is unreachable or generates an error, then NIOS sends the request to the next server in the list that has not been previously queried, and so on. NIOS logs an error message in syslog if all servers have been queried and they all generate errors or are unreachable.

To configure a TACACS+ authentication server group:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **TACACS+ Services** subtab and click the Add icon.
3. In the *Add TACACS+ Service* wizard, complete the following:
 - **Name:** Enter a name for the server group.
 - **TACACS+ Servers:** Click the Add icon and complete the following:
 - **Server Name or IP address:** The name or IP address of the TACACS+ server.
 - **Comment:** You can enter additional information about the server.
 - **Port:** The TCP destination port for TACACS+ communication. This port is used for authentication, accounting and authorization packets. The default is port 49.
 - **Authentication Type:** Select **ASCII**, **PAP** or **CHAP**. The default is **CHAP**.

- **Shared Secret:** The shared key that the NIOS appliance and the TACACS+ server use to encrypt and decrypt messages.
- **Enable Accounting:** Select this to enable NIOS to send accounting information to the TACACS+ server.
- **Connect through Management Interface:** Select this check box to enable the appliance to use the MGMT port to communicate with the TACACS+ server. Ensure that the MGMT port is configured. Otherwise, the appliance will use the LAN interface
- **Disable Server:** Select this to prevent queries from being sent to this server. You can retain the configuration, but disable the service.

Click **Test** to test the configuration. Click **Add** to add the TACACS+ server to the list.

When you add multiple TACACS+ servers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact a TACACS+ server. You can move a server up or down the list by selecting it and clicking the up or down arrow.

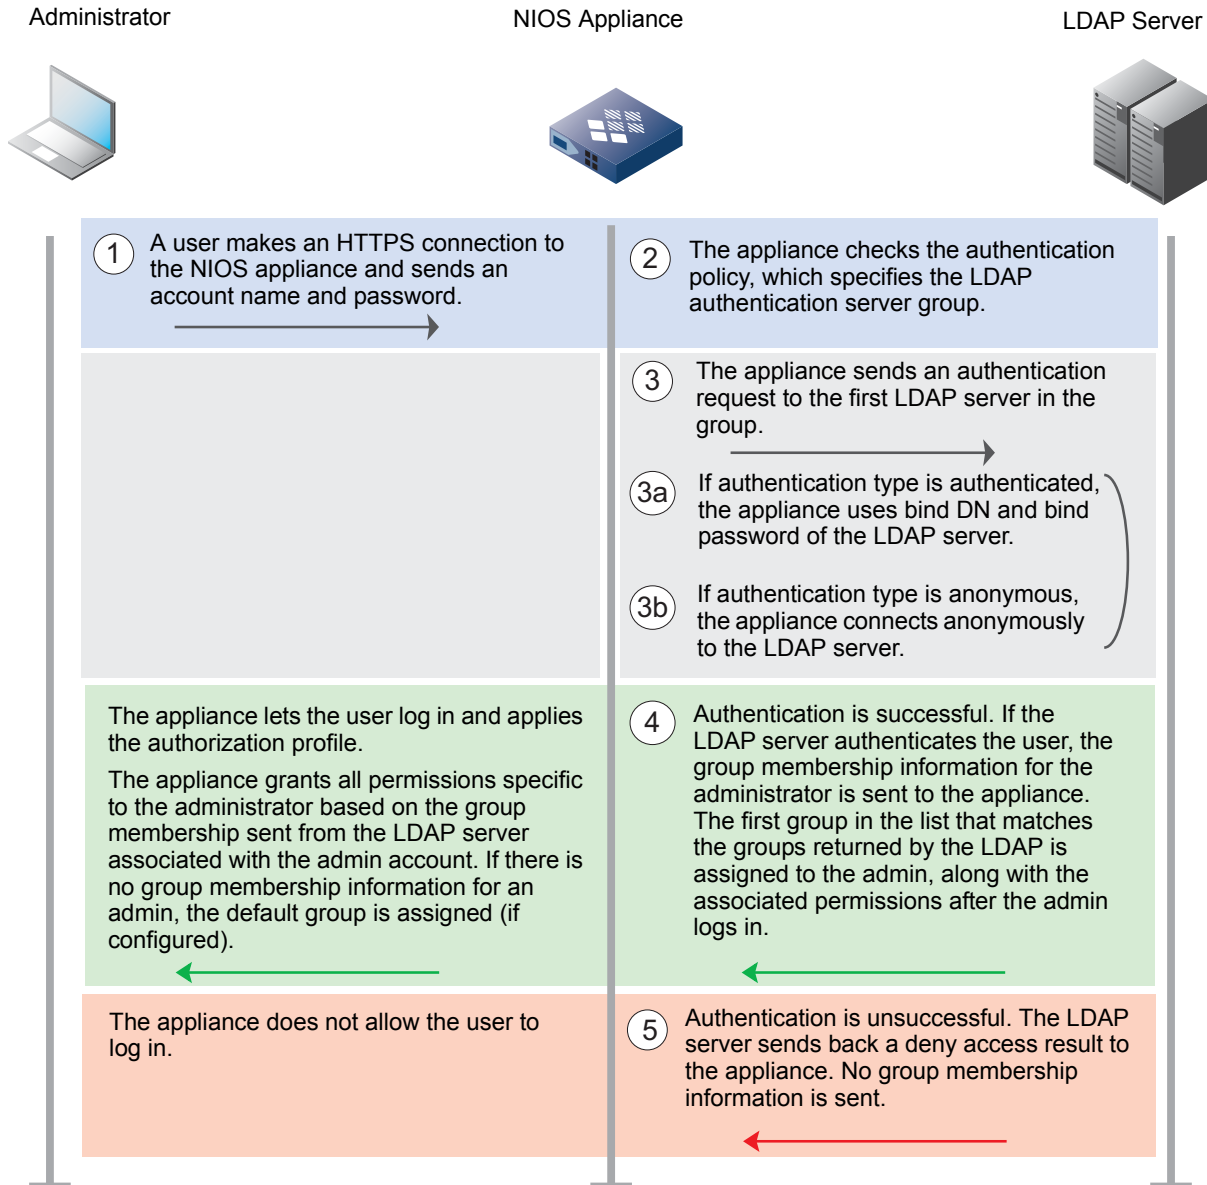
- **Authentication/Authorization:** Optionally, modify the authentication and authorization settings. These settings apply to all TACACS+ servers that you configure on the NIOS appliance.
 - **Timeout(s):** Specify the number of seconds or milliseconds that the appliance waits for a response from the TACACS+ server before it tries to contact it again. The amount of time before the server is retried. The default and minimum is 5000, and the maximum is 60000.
 - **Retries:** Specify how many times NIOS attempts to contact a TACACS+ server and fails before it tries to contact the next server on the list. The default is 0. The maximum is 5.
- **Accounting:** Optionally, modify the Accounting settings.
 - **Timeout(s):** Specify the number of seconds or milliseconds that the appliance waits for a response from the TACACS+ server. The amount of time before the server is retried. The default and minimum is 1000, and the maximum is 30000.
 - **Retries:** Specify how many times the appliance attempts to contact an accounting TACACS+ server and fails before it tries to contact the next accounting server on the list. The default is 0. The maximum is 5.
- **Comment:** Enter additional information about the service.
- **Disable:** Select this to retain an inactive TACACS+ authentication service profile.

4. Save the configuration.

AUTHENTICATING ADMINS USING LDAP

LDAP (Lightweight Directory Access Protocol) is an internet protocol for accessing distributed directory services. The NIOS appliance can authenticate admin accounts by verifying user names and passwords against LDAP. The NIOS appliance queries the LDAP server for the group membership information of the admin. The appliance matches the group names from the LDAP server with the admin groups in its local database. It then authorizes services and grants the admin privileges, based upon the matching admin group on the appliance. [Figure 4.8](#) illustrates the LDAP authentication process.

Figure 4.8 Authenticating using an LDAP server



Authentication Protocols

When you configure the NIOS appliance to authenticate admins against an LDAP server group, you must specify the authentication protocol of each LDAP server, which can be either anonymous or authenticated. The NIOS appliance connects anonymously to the LDAP server when the authentication type is anonymous. With authenticated type, the NIOS appliance connects using the bind DN and bind password defined for that server.

You can configure one of the following modes to send the authentication request to the LDAP server:

- **Ordered:** In this mode, the authentication request is sent to the first server in the list. The authentication request is sent to the next server only when the first server is out of service or unavailable.
- **Round Robin:** In this mode, the first authentication request is sent to a server chosen randomly in a group. If there is no response from the server, continued attempts are performed sequentially until it selects the last server in the list. Then it starts with the first server in the list and continues the selection process until all the servers have been attempted.

Configuring LDAP

Do the following to configure NIOS to use one or more LDAP server groups to authenticate administrators:

- Configure at least one LDAP authentication server group. For more information, see [Configuring an LDAP Server Group](#) on page 183.
- Define admin groups for the admins that are authenticated by the LDAP servers and specify their privileges and settings. The group names in NIOS must match the admin group names on the LDAP server. For more information about defining admin groups, see [About Admin Groups](#) on page 154.
- In the authentication policy, add the LDAP server groups and the admin groups that match those on the LDAP server. You can also designate an admin group as the default group for remote admins. NIOS assigns admins to this group when it does not find a matching group for a remote admin. For more information about configuring the policy, see [Defining the Authentication Policy](#) on page 185.

Configuring an LDAP Server Group

You can add one or more LDAP servers to an LDAP group for redundancy. The NIOS appliance tries to connect with the LDAP server based on the method you configure for the authentication request. If it does not receive a response within the configured timeout period and has tried to connect the specified retry value, then it tries the next LDAP server on the list. The appliance makes a syslog entry when it fails to connect to any of the servers in the group and sends an SNMP trap and an email notification if configured.

To configure an LDAP server group on the NIOS appliance:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the Add icon in the **LDAP Services** subtab.
3. In the *Add LDAP Authentication Service* wizard, complete the following:
 - **Name:** Enter the name of the server group.
 - **LDAP Servers:** Click the Add icon and enter the following:
 - **Server Name or IP Address:** Enter the FQDN (fully-qualified domain name) of the server or enter the IPv4/IPv6 address.
 - **LDAP Version:** Select the LDAP version. The NIOS appliance supports both LDAPv2 and LDAPv3. The default LDAP version is v3.
 - **Base DN:** Enter the base DN (Distinguished Name) value. All entries stored in an LDAP directory have a unique DN.
 - **Authentication Type:** Select the authentication type from the drop-down list. The supported authenticated types are as follows:
 - **Anonymous:** Select this to connect to the LDAP server anonymously. This is selected by default.
 - **Authenticated:** Select this to connect using the bind DN and bind password defined for that server.

- **Bind User DN:** Enter the bind user DN.
 - **Bind Password:** Enter the bind password.
 - **Encryption:** Select the encryption type from the drop-down list.
 - **SSL:** This is selected by default. All the network traffic is encrypted through an SSL (Secure Sockets Layer) protocol. The appliance automatically updates the authentication port to 636 for SSL. You must upload a CA certificate that verifies the LDAP server certificate. Click **CA Certificates** to upload the certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
 - **NONE:** Select this to unencrypt the connection. Note that Infoblox strongly recommends that you select the SSL option to ensure the security of all communications between the server and the member.
 - **Network Port:** Enter the authentication port number on the LDAP server to which the appliance sends authentication requests. The default value is 636. When you select NONE from the Encryption drop-down list, the appliance automatically updates the authentication port to 389.
 - **Comment:** Enter useful information about the LDAP server.
 - **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for administrator authentication communications with just this LDAP server.
 - **Disable Server:** Select this to disable the LDAP server if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server. You cannot disable the only server in a group if it is already being used by the remote authentication policy.
 - Click **Test** to test the configuration. If the NIOS appliance connects to the LDAP server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the server, the appliance displays a message indicating an error in the configuration.
 - Click **Add** to add the LDAP server to the group.
- When you add multiple LDAP servers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact an LDAP server. You can move a server up or down the list by selecting it and clicking the up or down arrow.
- You can also delete a server by selecting it and clicking the Delete icon.
- **Server Timeout(s):** Specify the number of seconds that the appliance waits for a response from the LDAP server. The default value is 5 seconds.
 - **Retries:** Specify how many times the appliance attempts to contact an authentication LDAP server. The default value is 5.
- If you have configured multiple LDAP servers for authentication and the NIOS appliance fails to contact the first server in the list, it tries to contact the next server after completing the specified number of attempts, and so on.
- **Mode:** Specifies the order in which a Grid member connects to an LDAP server.
 - **Ordered List:** The Grid member always selects the first LDAP server in the list when it sends an authentication request. It queries the next server only when the first server is considered down. This is the default.
 - **Round Robin:** The Grid member sends the first authentication request to a server chosen randomly in a group. If there is no response from the server, the Grid member selects the next server in the group. Continued attempts are performed sequentially until it selects the last server in the group. Then it starts with the first server in the group and continues the selection process until all the servers have been attempted.
 - **Recovery Interval:** Specify the number of seconds that the appliance waits to recover from the last failed attempt in connecting to an LDAP server. Select the time unit from the drop-down list. The default is 30 seconds. This is the time interval that NIOS waits before it tries to contact the server again since the last attempt when the appliance could not connect to the LDAP server or when the LDAP server did not send a reply within the configured response timeouts and retry attempts.

- **Group Membership Attribute:** Specify the LDAP group attribute (such as "memberOf" and "isMemberOf"). This is used to query the server and retrieve the group names to which the admin belongs. The default value is memberOf.
- **LDAP Search Scope:** To search for an admin user name in the LDAP directory, select one of the following LDAP search scope:
 - **Base:** Specify Base to perform search only on base in the LDAP directory. This is the top level of the LDAP directory tree.
 - **One Level:** Specify One Level to perform search on base DN and one level below the base in the LDAP directory.
 - **Sub tree:** Specify Sub tree to perform search on base and all the entries below the base DN in the LDAP directory.

The default value is One Level.

- **User ID:** Specify the attribute associated with the user object in the LDAP server, such as "uid" and "cn". This attribute is used to match the NIOS user name.
- **Map LDAP Field to Extensible Attribute (for Captive Portal Users only):** If you configure the LDAP authentication server group to authenticate the captive portal users, you can map an LDAP attribute value to an existing extensible attribute. This mapping is optional. By doing so, the LDAP attribute value will be queried from the LDAP server once the captive portal user authentication is successful. The attribute value received from the LDAP server is mapped to the corresponding extensible attribute. NIOS updates or creates a MAC address filter depending on the captive portal user or the client's hardware and name.

Click the Add icon and enter the following:

- **LDAP Field:** Enter the LDAP attribute. This attribute is queried in the LDAP directory server.
- **Extensible Attributes:** Select an attribute from the drop-down list. The drop-down list displays only the extensible attributes configured with attribute type as string. Infoblox recommends that you avoid confidential data while mapping extensible attribute to an LDAP attribute because this data is visible in the extensible attribute field of the corresponding MAC address filter.

Note: Mapping an extensible attribute to an LDAP attribute must be unique for a given LDAP server. Attribute not defined in the LDAP directory for a given user is considered as null and is mapped to the corresponding extensible attribute with a default value. The default value of extensible attribute is Not Found. This default value is not configurable and they do not cause the authentication to fail.

- **Comment:** Enter useful information about the LDAP server group.
- **Disable:** Select this to disable the LDAP authentication server group. Note that you cannot disable an LDAP group if it is already being used to authenticate one or more administrators and/or captive portal users.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

DEFINING THE AUTHENTICATION POLICY

The authentication policy defines which authentication server groups the appliance uses to authenticate admins and lists the local admin groups that map to the remote admin groups.

By default, the appliance provides the "Local Admin" service for authenticating users against the local database. You cannot modify or delete this default service.

Configuring a List of Authentication Server Groups

To enable NIOS to use multiple authentication server groups, define a prioritized list as follows:

1. From the **Administration** tab, select the **Administrators** tab -> **Authentication Policy** tab.

2. From the **Authenticate users against these services in this order** section, click the Add icon to add an authentication server group.
3. Select one of the following in the *Add Authentication Service* section:
 - **Active Directory:** Select this to add an AD authentication server group, and then select a group from the drop-down list.
 - **RADIUS:** Select this to add a RADIUS authentication server group, and then select a group from the drop-down list.
 - **TACACS+:** Select this to add the TACACS+ authentication server group, and then select a group from the drop-down list.
 - **LDAP:** Select this to add the LDAP authentication server group, and then select a group from the drop-down list.
4. Click **Add**.
 You can reorder the list by selecting an authentication server group and moving it up or down the list using the arrow keys.

Configuring a List of Remote Admin Groups

In order for NIOS to assign a remote admin to the correct group, you must list the admin groups in the local database that match the remote admin groups. You can also define a default admin group to which NIOS assigns remote users with no admin groups listed.

The appliance matches a remote admin to a group in the order the groups are listed. When the appliance receives information that an admin belongs to one or more groups, the appliance assigns the user to the first group in the list that matches. It assigns the admin to the default group, if specified, if no groups are returned by the authentication server, or if the appliance does not find a group in the local database that matches the group returned by the authentication server.

To configure the remote admin group list:

1. From the **Administration** tab, select the **Administrators** tab -> **Authentication Policy** tab.
2. From the **Map the remote admin group to the local group in this order** section, click the Add icon.
3. In the *Admin Group Selector* dialog box, select an admin group, and then click the Select icon. Use Shift+click and Ctrl+click to select multiple admin groups.

You can reorder the list by selecting an admin group and using the arrow keys to move it up or down the list.

To assign a user to a specific admin group if the remote admin group is not found, select **Assign User to this Group if Remote Admin Group cannot be found**, and then click **Select**. In the *Admin Group Selector* dialog box, select an admin group, and then click the Select icon.

AUTHENTICATING ADMINS USING TWO-FACTOR AUTHENTICATION

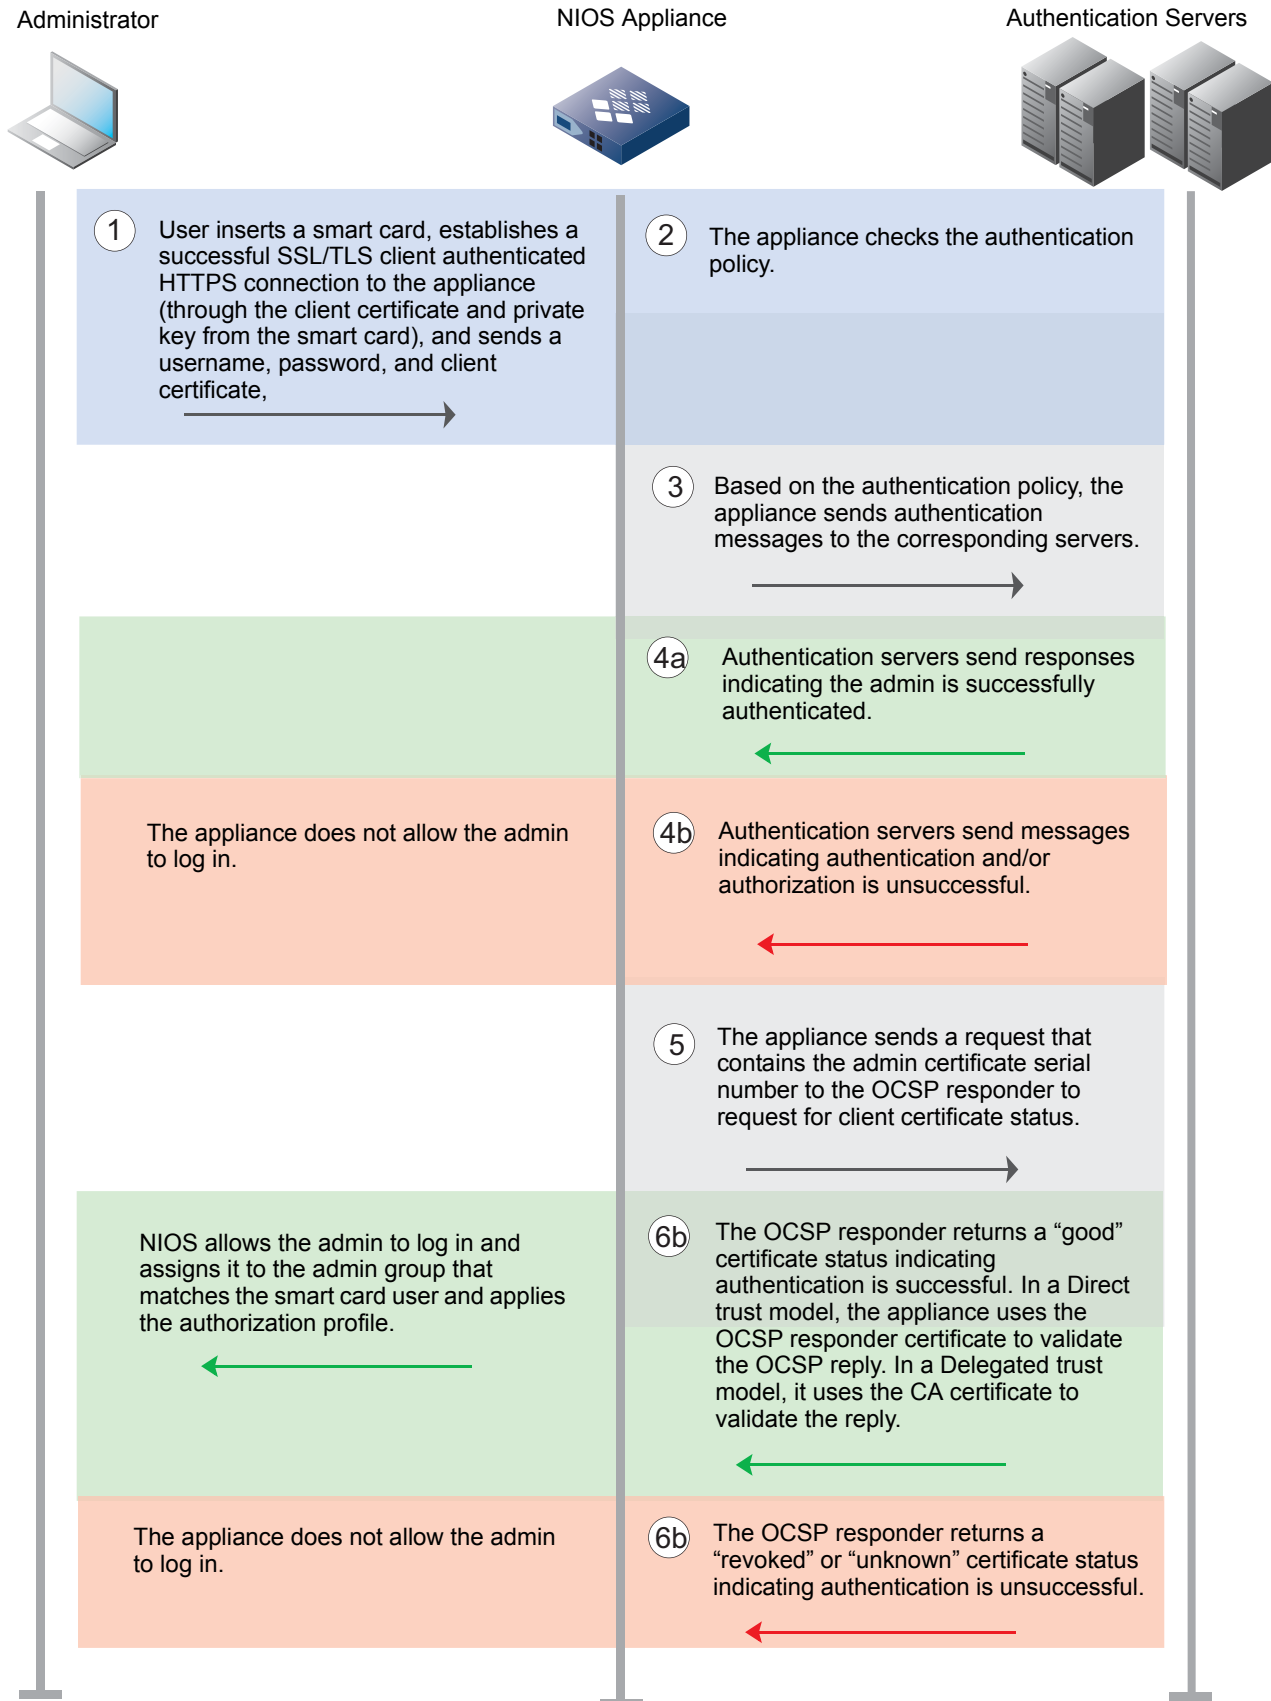
NIOS can authenticate users, such as U.S. Department of Defense CAC users, with smart cards that contain X.509 client certificates. The status of these certificates is stored remotely on OCSP responders. You can configure NIOS to use the two-factor authentication method to authenticate these users. In two-factor authentication, NIOS first negotiates SSL/TLS client authentication to validate client certificates. It then authenticates the admins based on the configured authentication policy. Finally, NIOS validates the status of the client certificate through the OCSP service. Note that you cannot add OCSP validation as part of the authentication policy. You must first configure the authentication policy, and then configure and enable the OCSP service for the two-factor authentication to take effect. For information about how to set up an authentication policy, see [Defining the Authentication Policy](#) on page 185.

OCSP is an internet protocol that validates certificate status for X.509 digital certificates that are assigned to specific admins. For more information about OCSP, refer to RFC 2560 at <http://tools.ietf.org/html/rfc2560>. The status of these client certificates is stored on OCSP responders to which NIOS sends requests about certificate status. A certificate status can be “good,” “revoked,” or “unknown.” After a successful SSL/TLS client authentication, NIOS authenticates the admin based on the configured authentication policy. If the authentication fails at this point, the appliance denies access to the admin. If the authentication policy has passed, the appliance sends a request to the OCSP responder for client certificate status about the admin. If the appliance receives a “good” status from the OCSP responder, the two-factor authentication is successful. The admin can now access the appliance. If the appliance receives a “revoked” or “unknown” status from the OCSP responder, the two-factor authentication fails. The admin cannot access the appliance even though the admin authentication policy has passed.

When there are multiple OCSP responders configured, the appliance contacts the responders based on their configured order. For the same client certificate, the appliance always takes the status reported by the first responder on the list that actually responds, even when there are different OCSP replies from different responders. When the appliance cannot contact the first responder or if the first responder does not reply, the appliance then takes the OCSP reply from the second responder and so on.

Note: Authentication for both the admin authentication policy and OCSP validation must be successful before a smart card admin can access the appliance.

Figure 4.7 illustrates the two-factor authentication and authorization process.



Best Practices for Configuring Two-Factor Authentication

Only superusers and limited-access users with the correct permissions can configure two-factor authentication. For information about admin roles and permissions, see [Managing Admin Groups and Admin Roles](#) on page 158. To configure two-factor authentication, consider the following:

- You must first set up an OSCP authentication server group and enable it.
- You can configure only one OSCP authentication server group that contains one or multiple OSCP responders to which NIOS sends requests about client certificate status. The appliance supports IPv4 and IPv6 OSCP responders.
- When you configure multiple OSCP responders, you can put them in an ordered list. The appliance contacts the first responder on the list. If the connection fails, it moves on to the second one, and so on. The result of the status check for a client certificate is based on the status reported by the first responder that replies.
- You can configure the timeout value and retry attempts that the appliance waits and tries before it moves on to the next OSCP responder.
- You can upload server certificates for each responder for OSCP response validation. You must upload an OSCP server certificate if you select the direct trust model.
- You can disable a specific responder if the server is out of service for a short period of time.
- Before you add an OSCP responder to the server group, you can test the server credentials.

To configure and enable two-factor authentication, complete the following tasks:

1. For local and remote authentication, ensure that the admin names for smart card users match the CNs (Common Names) used in the client certificates. For information about local and remote authentication, see [About Admin Accounts](#) on page 152.
2. Upload the CA (Certificate Authority) certificate, as described in [About CA Certificates](#) on page 56. The CA-signed certificates are used to validate OSCP server certificates and admin OSCP client certificates. Ensure that the CA certificate is in .PEM format. The .PEM file can contain more than one certificate.

Note: The uploaded CA certificates must be the ones that issued the client certificates to be authenticated. Otherwise, clients such as browsers, cannot establish a successful SSL/TLS client authenticated HTTPS session to the appliance.

3. Configure an OSCP authentication server group and enable it, as described in [Configuring the OSCP Authentication Server Group](#) on page 189.

Note that once you save the OSCP authentication server group configuration, the appliance terminates administrative sessions for all admin users. After you enable the OSCP service, you can verify whether two-factor authentication is enabled. Go to the **Administration** -> **Administrators** -> **Authentication Policy** tab, Grid Manager displays the “Two-Factor Authentication Enabled” banner in this tab.

Configuring the OSCP Authentication Server Group

To configure and enable the OSCP authentication service, complete the following:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **OCSP Services** subtab and click the Add icon.
3. In the *Add OSCP Service* wizard, complete the following:
 - **Name:** Enter a name for the service.
 - **OCSP Responders:** Click the Add icon and complete the following in the Add OSCP Responder section:
 - **Server Name or IP Address:** Enter the FQDN or the IP address of the OSCP responder that is used for authentication. The appliance supports IPv4 and IPv6 OSCP responders.
 - **Comment:** Enter useful information about the OSCP responder.
 - **Port:** Enter the port number on the OSCP responder to which the appliance sends authentication requests. The default is 80.

- **Server Certificate:** Click **Select** to upload a server certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**. The appliance validates the certificate when you save the configuration. A server certificate is required for the direct trust model.
- **Disable:** Select this check box to disable the OCSP responder if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server.

Note: You cannot save the OCSP configuration when you disable all OCSP responders, thus the OCSP service is disabled and two-factor authentication is no longer in effect.

Click **Add** to save the configuration and add the responder to the table. You can add multiple OCSP responders for failover purposes. You can use the up and down arrows to place the responders in the order you desire. The appliance tries to connect with the first responder on the list. If the connection fails, it tries the next responder on the list, and so on. Grid Manager displays the following for each responder:

- **Responder:** The FQDN or the IP address of the OCSP responder.
- **Comment:** Information you entered about the OCSP responder.
- **Port:** The port number on the OCSP responder to which the appliance sends authentication requests.
- **Disable:** Indicates whether the responder is disabled or not. Note that you must enable at least one responder to enable the OCSP service.

You can also click **Test** to test the configuration. If the appliance connects to the responder using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the responder, the appliance displays a message indicating an error in the configuration.

- **Response Timeout(s):** Enter the time the appliance waits for a response from the specified OCSP responder. The default is 1 second. You can select the time unit from the drop-down list.
- **Retries:** Enter the number of times the appliance tries to connect to the responders after a failed attempt. The default is 5.
- **Recovery Interval:** Enter the time the appliance waits to recover from the last failed attempt in connecting to an OCSP responder. Select the time unit from the drop-down list. The default is 30 seconds. This is the time interval that NIOS waits before it tries to contact the responder again since the last attempt when the appliance could not connect with the responder or when the responder did not send a reply within the configured response timeouts and retry attempts.
- **Trust Model:** From the drop-down list, select **Direct** or **Delegated** as the trust model for OCSP responses. In a direct trust model, OCSP responses are signed with an explicitly trusted OCSP responder certificate. You must upload the OCSP responder certificate if you select **Direct**. In a delegated trust model, OCSP responses are signed with a trusted CA certificate. A server certificate is not required when you select **Delegated**. The default is **Direct**.
- **SSH Remote Console Authentication:** Select a method for admins to log in to the appliance through the CLI. From the drop-down list, select **No Login** to disable SSH remote connection through CLI, or select **Login no Certificate** to authenticate admins using their user names and passwords without authenticating their client certificates through the OCSP service. The default is **No Login**.

Note: To enable SSH remote console authentication, you must also enable remote console access in the Grid or member settings.

- **Enable Superuser login when all responders are unavailable:** Select this check box to enable superuser login when all OCSP responders are unavailable. As long as the superusers are authenticated through the configured authentication policy, enabling this allows superusers to log in to the appliance if all OCSP responders were disconnected or did not reply within the configured response timeouts and retry attempts.
- **Comment:** Enter useful information about the OCSP authentication service.
- **Disable:** Select this to retain an inactive OCSP authentication service profile.

Note that enabling the OCSF authentication service terminates administrative services for all users. Ensure that you have uploaded the correct CA certificates before enabling the service. Your login names must also match the CN (Common Name) used in the certificate. When you configure multiple OCSF responders, ensure that you place them in the correct order because the status check for a client certificate is based on the OCSF reply sent by the first OCSF responder that replies.

Viewing the OCSF Authentication Server Group

To view the OCSF authentication server group, complete the following:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Grid Manager displays the following about the OCSF authentication server group:
 - **Name:** The name of the OCSF server group.
 - **Comment:** Comments about the OCSF server group

You can also display the following column:

- **Disabled:** Indicates if the OCSF server group is enabled or disabled.

You can do the following in this tab:

- Sort the data in ascending or descending order by column.
- Select the OCSF server group and click the Edit icon to modify data, or click the Delete icon to delete it.
- Print and export the data in this tab.
- Create a bookmark for this page.

CHANGING PASSWORD LENGTH REQUIREMENTS

Password length requirements control how long a password must be for a NIOS appliance admin account. Increasing this value reduces the likelihood of hackers gaining unauthorized access.

To change password length requirements:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Password** tab.
3. Enter a number from 4 to 64 in the **Minimum Password Length** field.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

NOTIFYING ADMINISTRATORS

You can notify individual administrators about system status through email, or notify a group of people using an alias email address. If you have configured DNS resolution on your network, the **E-mail relay configuration** function is not required. If you did not configure the settings on the **DNS Resolver** section, you must enter a static IP address of the target system in the **Relay Name/IP Address** field. The appliance sends e-mail to administrators when certain events occur. Here is a list of events that trigger e-mail notifications:

- Changes to link status on ports and online/offline replication status
- Events that generate traps, except for upgrade failures (ibUpgradeFailure). For a list of events, see [Infoblox MIBs](#) on page 1051

The appliance attempts to send the email notification once after an event. It does not try to send the notification again, if the first attempt fails. Infoblox recommends that you use the **Test Email settings** button to test the email settings and to verify that the recipient received the notification.

You can define the email settings at the Grid and member levels.

Grid Level

To notify an administrator of an independent appliance or a Grid:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Email** tab, and then complete the following:
 - **Enable Email notification:** Select this.
 - **Email address:** Enter the email address of the administrator. Use an email alias to notify multiple people.
 - **Use SMTP Relay:** Select this if the NIOS appliance must send email to an intermediary SMTP (Simple Mail Transfer Protocol) server that relays it to the SMTP server responsible for the domain name specified in the email address. Some SMTP servers only accept email from certain other SMTP servers and might not allow email from the NIOS appliance. In this case, specify the DNS name or IP address of a different SMTP server that does accept email from the NIOS appliance and that will then relay it to the SMTP server that can deliver it to its destination.

Clear this if it is unnecessary to use an email relay server.

 - **SMTP Relay Name or Address:** If you have configured DNS resolution, enter the DNS name of the relay server.

If DNS resolution is not configured, enter the IP address of the relay server.
3. Optionally, click **Test Email settings** to confirm this feature is operating properly.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Member Level

To define email settings for a member:

1. From the **Grid** tab, select the **Grid Manager** tab -> *member* check box, and then select the Edit icon.
2. In the *Grid Member Properties* editor, select the **Email** tab, and then click **Override** to override Grid-level settings.
3. Complete the email configuration as described in [Grid Level](#) on page 192.

ADMINISTRATIVE PERMISSIONS FOR COMMON TASKS

[Table 4.6](#) lists some of the common tasks admins can perform and their required permissions.

All the permission tables in this chapter use the following definitions:

- **RW** = Read/Write permission
- **RO** = Read-only permission

Table 4.6 Permissions for Common Tasks

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery
For Grid and Members																			
Restart services for the entire Grid	RO																		
Configure Grid DNS properties, configure member DNS properties, assign members to DNS objects, and restart DNS service on members		RW																	
Configure Grid DHCP properties, configure member DHCP properties, assign members to DHCP objects, and restart DHCP service on members			RW																
Configure a Grid member				RW															
Restart services on a Grid member				RW															
Configure member DNS properties, assign member to DNS objects, and restart DNS service on member					RW														
Configure member DHCP properties, assign member to DHCP objects, and restart DHCP service on member						RW													
Restart member DNS service							RW												
Restart member DHCP service								RW											
Initiate and control network discovery on all networks														RW					RW
Scheduling tasks for all supported objects									RW				RW					RW	

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery
For DNS resources																			
Create, modify, and delete DNS views									RW										
View and search for DNS views									RO										
Create, modify, and delete DNS zones with assigned members	RW									RW									
View and search for DNS zones with assigned members	RO									RO									
Create, modify, and delete all resource records	RW											RW							
View and search for all resource records	RO											RO							
Assign member to DNS objects							RW												
For DHCP Resources																			
Create, modify, and delete network views and their associated DNS views	RW								RW				RW						
View network properties and statistics	RO												RO						
Create, modify, and delete networks with assigned members				RW										RW					
Create, modify, and delete networks without assigned members															RW				
Create, modify, and delete DHCP ranges in a specific network with assigned members				RW											RW	RW			
Create, modify, and delete fixed addresses in a specific network without assigned members															RW		RW		
Assign member to DHCP objects								RW											

ADMINISTRATIVE PERMISSION FOR THE GRID

By default, the Grid Master denies access to Grid members when a limited-access admin group does not have defined permissions. You can grant an admin group read-only or read/write permission, or deny access to all Grid members or you can grant permission to specific Grid members, as described in [Applying Permissions and Managing Overlaps](#) on page 166.

Note: Only superusers can modify DNS and DHCP Grid properties.

The following sections describe the types of permissions that you can set with Grid permissions:

- [Administrative Permissions for Grid Members](#) on page 195
- [Administrative Permissions for Network Discovery](#) on page 196
- [Administrative Permissions for Scheduling Tasks](#) on page 196
- [Administrative Permissions for Microsoft Servers](#) on page 197

Administrative Permissions for Grid Members

[Table 4.7](#) lists the tasks admins can perform and the required permissions for Grid members.

Table 4.7 rid Member Permissions

Tasks	Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	DNS Views	DNS Zones	Networks	DHCP Ranges
Assign member to DNS zones				RW			RW		
Assign member to networks					RW			RW	
Assign member to DHCP ranges									RW
Configure member properties	RW								
Add a member to a Match Members list of a DNS view	RW								
Delete a view with members in a Match Members list						RO			
View DNS and DHCP member properties		RO	RO						
View and download syslog	RO								
View DNS and DHCP configuration file		RO	RO						
View network statistics	RO								
Restart DNS service on the member				RW					
Restart DHCP service on the member					RW				

Administrative Permissions for Network Discovery

Limited-access admin groups can initiate a discovery and manage discovered data based on their administrative permissions.

You can set global permissions for network discovery as described in [Defining Global Permissions](#) on page 161. The following table lists the tasks admins can perform and the required permissions for network discovery.

Table 4.8 Permissions for Network Discovery

Tasks	Network Discovery	DNS Zones	Networks Selected for Discovery
Initiate and control a discovery on selected networks	RW		RW
View discovered data			RO
Add unmanaged data to existing hosts, and resolve conflicting IP addresses			RW
Convert unmanaged data to a host, fixed address, reservation, A record, or PTR record		RW	RW

Administrative Permissions for Scheduling Tasks

You can schedule tasks, such as adding hosts or modifying fixed addresses, for a future date and time. To schedule tasks, you must first enable the scheduling feature at the Grid level, and then define administrative permissions for admin groups and admin roles. For information, see [About Extensible Attributes](#) on page 322. Only superusers can enable and disable this feature and grant scheduling permissions to admin groups. Limited-access admin groups can schedule tasks only when they have scheduling permissions.

Superusers can do the following:

- Enable and disable task scheduling at the Grid level
- Grant and deny scheduling permissions to admin groups and admin roles
- Schedule tasks for all supported object types
- Reschedule and delete any scheduled task

You can set global permissions to schedule tasks as described in [Defining Global Permissions](#) on page 161. The following table lists the tasks admins can perform and the required permissions. Users with read/write permission to scheduling can view, reschedule, and delete their own scheduled tasks.

Table 4.9 Scheduling Task Permissions

Tasks	Scheduling Task	All Networks	All DNS Views	All Shared Record Groups
Schedule the addition, modification, and deletion of all supported object types	RW	RW	RW	RW
View, reschedule, and delete scheduled tasks	RW	RW	RW	RW
Convert unmanaged data to a host, fixed address, reservation, A record, or PTR record	RW	RW	RW	

To schedule tasks for specific resources, admins must have Read/Write permission to scheduling tasks, plus the required permissions to the supported resources. For information about permissions for specific resources, see the following:

- Grid members—See [Administrative Permission for the Grid](#) on page 195.
- DNS resources—See [Administrative Permissions for DNS Resources](#) on page 199.
- DHCP resources—See [Administrative Permissions for DHCP Resources](#) on page 205.

Note that the appliance deletes all pending scheduled tasks when superusers disable task scheduling at the Grid level. The appliance deletes an admin's scheduled tasks when superusers do the following:

- Set the scheduling permission of admin groups and roles to “Deny”
- Delete or disable an admin group or an admin role
- Delete or disable local admins
- Delete the scheduling permission from any admin group or admin role that contains users with pending scheduled tasks
- Change the admin group of a limited-access admin

Administrative Permissions for Microsoft Servers

By default, only superusers can add Microsoft servers as managed servers to the Grid. Limited-access admins can add and manage Microsoft servers from the Grid based on their administrative permissions.

The following table lists the tasks admins can perform and the required permissions. Note that only superusers can add a Microsoft server to a name server group.

Table 4.10 Microsoft Server Permissions

Tasks	Microsoft Server(s)	Grid Member(s)	Network Views	DNS Views	DNS Zones	Resource Records	Networks	DHCP Ranges	Superscopes
Assign Microsoft server to member	RW	RW							
Assign a network view to the Microsoft server	RW	RW	RW						
Assign a DNS view to the Microsoft server	RW	RW		RW					
Assign Microsoft server as primary or secondary to DNS zones	RW			RW	RW				
Remove a Microsoft server as the primary or secondary server of a zone					RW				
Remove a zone from a Microsoft server					RW				
Edit zones and resource records of Microsoft servers					RW	RW			
Assign a Microsoft server to a network	RW						RW		
Assign a Microsoft server to a DHCP range	RW							RW	
Remove a network served by a Microsoft server	RW						RW		
Remove a DHCP range (scope) from a Microsoft server							RW	RW	
Add, modify and remove Microsoft superscopes	RW							RW	RW
Clear leases from Microsoft server	RW							RW	
Edit Microsoft server properties	RW								
View Microsoft server properties	RO								

Tasks	Microsoft Server(s)	Grid Member(s)	Network Views	DNS Views	DNS Zones	Resource Records	Networks	DHCP Ranges	Superscopes
View and download Microsoft logs	RO								
Start/Stop DNS or DHCP on the Microsoft server	RW								
Remove a Microsoft server from the Grid	RW								

ADMINISTRATIVE PERMISSIONS FOR IPAM RESOURCES

Limited-access admin groups can access certain IPAM resources only if their administrative permissions are defined. By default, the appliance denies access when a limited-access admin group does not have defined permissions. You can grant admin groups read-only or read/write permission, or deny access to the following IPAM resources:

- Network views
- IPv4 networks
- IPv6 networks
- Hosts

The appliance applies permissions for IPAM resources hierarchically. Permissions to a network view apply to all networks and resources in that view. You can also grant an admin group broad permissions to IPAM resources, such as read/write permission to all IPv4 networks and IPv6 networks in the database. In addition, you can grant permission to a specific host in a network. Permissions at more specific levels override global permissions.

The following sections describe the types of permissions that you can set for IPAM resources:

- [Administrative Permissions for Network Views](#) on page 206
- [Administrative Permissions for IPv4 and IPv6 Networks](#) on page 198
- [Administrative Permissions for Hosts](#) on page 199

Administrative Permissions for IPv4 and IPv6 Networks

Limited-access admin groups can access IPv4 and IPv6 networks only if their administrative permissions are defined. Permissions for a network apply to all its DNS and DHCP resources, if configured. To override network-level permissions, you must define permissions for specific objects within the networks. You can also define permissions for specific DHCP objects and Grid member to restrict admins to perform only the specified DHCP tasks on the specified member. For more information, see [Defining DNS and DHCP Permissions for Grid Members](#) on page 164.

You can grant read-only or read/write permission, or deny access to networks, as follows:

- All IPv4 or IPv6 networks—Global permission that applies to all networks in the database.
- A specific network—Network permissions apply to all objects in the network. This overrides global permissions.
- A specific network on a specific member—Network permissions apply to all objects in the network and member permissions apply to the specific member. For information about member permissions, see [Modifying Permissions on a Grid Member](#) on page 165.

Administrative Permissions for Hosts

A host record can contain both DNS and DHCP attributes if you configure them. When applying administrative permissions to host records, the permissions apply to all relevant DNS and DHCP resources within the host records. You can define global permissions to all hosts. To override global permissions, you must define permissions for specific hosts.

You can grant read-only or read/write permission, or deny access to host records, as follows:

- All hosts—Global permission that applies to all host records in the Grid.
- A specific host—Object permission that applies only to a selected host.

ADMINISTRATIVE PERMISSIONS FOR DNS RESOURCES

You can grant roles and admin groups read-only or read/write permission, or deny access to the following DNS resources:

- DNS Views
- DNS Zones
- Response Policy Zones
- All RPZ Rules
- Hosts
- Bulk Hosts
- A records
- AAAA records
- CNAME records
- DNAME records
- MX records
- PTR records
- SRV records
- TXT records
- Hosts
- Bulk Hosts
- Shared Record Groups
- Shared A records
- Shared AAAA records
- Shared MX records
- Shared SRV records
- Shared TXT records
- DNS64 synthesis groups

The appliance applies permissions for DNS resources hierarchically. Permissions to a DNS view apply to all zones and resource records in that view. Permissions for a zone apply to all its subzones and resource records, and resource record permissions apply to those resource records only. To override permissions set at higher level, you must define permissions at a more specific level. To assign permissions, see [Applying Permissions and Managing Overlaps](#) on page 166.

You can also define permissions for specific DNS objects and Grid member to restrict admins to perform only the specified DNS tasks on the specified member. For more information, see [Defining DNS and DHCP Permissions for Grid Members](#) on page 164.

The following sections describe the different types of permissions that you can set for DNS resources:

- [Administrative Permissions for DNS Views](#) on page 200
- [Administrative Permissions for Zones](#) on page 201
- [Administrative Permissions for Resource Records](#) on page 202

Administrative Permissions for DNS Views

Limited-access admin groups can access DNS views, including the default view, only if their administrative permissions are defined. Permissions to a DNS view apply to all its zones and resource records. To override view-level permissions, you must define permissions for its zones and resource records. For example, you can grant an admin group read-only permission to a view and read/write permission to all its zones. This allows the admins to display the view properties, but not edit them, and to create, edit and delete zones in the view.

You can grant read-only or read/write permission, or deny access to DNS views, as follows:

- All views—Global permission that applies to all DNS views in the database.
- A specific view—Applies to its properties and its zones, if you do not define zone-level permissions. This overrides the global view permissions.
- All zones in a view—If you do not define permissions for zones, they inherit the permissions of the view they are in.

For information on setting permissions for a view and its zones, see [Applying Permissions and Managing Overlaps](#) on page 166.

The following table lists the tasks admins can perform and the required permissions for DNS views.

Table 4.11 Permissions for DNS Views

Tasks	Grid Member(s)	All DNS Views	Specific DNS View	All DNS Zones
Create, modify, and delete DNS views		RW		
Create, modify, and delete DNS zones with assigned members	RW			RW
Create, modify, and delete DNS zones without assigned members				RW
Modify and delete a specific DNS view			RW	
Create, modify, and delete DNS zones, subzones, and resource records in a specific DNS view			RW	RW
Add Grid members to a Match Members list of a DNS view	RW		RW	
Delete a DNS view with Grid members in a Match Members list	RW		RW	
View DNS view properties, DNS zones, and resource records		RO		
View DNS zone properties, subzones, and resource records				RO
Restart services from the DNS tab	RO		RW	

Administrative Permissions for Zones

By default, zones inherit administrative permissions from the DNS view in which they reside. You can override view-level permissions by setting permissions for specific zones. Permissions set for a zone are inherited by its subzones and resource records. To override zone-level permissions, set permissions for specific subzones and resource records.

For example, you can grant an admin group the following permissions:

- Read-only to a zone and to all its A, AAAA, and PTR records (in reverse and forward-mapping zones)
- Read/Write permission to all MX and SRV records in the zone
- Deny to all the other resource records—CNAME, DNAME, TXT, host, and bulk host

You can grant read-only or read/write permission, or deny access to zones as follows:

- All zones—Global permission that applies to all zones in all views.
- All zones in a view—Permissions at this level override the global permissions.
- A specific zone—Applies to the zone properties and resource records, if you do not define permissions for its resource records. This overrides global and view-level permissions. If you delete a zone and reparent its subzone, the subzone inherits the permissions of the new parent zone.
- All Response Policy Zones—Global permission that applies to all the Response Policy Zones.
- All Response Policy Rules—Global permission that applies to all the local Response Policy Zone rules.

Note: Object permissions are not applicable to Response Policy Zone rules.

- Each resource record type in a zone—For example, you can define permissions for all A records and for all PTR records in a zone. If you do not define permissions for resource records, they inherit the permissions of the zone in which they reside.

For information on setting permissions for zones and resource records, see [Applying Permissions and Managing Overlaps](#) on page 166.

The following table lists the tasks admins can perform and the required permissions for zones.

Table 4.12 DNS Zone Permissions

Tasks	Grid Member(s)	Specific DNS View	All DNS Zones	Specific DNS Zone	Resource Records	Shared Record Group
Create, modify, and delete zones, subzones and resource records with assigned members	RW		RW			
Create, modify, and delete zones, subzones and resource records without assigned members			RW			
Lock and unlock a zone				RW		
Delete a zone with assigned Grid members	RW			RW		
Create, modify, and delete all zones, subzones, and resource records in a specific view		RW	RW			
Assign a name server group (member) to a zone	RW			RW		
Delete a zone with name server groups assigned	RW			RW		
Assign a shared record group to a zone				RW		RW
View zone properties, subzones, and resource records of a specific zone				RO		

Tasks	Grid Member(s)	Specific DNS View	All DNS Zones	Specific DNS Zone	Resource Records	Shared Record Group
Search for zones, subzones, and resource records in a specific DNS view		RO	RO			
Copy resource records from one zone to another: Source zone				RO	RO	
Copy resource records from one zone to another: Destination Zone				RW	RW	

Administrative Permissions for Resource Records

Resource records inherit the permissions of the zone to which they belong. You can override zone-level permissions by setting permissions for specific resource records.

You can grant read-only or read/write permission, or deny access to resource records as follows:

- Each resource record type in all zones and in all views—Global permission that applies to all resource records of the specified type; for example, all A records in the database.
- Each resource record type in a zone—Permissions at this level override global permissions.
- A specific resource record—Overrides zone-level permissions.

For information on setting permissions for resource records, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for resource records.

Table 4.13 DNS Resources

Tasks	Resource Record Type	Specific Resource Record
Create, modify, and delete resource records for a specified type, such as all A records or all PTR records	RW	
View resource records for a specified type only	RO	
Search for records of a specified type	RO	
View a specific resource record		RO
View, modify, and delete a specific resource record		RW

The following are additional guidelines:

- Only admins with read/write permission to bulk host records and read/write permission to reverse zones can create bulk host records and automatically add reverse-mapping zones.
- To create host records, admins must have read/write permission to the network and zone of the host.
- Admins must have read-only permission to the host records in a zone to view the Host Name Compliance Report. Admins must have read/write permission to the resource records in a zone to modify host names that do not comply with the host policy.

Administrative Permissions for Shared Record Groups

By default, only superusers can add, edit, and delete shared record groups. Limited-access admin groups can access shared record groups, only if their administrative permissions are defined.

You can set different permissions for a shared record group and for each type of shared resource record in the group. For example, you can grant a role or an admin group the following permissions:

- Read-only to a shared record group and to all its shared A and AAAA records
- Read/Write permission to all the shared MX and SRV records in the shared record group
- Deny to the TXT records

You can grant read-only or read/write permission, or deny access to shared record groups, as follows:

- All shared record groups—Global permission that applies to all shared record groups in the database.
- A specific shared record group—Overrides global permissions.
- Each shared record type in all shared record groups — The shared resource record types include shared A records, shared AAAA records, shared MX records, shared SRV records, and shared TXT resource records.
- Each shared record type in a shared record group— Permissions at this level override global permissions.
- A specific shared record—Overrides zone-level permissions.

Note the following guidelines:

- Shared record group permissions override zone permissions.
- Even if a zone is locked, superusers and limited-access users with read/write access can still edit or delete a shared record in the zone.

For information on setting permissions for shared record groups, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for shared record groups.

Table 4.14 Permissions for Shared Record Groups

Tasks	All Shared Record Groups	Specific Shared Record Group	Shared Record Type	Specific DNS Zone	Specific Shared Record
Create, modify, and delete shared record groups	RW				
Modify and delete a shared record group		RW			
View a shared record group		RO			
Create, modify, and delete shared records for a specific type			RW		
View or search for shared records of a specific type			RO		
Create, modify, and delete shared records for a specific type in a specified shared record group		RW	RW		
View shared records for a specific type in a specified shared record group only		RO	RO		
Create, modify, and delete a shared record					RW
View a specific shared record					RO
Assign a shared record group to DNS zones		RW		RW	
Change the DNS zones associated with a shared record		RW		RW	

Tasks	All Shared Record Groups	Specific Shared RRecord Group	Shared Record Type	Specific DNS Zone	Specific Shared Record
Delete zones with a shared record group assigned. Before you delete a shared record group, you must remove all zones associated with it.		RW		RW	

Administrative Permissions for DNS64 Synthesis Groups

By default, only superusers can add, edit, and delete DNS64 synthesis groups. Limited-access admin groups can access synthesis groups, only if their administrative permissions are defined.

You can grant read-only or read/write permission, or deny access to synthesis groups, as follows:

- All synthesis groups—Global permission that applies to all shared record groups in the database.
- A specific synthesis group—Overrides global permissions.

For information on setting permissions for synthesis groups, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for synthesis groups.

Table 4.15 Permissions for DNS64 Synthesis Groups

Tasks	All Synthesis Groups	Specific Synthesis Group	Grid	Specific Member	Specific DNS View
Create, modify, and delete synthesis groups	RW				
Modify and delete a specific synthesis group		RW			
View a synthesis group		RO			
Apply a synthesis group to the Grid		RO	RW		
Apply a synthesis group to a member		RO		RW	
Apply a synthesis group to a DNS view		RO			RW

ADMINISTRATIVE PERMISSIONS FOR DHCP RESOURCES

Limited-access admin groups can access certain DHCP resources only if their administrative permissions are defined. By default, the appliance denies access when a limited-access admin group does not have defined permissions. You can grant admin groups read-only or read/write permission, or deny access to the following DHCP resources:

- Network views
- IPv4 networks
- Hosts
- IPv4 DHCP ranges
- IPv4 DHCP fixed addresses
- IPv4 DHCP reservations
- MAC address filters
- IPv4 shared networks
- IPv4 network templates
- IPv4 DHCP range templates
- IPv4 fixed address templates
- IPv4 DHCP enabled host addresses
- IPv4 DHCP lease history
- Roaming hosts
- IPv6 networks
- IPv6 DHCP ranges
- IPv6 DHCP fixed addresses
- IPv6 DHCP enabled host addresses
- IPv6 shared networks
- IPv6 network templates
- IPv6 DHCP range templates
- IPv6 fixed address templates
- IPv6 DHCP lease history

You can grant an admin group broad permissions to DHCP resources, such as read/write permission to all IPv4 or IPv6 networks and shared networks in the database. In addition, you can grant permission to specific resources, such as a specific IPv4 or IPv6 network or DHCP range, or an individual address in an IPv4 or IPv6 network. Permissions at more specific levels override global permissions.

You can also define permissions for specific DHCP objects and Grid member to restrict admins to perform only the specified DHCP tasks on the specified member. For more information, see [Defining DNS and DHCP Permissions for Grid Members](#) on page 164.

The following sections describe the different types of permissions that you can set for DHCP resources:

- [Administrative Permissions for Network Views](#) on page 206
- [Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks](#) on page 207
- [Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations](#) on page 208
- [Administrative Permissions for IPv4 and IPv6 DHCP Ranges](#) on page 210
- [Administrative Permissions for IPv4 or IPv6 DHCP Templates](#) on page 211
- [Administrative Permissions for Roaming Hosts](#) on page 212
- [Administrative Permissions for MAC Address Filters](#) on page 212
- [Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories](#) on page 213

Administrative Permissions for Network Views

Limited-access admin groups can access network views, including the default network view, only if they have read-only or read/write permission to a specific network view or to all network views. Permissions granted to a network view apply to all its IPv4 and IPv6 networks, shared networks, DHCP ranges and fixed addresses.

You can grant admin groups read-only or read/write permission, or deny access to network views as follows:

- All network views—Global permission that applies to all network views in the database.
- A specific network view—Permission to a specific network view applies to the properties you set in the *Network View* editor, and to all the IPv4 and IPv6 networks and shared networks in the network view. This overrides the global permission to all network views. When you configure permissions for a network view, you can also set permissions for the following:
 - All IPv4 and IPv6 networks in the selected network view—If you do not define permissions for IPv4 or IPv6 networks, they inherit the permissions of their network view.
 - All IPv4 and IPv6 shared networks in a specific network view—If you do not define permissions for IPv4 or IPv6 shared networks, they inherit the permissions of their network view.

Note that you can grant an admin group read-only or read/write permission to specific IPv4 or IPv6 networks in a network view, without granting them permission to that network view. For information, see [Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks](#) on page 207.

For information on how to define permissions for network views, see [Applying Permissions and Managing Overlaps](#) on page 166.

The following table lists the tasks admins can perform and the required permissions for network views.

Table 4.16 Network View Permissions

Tasks	All DNS Views	Specific DNS View	All Network Views	Specific Network View	All IPv4 or IPv6 Networks	All IPv4 or IPv6 Shared Networks
Create and delete network views and their associated DNS views	RW		RW			
Create and delete a network view and its associated DNS views		RW		RW		
Create, modify, and delete IPv4 and IPv6 networks and shared networks in all network views			RW			
Create, modify, and delete IPv4 and IPv6 networks and shared networks in a network view				RW		
View the properties of all network views			RO			
View network statistics of all network views			RO			
View and search for all IPv4 and IPv6 networks and shared networks			RO			
View the properties of a network view				RO		
View and search for IPv4 and IPv6 networks and shared networks in a network view				RO		
Expand and join IPv4 and IPv6 networks			RW			
Expand and join IPv4 and IPv6 networks in a specific network view				RW		
Create, modify, and delete IPv4 and IPv6 networks, DHCP ranges and fixed addresses in a specific network view				RW		

Tasks	All DNS Views	Specific DNS View	All Network Views	Specific Network View	All IPv4 or IPv6 Networks	All IPv4 or IPv6 Shared Networks
View network statistics and properties of all networks in a network view				RO		
Search for IPv4 and IPv6 networks in a network view				RO		
Create, modify, and delete all IPv4 or IPv6 shared networks						RW
View the properties of all IPv4 or IPv6 shared networks						RO
View and search for IPv4 and IPv6 shared networks in a network view				RO		
Restart services from the DHCP tab	RO			RW		

Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks

Limited-access admin groups can access IPv4 and IPv6 networks, including shared networks, only if their administrative permissions are defined. Permissions for a network apply to all its DHCP ranges and fixed addresses. To override network-level permissions, you must define permissions for specific DHCP ranges and fixed addresses. For example, you can grant an admin group read-only permission to a network, read/write permission to its DHCP ranges, and read-only permission to its fixed addresses.

You can grant read-only or read/write permission, or deny access to networks, as follows:

- All IPv4 or IPv6 networks—Global permission that applies to all IPv4 or all IPv6 networks in the database.
- All IPv4 or IPv6 shared networks—Global permission that applies to all IPv4 or all IPv6 shared networks in the database.
- A specific IPv4 or IPv6 network—Network permissions apply to its properties and to all DHCP ranges, fixed addresses and hosts in the network, if they do not have permissions defined. This overrides global permissions.
- All IPv4 or IPv6 DHCP ranges in a network—If you do not define permissions for DHCP ranges, they inherit the permissions of the network in which they reside.
- All IPv4 or IPv6 fixed addresses in a network—If you do not define permissions for fixed addresses, they inherit the permissions of the network in which they reside.

To define permissions for a specific IPv4 or IPv6 network and its DHCP ranges and fixed addresses, see [Applying Permissions and Managing Overlaps](#) on page 166.

The following table lists the tasks admins can perform and the required permissions for IPv4 and IPv6 networks.

Table 4.17 Network Permissions

Tasks	Grid Member(s)	All IPv4 or IPv6 Networks	Specific IPv4 or IPv6 Network	All IPv4 or IPv6 Shared Networks	Specific DNS Zone	All IPv4 or IPv6 DHCP Ranges	All IPv4 or IPv6 Fixed Addresses	IPv4 or IPv6 Network Template
Create, modify, and delete IPv4 or IPv6 networks, DHCP ranges, and fixed addresses without assigned Grid members		RW						
Create, modify, and delete IPv4 or IPv6 networks, DHCP ranges, and fixed addresses with assigned Grid members	RW	RW						
Assign a Grid member to a specific IPv4 or IPv6 network and its DHCP ranges	RW		RW					
Expand and join IPv4 or IPv6 networks		RW						
Create IPv4 or IPv6 networks from templates		RW						RO
Create, modify, and delete an IPv4 or IPv6 network		RW						
View IPv4 or IPv6 network properties and statistics, and search for DHCP ranges and fixed addresses in a specific network			RO					
Create, modify, and delete IPv4 or IPv6 DHCP ranges and fixed addresses in a specific network			RW					
Create and split an IPv4 or IPv6 network and automatically create a reverse DNS zone			RW		RW			
Create, modify, and delete IPv4 or IPv6 shared networks				RW				
View IPv4 or IPv6 shared networks				RO				
Create, modify, and delete IPv4 or IPv6 DHCP ranges with an assigned member in a specific network	RW		RW					
Create, modify, and delete IPv4 or IPv6 DHCP ranges						RW		
View and search for IPv4 or IPv6 DHCP ranges in a specific network			RO					
Create, modify, and delete IPv4 or IPv6 fixed addresses							RW	
View and search for IPv4 or IPv6 fixed addresses in a specific network			RO					

Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations

IPv4 and IPv6 fixed addresses and IPv4 reservations inherit the permissions of the networks in which they reside. You can override network-level permissions by defining permissions for fixed addresses.

You can grant read-only or read-write permission, or deny access to fixed addresses, as follows:

- All IPv4 fixed addresses/reservations—Global permission that applies to all IPv4 fixed addresses and reservations in the database.
- All IPv6 fixed addresses—Global permission that applies to all IPv6 fixed addresses in the database.
- All IPv4 fixed addresses/reservations in a network—Permissions at this level override global permissions. If you do not define permissions for fixed addresses and reservations, they inherit the permissions of the network in which they reside.

- All IPv6 fixed addresses in a network—Permissions at this level override global permissions. If you do not define permissions for IPv6 fixed addresses, they inherit the permissions of the network in which they reside.
- A single IPv4 fixed address/reservation—Overrides global and network-level permissions.
- A single IPv6 fixed address—Overrides global and network-level permissions.

For information on setting permissions for fixed addresses, see [Applying Permissions and Managing Overlaps](#) on page 166.

The following table lists the tasks admins can perform and the required permissions for IPv4 and IPv6 fixed addresses.

Table 4.18 Permissions for Fixed Addresses/Reservations

Tasks	Specific IPv4 or IPv6 Network	All IPv4 or IPv6 fixed Addresses/ IPv4 Reservations	Specific IPv4 or IPv6 Fixed Address/ IPv4 Reservation
Create, modify, and delete IPv4 fixed addresses/reservations or IPv6 fixed addresses		RW	
Create, modify, and delete IPv4 fixed addresses/reservations or IPv6 fixed addresses in a specific network	RW		
Modify and delete an IPv4 fixed address/reservation or IPv6 fixed address			RW
View and search for all IPv4 fixed addresses/reservations or IPv6 fixed addresses		RO	
View and search for IPv4 fixed addresses/reservations or IPv6 fixed addresses in a network	RO	RO	
View and search for an IPv4 fixed address/reservation or IPv6 fixed address			RO

Administrative Permissions for IPv4 or IPv6 DHCP Enabled Host Addresses

A read-write permission to IPv4 or IPv6 Host Address gives limited-access users the ability to create, modify, and delete IPv4 and IPv6 DHCP enabled host addresses in a specified network. Admin users with a read-write permission can create, modify, and delete IPv4 or IPv6 DHCP enabled host addresses only in the specified network. They do not have the ability to create, modify or delete any networks or objects, such as fixed addresses, in those networks.

You can also grant admin users read-only permission or deny access to the following:

- IPv4 Host Address—Object permission that applies to all IPv4 DHCP enabled host addresses in a specified network.
- IPv6 Host Address—Object permission that applies to all IPv6 DHCP enabled host addresses in a specified network.

For information about setting permissions for DHCP enabled host addresses, see [Applying Permissions and Managing Overlaps](#) on page 166.

The following table lists tasks that admins can perform and the required permissions for IPv4 and IPv6 DHCP enabled host addresses.

Table 4.19 Permissions for DHCP Enabled Host Addresses

Tasks	Specific IPv4 or IPv6 Network	All IPv4 or IPv6 DHCP enabled host Addresses
	Create, modify, and delete IPv4 or IPv6 DHCP enabled host addresses in a specified network	RW
	Modify and delete a specific IPv4 or IPv6 DHCP enabled host address	RW
	View and search for all IPv4 or IPv6 DHCP enabled host addresses	RO
	View and search for IPv4 or IPv6 DHCP enabled host addresses in a specified network	RO

Administrative Permissions for IPv4 and IPv6 DHCP Ranges

DHCP ranges inherit the permissions of the networks in which they reside. You can override network-level permissions by defining permissions for DHCP ranges. You can read-only or read/write permission, or deny access to DHCP address ranges, as follows:

- All IPv4 or IPv6 DHCP ranges—Global permission that applies to all IPv4 or IPv6 DHCP ranges in the database.
- All IPv4 or IPv6 DHCP ranges in a network—Permissions at this level override global permissions. If you do not define permissions for DHCP ranges, they inherit the permissions of the network in which they reside.
- A single IPv4 or IPv6 DHCP range—Overrides global and network-level permissions.

For information on setting permissions for DHCP ranges, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admin can perform and the required permissions for DHCP ranges.

Table 4.20 DHCP Ranges

Tasks	Grid Member(s)	Specific IPv4 or IPv6 Network	All DHCP IPv4 or IPv6 Ranges	Specific IPv4 or IPv6 DHCP Range	MAC Address Filter
	Create, modify, and delete IPv4 or IPv6 DHCP ranges with an assigned member or a failover association	RW		RW	
	Create, modify, and delete IPv4 or IPv6 DHCP ranges in a network with assigned members	RW	RW		
	Modify and delete an IPv4 or IPv6 DHCP range with an assigned member	RW		RW	
	View and search for all IPv4 or IPv6 DHCP ranges with an assigned member	RO		RO	
	View and search for IPv4 or IPv6 DHCP ranges in a network with assigned members	RO	RO		

Tasks	Grid Member(s)	Specific IPv4 or IPv6 Network	All DHCP IPv4 or IPv6 Ranges	Specific IPv4 or IPv6 DHCP Range	MAC Address Filter
View and search for an IPv4 or IPv6 DHCP range with an assigned member	RO			RO	
View and search for an IPv4 or IPv6 DHCP range without an assigned member				RO	
Apply relay agent and option filters to an IPv4 DHCP range				RW	
Apply a MAC address filter to an IPv4 DHCP range				RW	RO

Administrative Permissions for IPv4 or IPv6 DHCP Templates

There are three types of DHCP templates for IPv4 and IPv6 objects—network, DHCP range, and fixed address/reservation templates. To access any of these templates, a limited-access admin group must have read-only permission to the template. Limited-access admin groups cannot have read/write permission to the templates. Only superusers can create, modify and delete network, DHCP range, and fixed address templates. An admin group with read-only permission to the DHCP templates can view them and use them to create networks, DHCP ranges and fixed addresses, as long as they have read/write permissions to those DHCP resources as well.

You can set global read-only permission that applies to all DHCP templates, and you can set permissions to specific templates as well.

For information on setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for DHCP templates.

Table 4.21 Permissions for DHCP Templates

Tasks	IPv4 or IPv6 DHCP Templates	All IPv4 or IPv6 Networks	All IPv4 or IPv6 DHCP Ranges	All IPv4 or IPv6 Fixed Addresses/IPv4 Reservations
Create IPv4 or IPv6 networks from templates	RO	RW		
Create IPv4 or IPv6 DHCP ranges from templates	RO		RW	
Create IPv4 fixed addresses/reservations or IPv6 fixed addresses from templates	RO			RW
View templates	RO			

Note the following additional guidelines:

- DHCP range templates and fixed address templates do not inherit their permissions from network templates. You must set permissions for each type of template.

- An admin group can create a network using a network template that includes a DHCP range template and a fixed address template, even if it has no permission to access the DHCP range and fixed address templates.

Administrative Permissions for Roaming Hosts

Limited-access admin groups can access roaming hosts only if their administrative permissions are defined. The appliance denies access to roaming hosts for which an admin group does not have defined permissions.

You can grant read-only or read/write permission, or deny access to roaming hosts as follows:

- All roaming hosts in the database—Global permission that applies to all the roaming hosts in the database.
- A specific roaming host—Permissions that applies to specific roaming host.

For information on setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for roaming host.

Table 4.22 Permissions for Roaming Hosts

Tasks	Grid DHCP Properties	Specific IPv4 or IPv6 Roaming Host	All Roaming Host
Enable roaming hosts	RW		
View roaming host	RO	RO	RO
Create, modify, and delete roaming hosts	RO		RW
Modify and delete roaming host	RO	RW	

Administrative Permissions for MAC Address Filters

Limited-access admin groups can access MAC address filters only if their administrative permissions are defined. The appliance denies access to MAC address filters for which an admin group does not have defined permissions.

You can grant read-only or read/write permission, or deny access to MAC address filters as follows:

- All MAC address filters in the database
- A specific MAC address filter

For information on setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for MAC address filters.

Table 4.23 Permissions for MAC Filters

Tasks	All MAC Address Filters	Specific MAC Address Filter	Specific IPv4 DHCP Ranges
Create, modify, and delete MAC address filters	RW		
Create, modify, and delete MAC address entries for a MAC address filter		RW	
Modify and delete a MAC address filter		RW	
Apply a MAC address filter to an IPv4 DHCP range		RO	RW
Delete a MAC address filter from an IPv4 DHCP range		RO	RW
View MAC address filters and their MAC address entries	RO		
View a MAC address filter and its MAC address entries		RO	

Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories

A limited-access admin group can view and export the IPv4 and IPv6 DHCP lease histories if it has read-only permission to the IPv4 and IPv6 DHCP lease history. Permissions to the IPv4 and IPv6 DHCP lease histories are different from the network permissions. Therefore, an admin group can access the IPv4 and IPv6 DHCP lease histories, regardless of its network permissions. Note that only superusers can import a DHCP lease history file.

To define permissions for the IPv4 and IPv6 DHCP lease histories:

- For an admin group: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin_group* in the Groups table, and then click the Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.
or
For an admin role: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin_role* in the Roles table, and then click Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.
- Complete the following in the *Manage Global Permissions* dialog box:
 - Permission Type:** Select **DHCP Permissions** from the drop-down list.
 - In the table, select **Read/Write**, **Read-only**, or **Deny** for **All IPv4 DHCP Lease History** and **All IPv6 DHCP Lease History**.
- Save the configuration and click **Restart** if it appears at the top of the screen.

ADMINISTRATIVE PERMISSIONS FOR FILE DISTRIBUTION SERVICES

You can restrict access to the TFTP, HTTP and FTP services provided by the appliance. By default, the appliance denies access to the TFTP, HTTP and FTP services, unless an admin group has their administrative permissions defined.

You can grant read-only or read/write permission, or deny access to the following resources:

- Grid File Distribution Properties—Applies to the Grid and its members, directories, and files. You can set this from the Administrators perspective only.
- Member File Distribution Properties—Applies to the Grid member properties only.
- A specific directory—Applies to the directory and its files.

For information on setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for file distribution services.

Table 4.24 Permissions for File Distribution Services

Tasks	Grid File Distribution Properties	Member Distribution Properties	Specific Directory
Create and remove directories and files	RW		
Modify the Grid and member file distribution properties	RW		
View the Grid and member file distribution properties, directories, and files	RO		
Modify the member file distribution properties		RW	
View the member file distribution properties		RO	
Add and delete a directory, subdirectories, and files in the directory			RW
View a directory and its subdirectories and files			RO

ADMINISTRATIVE PERMISSIONS FOR DASHBOARD TASKS

Limited-access admin groups can configure IPAM tasks on the Tasks Dashboard only if their administrative permissions are defined. The appliance denies access to IPAM tasks for which an admin group does not have defined permissions.

You can grant read-only or read/write permission, or deny access to IPAM tasks as follows:

- All IPAM tasks on the Tasks Dashboard
- A specific IPAM task

When you deny access to an IPAM task for an admin group, users cannot configure the task on their dashboards. Users must have at least read-only permission to a specific task to see it in the task pack. To perform a specific task, users must also have read/write permission to the objects associated with the task. For information about specific permissions for IPAM, DNS, and DHCP objects, see [Administrative Permissions for IPAM Resources](#) on page 198, [Administrative Permissions for DNS Resources](#) on page 199, and [Administrative Permissions for DHCP Resources](#) on page 205.

For information about setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for configuring IPAM tasks on the Tasks Dashboard.

Table 4.25 Permissions for IPAM Tasks

Tasks	All Dashboard Tasks	Add Networks	Add Hosts	Add Fixed Addresses	Add CNAME Record	Add TXT Record	Add MX Record
Configure all tasks in the IPAM task pack	RO RW						
Configure the Add Networks task		RO RW					
Configure the Add Hosts task			RO RW				
Configure the Add Fixed Addresses task				RO RW			
Configure the Add CNAME Record task					RO RW		
Configure the Add TXT Record task						RO RW	
Configure the Add MX Record task							RO RW

ADMINISTRATIVE PERMISSIONS FOR OCSP SERVER GROUPS AND CA CERTIFICATES

Limited-access admins can configure OCSP server groups and CA certificates only if their administrative roles and permissions are defined. If you want to allow admins to configure two-factor authentication, you can assign the PKI Admin role to limited-access admins or grant them read/write permissions to the following:

- All OCSP Groups
- All CA Certificates

For information about setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the admin tasks and required permissions for configuring an OCSP server group and managing CA certificates.

Table 4.26 Administration Permissions

Tasks	Grid Member(s)	All OSCP Server Groups	All CA Certificates			
Create, modify, and delete the OCSP server group		RW				
Create, modify, and delete CA certificates	RW		RW			

ADMINISTRATIVE PERMISSIONS FOR LOAD BALANCERS

Limited-access admins can view and manage GLBs (Global Load Balancers), load balancer synchronization groups, and their associated objects if their administrative roles and permissions are defined. If you want to allow admins to manage GLB objects, assign the Load Balancer Admin role to limited-access admins and grant them the following permissions:

- Read/write, read-only, or deny permission to NIOS managed GLB groups and independent load balancers
- Read/write, read-only, or deny permission to NIOS managed GLB objects
- Read-only or deny permission to GLB objects, such as DNS profiles and iRules, that are synchronized from the GLB but cannot be managed through NIOS

For information about setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the admin tasks and required permissions for configuring GLBs, load balancer synchronization groups, and their associated objects.

Table 4.27 Administration Permissions for Load Balancers and Load Balancer Synchronization Groups

Tasks	Grid Member(s)	All Load Balancer Objects	All Load Balancers	All Load Balancer Groups
View load balancer objects	RO	RO		
Add and modify synchronized load balancer objects	RW	RW		
Add, modify, and delete synchronized load balancer objects	RW	RW	RW	RW
Modify and delete synchronized load balancer groups	RW		RW	RW

ADMINISTRATIVE PERMISSIONS FOR NAMED ACLs

Only superusers and limited-access users with Read/Write permission to All Named ACLs and Read/Write permission to the corresponding objects and operations can manage named ACLs and their ACEs.

For information about access control and named ACLs, see [Configuring Access Control](#) on page 306. The following table lists the operations and required permissions for managing named ACLs.

Table 4.28 Administration Permissions

Tasks	Grid Member(s)	All Named ACLs	DNS Views	Related DNS objects	File Distribution
Create, modify, and delete named ACLs for all supported operations	RW	RW	RW	RW	RW
View named ACLs and ACEs	RW	RO	RO	RO	RO

ADMINISTRATIVE PERMISSIONS FOR DNS THREAT PROTECTION

You can grant read-only or read/write permission, or deny access to the following resources:

- Grid Security Properties—Applies to the Grid and its members.
- Member Security Properties—Applies to the Grid members only.

For information about setting permissions, see [Applying Permissions and Managing Overlaps](#) on page 166. The following table lists the tasks admins can perform and the required permissions for the threat protection service.

Table 4.29 Permissions for Threat Protection Service

Tasks	Grid Security Properties	Member Security Properties
View Grid security properties	RO	
Update Grid Security properties	RW	
View member security properties for specific Grid members	RO	RO
Update member security properties for specific Grid members	RW	RW
Start and stop threat protection service for a Grid member	RW	RW
Publish rules for a Grid member	RW	RW
View rule categories and rules for the Grid	RO	
Enable and disable rules for the Grid	RW	

Tasks	Grid Security Properties	Member Security Properties
Update rule versions for any rules on the Grid	RW	
Revert to a previous rule version for any rules on the Grid	RW	
Modify configuration parameters, such as action and severity, for rules on the Grid	RW	
Create custom rules from rule templates for the Grid	RW	
Delete custom rules for the Grid	RW	
View rule categories and rules on a Grid member	RO	RO
Enable and disable rules on a Grid member	RW	RW
Update rule versions for any rules on a Grid member	RW	RW
Revert to a previous rule version for any rules on a Grid member	RW	RW
Modify configuration parameters, such as action and severity, for rules on a Grid member	RW	RW
View threat protection related event statistics on a Grid member	RO	RO
Upgrade rulesets for a Grid	RW	



Chapter 5 Deploying a Grid

To deploy a Grid, it is important to understand what a Grid is, how to create a Grid Master and add members, and how to manage the Grid. This chapter explains these tasks in the following sections:

- [Introduction to Grids](#) on page 223
 - [Grid Communications](#) on page 225
 - [NAT Groups](#) on page 226
 - [Automatic Software Version Coordination](#) on page 229
 - [Grid Bandwidth Considerations](#) on page 231
- [About HA Pairs](#) on page 233
 - [Planning for an HA Pair](#) on page 233
 - [About HA Failover](#) on page 234
 - [VRRP Advertisements](#) on page 235
- [Creating a Grid Master](#) on page 236t
 - [Port Numbers for Grid Communication](#) on page 238
 - [Grid Setup Wizard](#) on page 238
 - [Creating an HA Grid Master](#) on page 238
 - [Creating a Single Grid Master](#) on page 242
- [Adding Grid Members](#) on page 245
 - [Adding a Single Member](#) on page 245
 - [Adding an HA Member](#) on page 247
 - [Joining Appliances to the Grid](#) on page 249
- [Auto-Provisioning NIOS Appliances](#) on page 250
 - [Joining Auto-Provisioned Appliances to the Grid](#) on page 250
- [Pre-Provisioning NIOS Appliances](#) on page 252
 - [Guidelines for Pre-provisioning Offline Grid Members](#) on page 252
 - [Configuring Pre-Provisioned Members](#) on page 253
 - [About Provisional Licenses](#) on page 253
 - [Joining Pre-Provisioned Members to the Grid](#) on page 254
 - [Configuration Example: Configuring a Grid](#) on page 255
- [Configuration Example: Configuring a Grid](#) on page 255
 - [Cable All Appliances to the Network and Turn On Power](#) on page 257
 - [Create the Grid Master](#) on page 257
 - [Define Members on the Grid Master](#) on page 259

- [*Join Appliances to the Grid*](#) on page 260
- [*Import DHCP Data*](#) on page 262
- [*Import DNS Data*](#) on page 263
- [*Using the Wizard*](#) on page 264
- [*After Using the Wizard*](#) on page 266
- [*Managing a Grid*](#) on page 267
 - [*Changing Grid Properties*](#) on page 267
 - [*Configuring Security Level Banner*](#) on page 268
 - [*Configuring Informational Level Banner*](#) on page 269
 - [*Configuring Recursive Deletions of Networks and Zones*](#) on page 269
 - [*Setting the MTU for VPN Tunnels*](#) on page 270
 - [*Removing a Grid Member*](#) on page 270
 - [*Promoting a Master Candidate*](#) on page 270
- [*About the Master Grid*](#) on page 271

Introduction to Grids

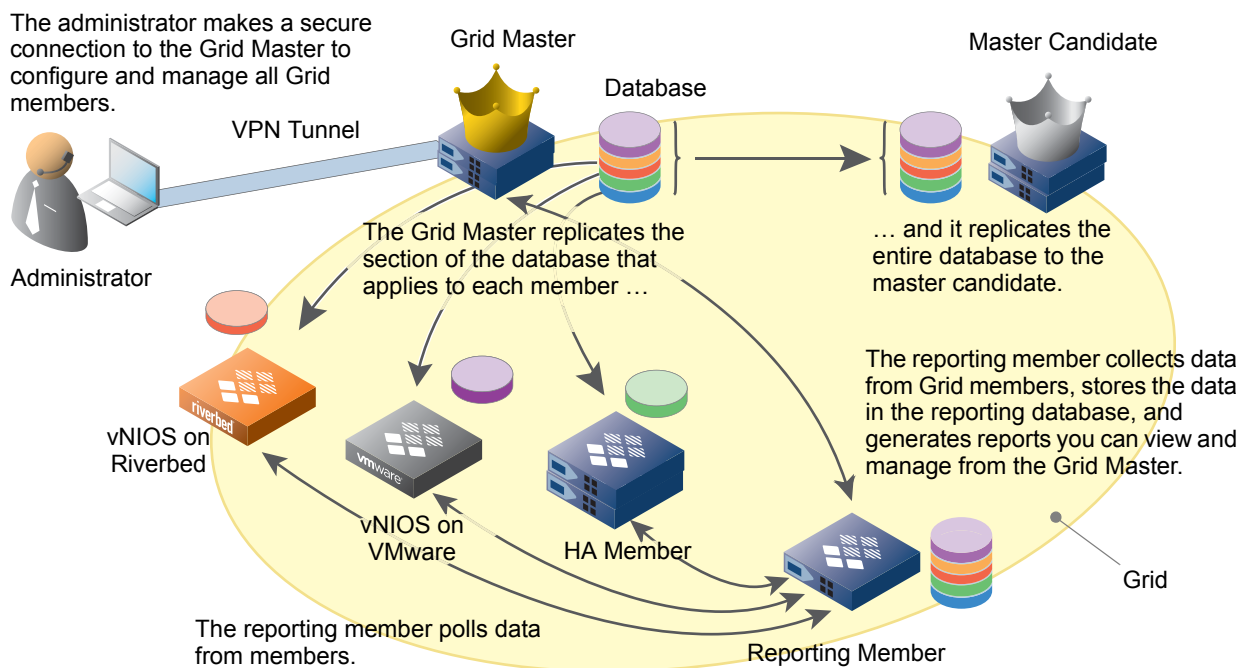
Note: Infoblox appliances support IPv4 and IPv6 networking configurations in most deployments cited in this chapter. You can set the LAN1 port to an IPv6 address and use that address to access Grid Manager. All HA operations can be applied across IPv6. Topics in this and following chapters generally use IPv4 examples. Also note that the LAN2 and MGMT ports also support IPv6. DNS and DHCP services are fully supported in IPv6 for the LAN2 and MGMT ports. Examples throughout this chapter use IPv4 addressing. Though interfaces on NIOS appliances support IPv4 and IPv6 transports, intra-Grid communications use IPv4.

A Grid is a group of two or more NIOS appliances that share sections of a common, distributed, built-in database and which you configure and monitor through a single, secure point of access: the Grid Master. A Grid can include Infoblox appliances and vNIOS appliances. A vNIOS appliance is a non-Infoblox hardware platform running the vNIOS software package. (Supported platforms include Riverbed Steelhead appliances running the Riverbed Services Platforms and VMware ESX and ESXi server platforms.) You can configure Infoblox appliances and vNIOS appliances for VMware as Grid Masters, Grid Master candidates, and Grid members. You can configure the other vNIOS appliances as Grid members only. For information, see [Supported vNIOS Appliance Configurations](#) on page 1305.

You can also add any of the supported Trinzic Reporting platforms as a logging and reporting device in your Grid. Infoblox provides a few Infoblox platforms that you can use as the logging and reporting device. For information about the supported appliances, see [Supported Platforms for Reporting](#) on page 1119. The reporting appliance collects data from members in the Grid and stores the data in the database. It then uses the data to generate predefined and user-defined reports that you can access through Grid Manager. These reports provide useful information about the IPAM, DNS, DHCP, and system activities and usage in your Grid. For more information about reporting, see [Infoblox Reporting Solution](#) on page 1113.

[Figure 5.1](#) shows the basic concept of a Grid, database distribution (or “replication”), and reporting.

Figure 5.1 Grid and Partitioned Database Replication

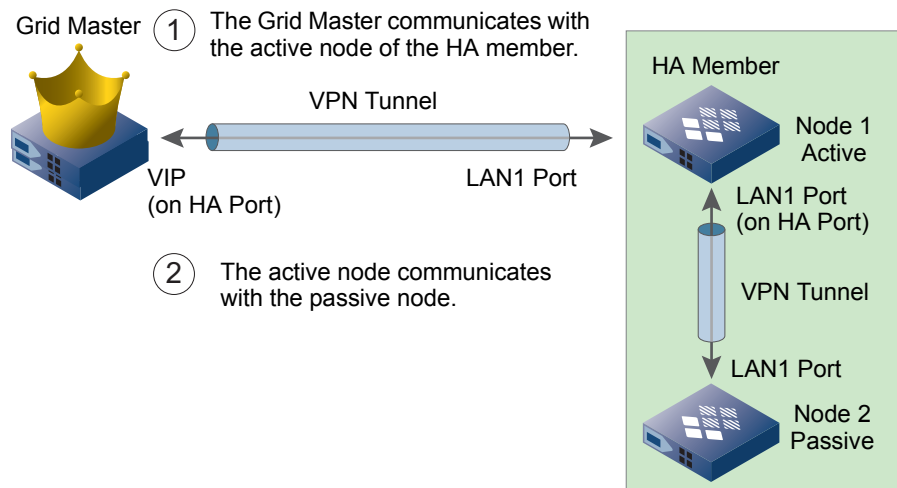


Note: In addition to the VPN tunnel securing administrative traffic to the Grid Master, all Grid communications between the Grid Master and Grid members pass through encrypted VPN tunnels (not shown).

The Grid Master can be either an HA master or a single master; that is, an HA (high availability) pair or a single appliance. Similarly, a Grid member can be either a single member or an HA member. You can add single appliances and HA pairs to a Grid, forming single members and HA members respectively. A single Grid member can be either an Infoblox appliance or a vNIOS appliance. An HA Grid member can be a pair of Infoblox appliances or vNIOS appliances. For information, see [Supported vNIOS Appliance Configurations](#) on page 1305.

The Grid Master communicates with every Grid member in a hub-and-spoke configuration. For an HA member, the Grid Master communicates with the active node, which in turn communicates with the passive node, as shown in [Figure 5.2](#).

Figure 5.2 Grid Communications to an HA Member



When adding vNIOS appliances to a Grid, you centralize the management of core network services of the virtual appliances through the Grid Master. vNIOS appliances support most of the features of the Infoblox NIOS software, with some limitations as described in [Appendix E, "vNIOS Appliance Limitations"](#), on page 1305.

For additional information specific to each platform, refer to the *Quick Start Guide for Installing vNIOS Software on Riverbed Services Platforms* and the *Quick Start Guide for Installing vNIOS Software on VMware Platforms*.

By default, Grid communications use the UDP transport with a source and destination port of 1194. This port number is configurable. For a port change to take effect, one of the following must occur: the HA master fails over, the single master reboots, or the Grid restarts services.

After adding an appliance or HA pair to a Grid, you no longer access the Infoblox GUI on that appliance. Instead, you access the GUI running on the Grid Master. Although you can create multiple administrator accounts to manage different services on various Grid members, all administrative access is through the Grid Master. So even if someone has administrative privileges to a single Grid member, that administrator must access the GUI running on the Grid Master to manage that member.

You can access the Infoblox GUI through an HTTPS connection to one of the following IP addresses and ports on the Grid Master:

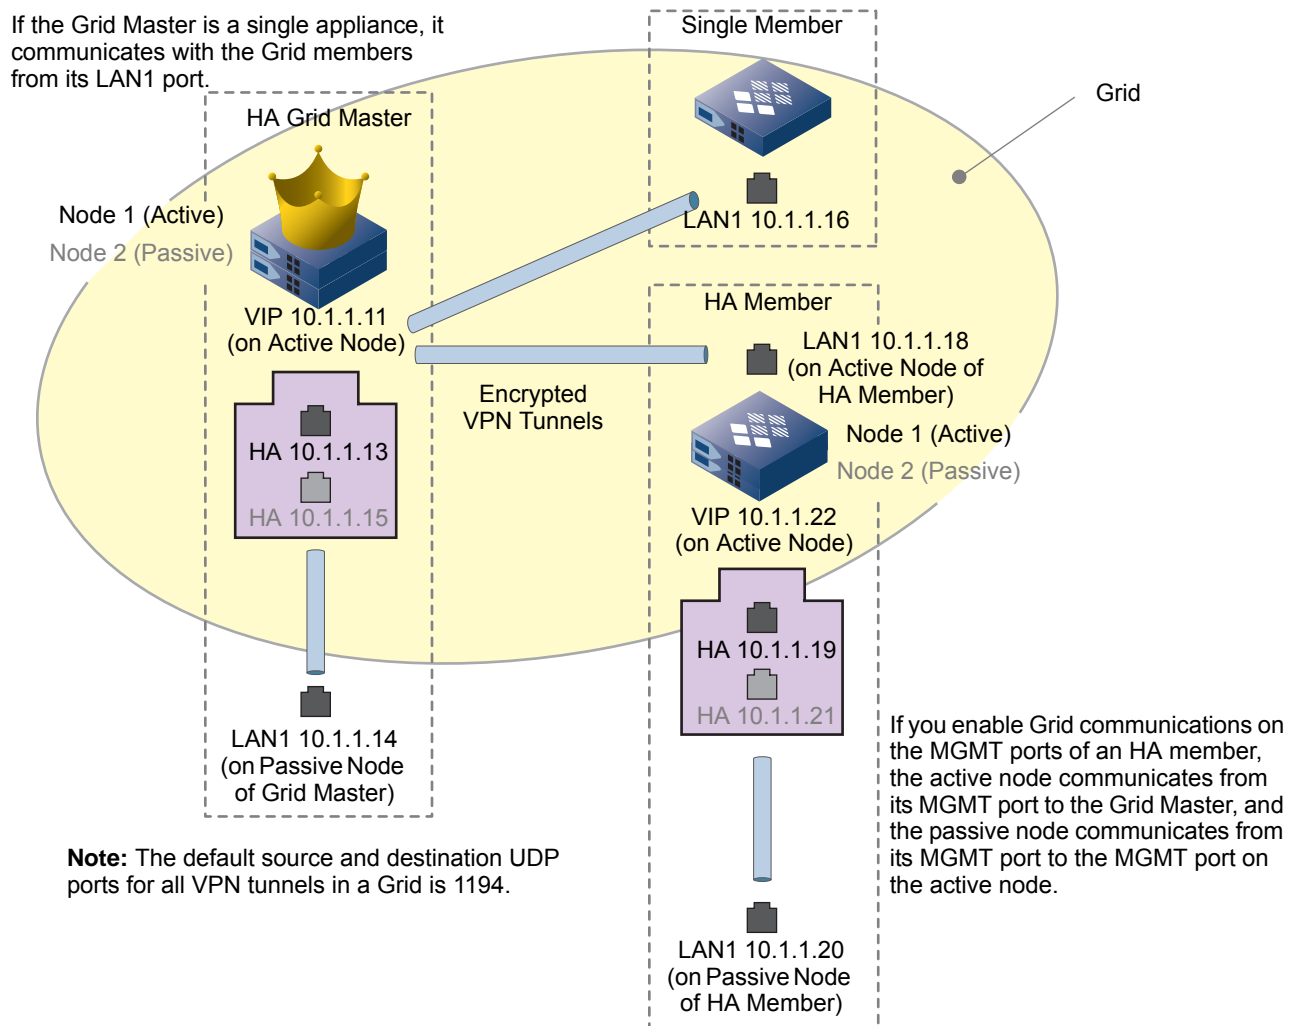
- The VIP address, which links to the HA port on the active node of an HA Grid Master
- The IP address of the LAN1 port on a single Grid Master
- The IP address of the MGMT port (if enabled) of the active node of an HA or single Grid Master. See [Using the MGMT Port](#) on page 359.

Grid Communications

The Grid Master synchronizes data among all Grid members through encrypted VPN tunnels. The default source and destination UDP port number for VPN tunnels is 1194. You can continue using the default port number or change it. For example, if you have multiple Grids, you might want each Grid to use a different port so that you can set different firewall rules for each. Whatever port number you choose to use for the VPN tunnels in a Grid, all the tunnels in that Grid use that single port number.

Before an appliance or HA pair forms a tunnel with the master, they first authenticate each other using the Challenge-Response Authentication Mechanism (CRAM). The source and destination port number for this traffic is 2114. During the CRAM handshake, the master tells the appliance or HA pair what port number to use when building the subsequent VPN tunnel.

Figure 5.3 VPN Tunnels within a Grid



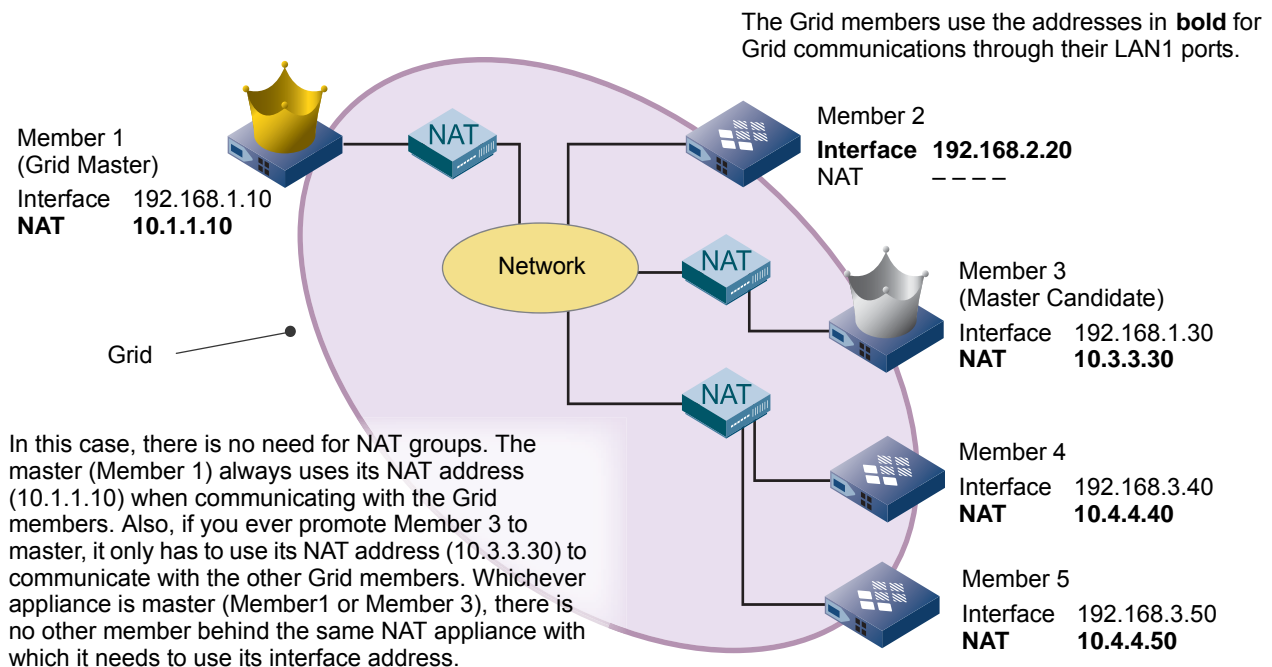
Another type of traffic, which flows outside the tunnels, is the VRRP (Virtual Router Redundancy Protocol) advertisements that pass between the active and passive nodes in an HA pair. The VRRP advertisements act like heartbeats that convey the status of each node in an HA pair. If the active node fails, the passive node becomes active. The VIP (virtual IP) address for that pair then shifts from the previously active node to the currently active node.

NAT Groups

Note: Infoblox NAT and NAT groups do not support NAT IPv6 operation.

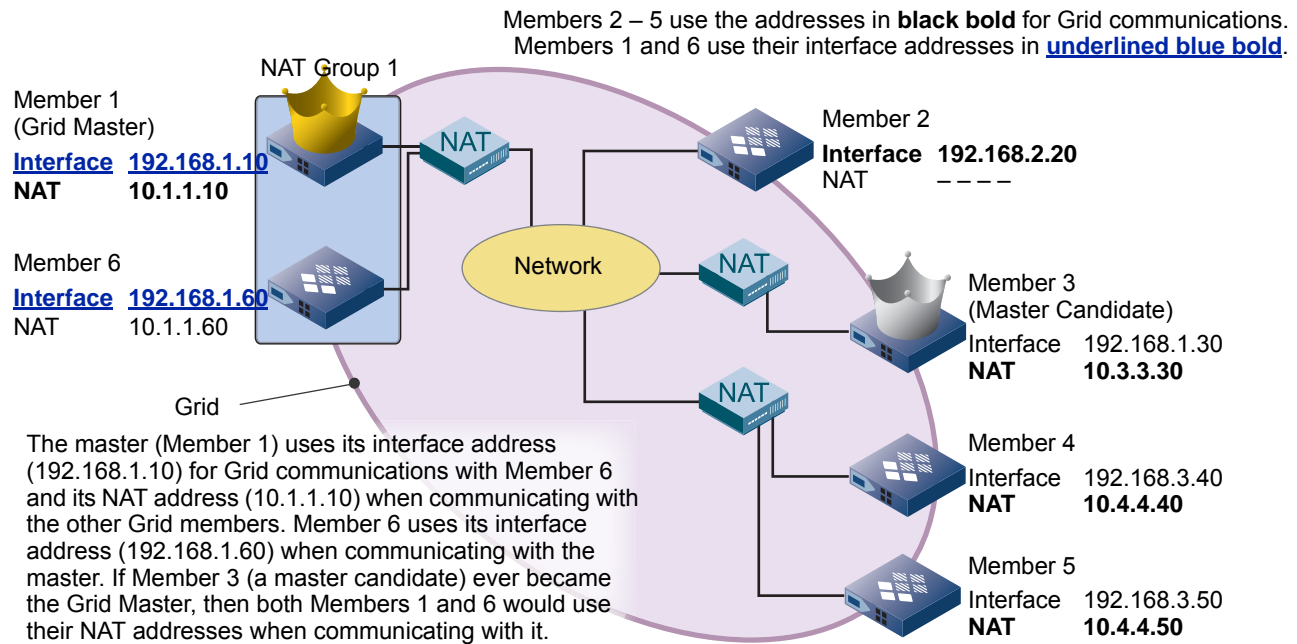
NAT groups are necessary if the Grid Master is behind a NAT appliance and there are members on both sides of that NAT appliance. Any members on the same side as the master go into the same NAT group as the master and use their interface addresses for Grid communications with each other. Grid members on the other side of that NAT appliance do not go in the same NAT group as the master and use the master's NAT address for Grid communications. These other members outside the NAT appliance can—but do not always need to be—in a different NAT group. To see when NAT groups become necessary for Grid communications, compare [Figure 5.4](#) below with [Figure 5.5](#) and [Figure 5.6](#) on page 228.

Figure 5.4 NAT without NAT Groups



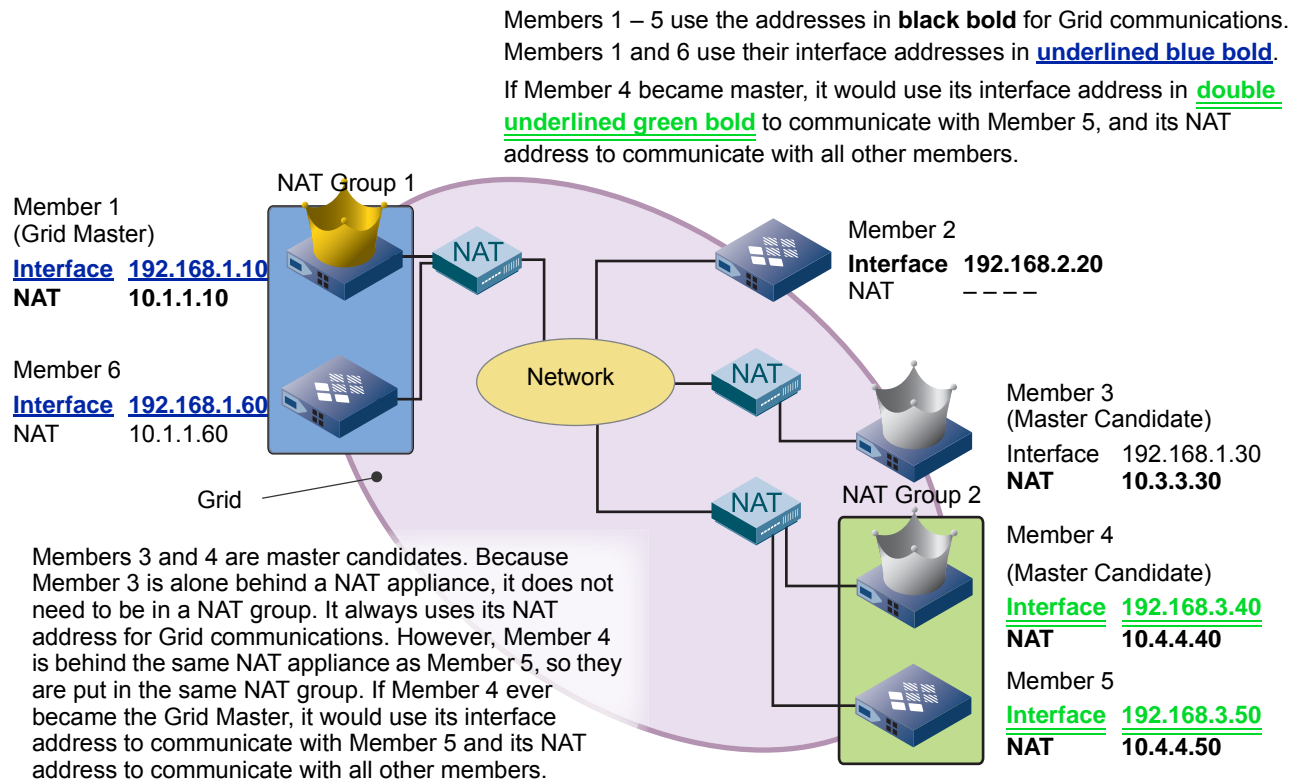
Note: A single or HA member using its MGMT port for Grid communications cannot be separated from the Grid Master behind a NAT appliance. For more information, see [Using the MGMT Port](#) on page 359.

Figure 5.5 Grid Master in NAT Group



The same use of NAT groups that applies to a Grid Master also applies to master candidates. If there are no other members behind the same NAT appliance as a master candidate, then the master candidate does not need to be in a NAT group. It always uses its NAT address for Grid communications. If another member is behind the same NAT appliance as the master candidate, then both the candidate and that member need to be in the same NAT group so that—if the candidate becomes master—they can use their interface addresses to communicate with each other (see [Figure 5.6](#)).

Figure 5.6 Grid Master and Master Candidate in NAT Groups



Although some members might not need to be in a NAT group, it is good practice to put all members in NAT groups in anticipation of adding or rearranging Grid members within the network. For example, in [Figure 5.4](#) – [Figure 5.6](#), Member 4 did not need to be in a NAT group until it became configured as a master candidate in [Figure 5.6](#). At that point, because Member 5 is also behind the same NAT appliance, it became necessary to create NAT Group 2 and add Members 4 and 5 to it. Similarly, if you add another member behind the NAT appliance in front of Member 3, then you must create a new NAT group and add Member 3 and the new member to it. Always using NAT groups can simplify such changes to the Grid and ensure that NAT appliances never interrupt Grid communications.

To create a NAT group:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **NAT Groups** tab.
4. Click the Add icon, and enter a name in the **Name** field and optionally, a comment in the Comment field.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

To add members to the NAT group:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Select a Grid member and click the Edit icon.
3. In the *Grid Member Properties* editor, select the **Network** -> **Advanced** tab and complete the following
 - **Enable NAT compatibility:** Select this check box.

- **NAT Group:** From the drop-down list, select the NAT group you previously created.
- **NAT Addresses:** For a single Grid Master or member, enter the address configured on the NAT appliance that maps to the interface address of the LAN1 port. A single master or member that serves DNS uses this NAT address for Grid communications and—if it serves DNS—DNS messages.

For an HA Grid Master or member, enter the address configured on the NAT appliance that maps to its VIP address. An HA master uses its VIP NAT address when communicating with Grid members. An HA member that serves DNS uses its VIP NAT address for its DNS messages. It uses its LAN1 port NAT address for Grid communications.

- **Node 1 (if HA)**
 - **NAT IP Address:** Enter the address configured on the NAT appliance that maps to the interface address of the LAN1 port on Node 1. When Node 1 of an HA member is active, it uses its NAT address for Grid communications.
- **Node 2 (if HA)**
 - **NAT IP Address:** Enter the address configured on the NAT appliance that maps to the interface address of the LAN1 port on Node 2. When Node 2 of an HA member is active, it uses its NAT address for Grid communications.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

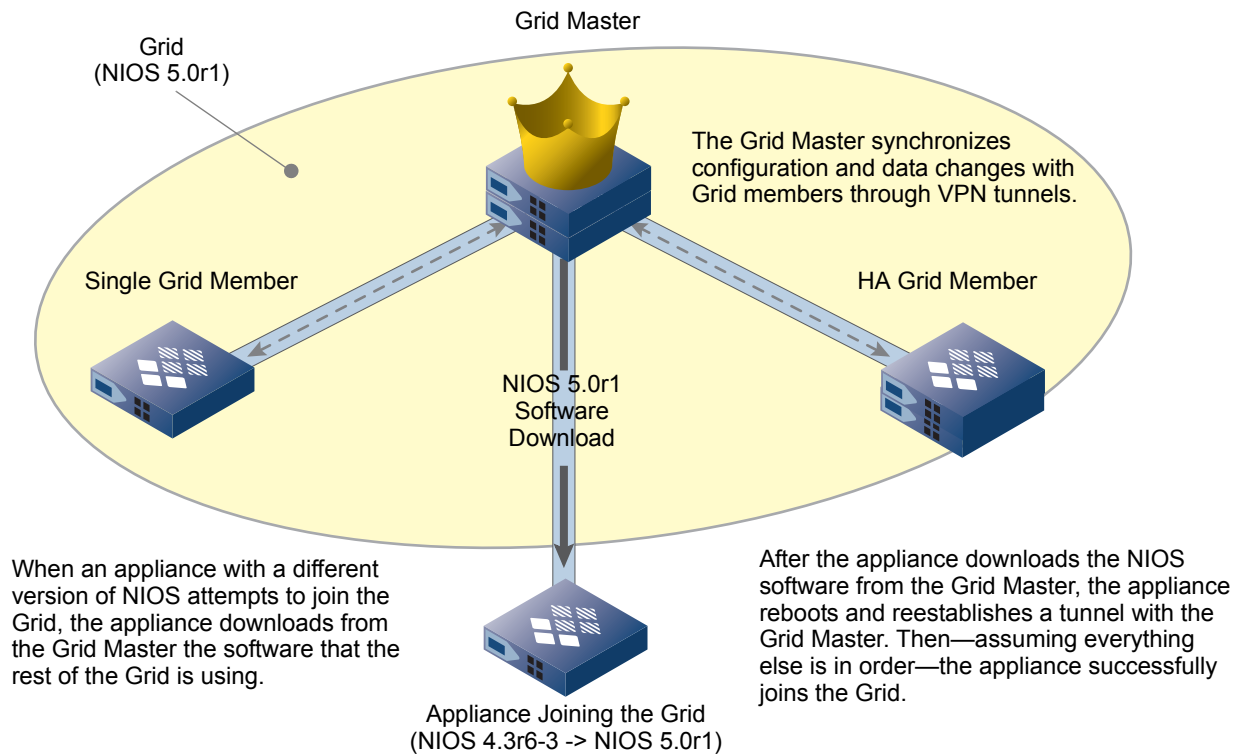
Automatic Software Version Coordination

When you add an appliance or HA pair to a Grid as a new member, it is important that it is running the same version of software as the other members in the Grid. Infoblox provides two methods for coordinating the software version:

- **Manual Upgrade and Downgrade:** Before adding an appliance or HA pair to a Grid, you can manually upgrade or downgrade the software on the appliance or HA pair to the version used by the rest of the Grid.
- **Automatic Upgrade and Downgrade:** The Grid Master automatically compares the software version of each appliance attempting to enter a Grid with that in use by the rest of Grid. If the versions do not match, the Grid Master downloads the correct version to the new appliance or HA pair.

Note: The Grid Master checks the software version every time an appliance or HA pair joins the Grid. The software version check occurs during the initial join operation and when a member goes offline and then rejoins the Grid.

Figure 5.7 Automatic Upgrade of An Appliance Joining a Grid



When a single appliance attempts to join the Grid for the first time, the following series of events takes place:

1. The appliance establishes an encrypted VPN tunnel with the Grid Master.
2. The master detects that the software version on the appliance is different from that in the rest of the Grid. For example, the appliance is running NIOS 4.3r6-3 software but the rest of the Grid is running NIOS 5.0r1 software.
3. The appliance downloads the NIOS 5.0r1 software from the Grid Master.
4. After the upgrade is complete, the NIOS application automatically restarts.
5. After the appliance reboots, it again contacts the Grid Master and step 1 is repeated. Because the software versions now match, the appliance can complete its attempt to join the Grid.

When an HA pair attempts to join the Grid for the first time, the following series of events takes place:

1. The active node of the HA pair establishes an encrypted VPN tunnel with the Grid Master.
2. The master detects that the software version on the node is different from that in the rest of the Grid. For example, the active node is running NIOS 4.3r6-3 software but the rest of the Grid is running NIOS 5.0r1 software.
3. The appliance downloads the NIOS 5.0r1 software from the Grid Master.
4. After the upgrade is complete, the NIOS application on the active node automatically restarts. This causes an HA failover.
5. The new active node (which was previously the passive node) attempts to join the Grid, repeating steps 1 – 4.
6. When the NIOS application on the currently active node restarts, there is another failover, and the currently passive node becomes active again.
7. The active node again contacts the Grid Master and step 1 is repeated. Because the software versions now match, it can complete its attempt to join the Grid.

Grid Bandwidth Considerations

Infoblox Grid technology relies upon database replication for its core functionality. When designing a Grid, it is important to consider the amount of traffic generated by this replication and the overall number of Grid members. Other communication between Grid members (such as log retrieval and monitoring functions) occurs as well. All of this traffic is securely communicated between the Grid Master and Grid members through encrypted VPN tunnels.

One component of the traffic through the tunnels is database replication traffic. There are three types to consider:

- **Complete database replication to a master candidate** — Occurs when a master candidate joins or rejoins a Grid. The Grid Master sends the complete database to a master candidate so that it has all the data it needs if it ever becomes promoted from member to master.
- **Partial database replication** — Occurs when an appliance or HA pair joins or rejoins the Grid as a regular member (which is not configured as a master candidate). The Grid Master sends it the section of the database that mainly applies just to the member.
- **Ongoing database updates** — Occurs as changes are made to the Grid configuration and data. The Grid Master sends all ongoing database updates to master candidates and individual member-specific updates to regular members.

If there are no or very few DNS dynamic updates, and no or very few DHCP lease offers and renewals issued, then this type of replication traffic is minimal.

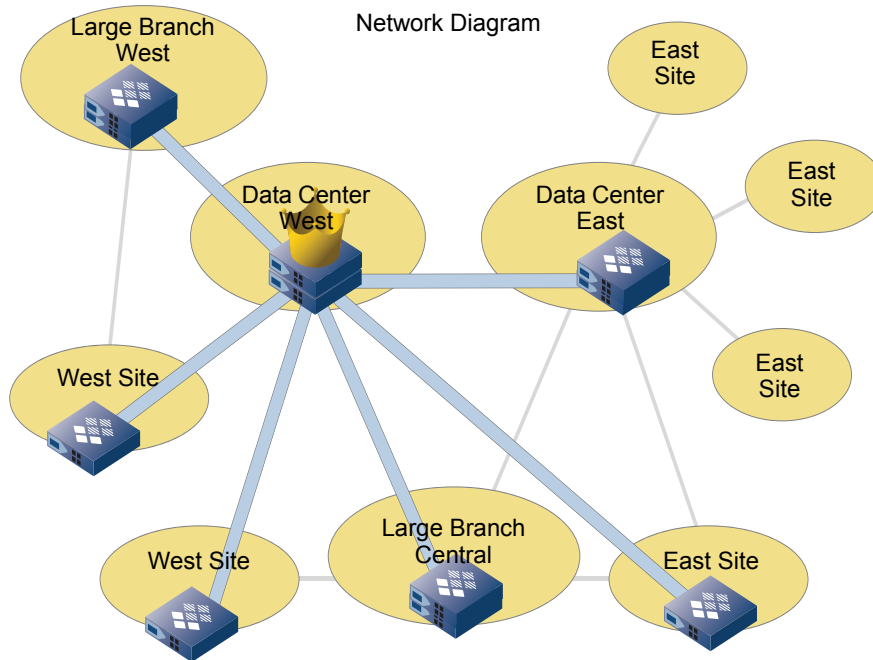
If there are many DDNS (dynamic DNS) updates (many per second) and/or many DHCP lease offers and renewals (many per second), then the replication traffic is the largest component of the VPN traffic among Grid members.

Note: A Grid Master replicates data to single members and to the active node of HA members. The active node then replicates the data to the passive node in the HA pair.

At a minimum, there must be 256 Kbps (kilobits per second) bandwidth between the Grid Master and each member, with a maximum round-trip delay of 500 milliseconds. For ongoing database updates, the amount of data sent or received is 15 Kb for every DDNS update, and 10 Kb for every DHCP lease -offer/renew. The baseline amount for heartbeat and other maintenance traffic for each member is 2 Kbps. Measure the peak DNS and DHCP traffic you see in your network to determine the bandwidth needed between the Grid Master and its members for this activity.

For example, you might decide to place your Grid members in the locations shown in [Figure 5.6](#) on page 228.

Figure 5.8 Grid Deployment



In this example, the Grid Master is optimally placed in the Data Center West. There are a total of seven members: the HA Grid Master, three HA members, and three single members. If all the members are master candidates, the Grid Master replicates all changes to the other six members. Assuming that the master receives 20 dynamic updates per minute and 40 DHCP lease renews per minute, the calculation for Grid bandwidth is:

20 DDNS updates/minute/60 secs = 0.333 DDNS updates/sec * 15 Kb = 5 Kbps * 6 members	= 30 Kbps
40 DHCP leases/minute/60 secs = 0.666 DHCP leases/sec * 10 Kb = 6.7 Kbps * 6 members	= 40.2 Kbps
2 Kbps of Grid maintenance traffic * 6 members	= 12 Kbps
Total	82.2 Kbps

Another component is the upgrade process. See [Upgrading NIOS Software](#) on page 411 for more information.

Bandwidth requirements, database size, and update rate determine the maximum size of the Grid you can deploy. Based on the various factors discussed above, you can determine the amount of bandwidth your Grid needs. If your calculations exceed the available bandwidth, then you might need to modify your deployment strategy, perhaps by splitting one large Grid into two or more smaller ones.

Note: This calculation does not take into account existing traffic other than DNS and DHCP services, so factor and adjust accordingly.

For international networks, because of bandwidth and delay requirements, a geographical grouping of Grid members might be the best approach. For example, if you have a global presence, it may make the most sense to have a North American Grid, a South American Grid, a European Grid, and an Asia/Pacific Grid.

About HA Pairs

You can configure two appliances as an HA (high availability) pair to provide hardware redundancy for core network services. An HA pair can be a Grid Master, a Grid member, or an independent appliance. The two nodes that form an HA pair—identified as Node 1 and Node 2—are in an active/passive configuration. The active node receives, processes, and responds to all service requests. The passive node constantly keeps its database synchronized with that of the active node, so it can take over services if a failover occurs. A failover is the reversal of the active/passive roles of each node; that is, when a failover occurs, the previously active node becomes passive and the previously passive node becomes active.

The appliance uses the following components in the HA functionality:

- **bloxSYNC:** An Infoblox proprietary mechanism for secure, real-time synchronization of the database that maintains the data, system configuration, and protocol service configuration between the two nodes. With bloxSYNC, the nodes continuously synchronize changes of their configurations and states. When a failover occurs, the passive node can quickly take over services. For information, see [About HA Failover](#) on page 234.
- **VRRP (Virtual Router Redundancy Protocol):** An industry-standard, MAC-level HA failover mechanism. VRRP utilizes the concept of an active and passive node that share a single VIP (virtual IP) address. When the active node that owns the VIP becomes unavailable, the passive node takes over the VIP and provides network core services. For information about VRRP, refer to *RFC3768, Virtual Router Redundancy Protocol (VRRP)* and *VRRP Advertisements* on page 235.

Using bloxSYNC and VRRP combined, if the active node fails or is taken offline for maintenance purposes, the passive node assumes the VIP and continues to respond to requests and services with minimal interruption. You can deploy an HA pair as a Grid Master, a Grid member, or an independent HA. To deploy an independent HA pair, see [Deploying an Independent HA Pair](#) on page 287. To deploy an HA Grid Master, see [Creating a Grid Master](#) on page 236.


Planning for an HA Pair

To achieve high availability, the HA and LAN1 (or VLAN) ports on both the active and passive nodes are connected to switches on the same network or VLAN. Both nodes in an HA pair share a single VIP address and a virtual MAC address so they can appear as a single entity on the network. You can also assign IPv6 addresses for each of the active and passive nodes, in addition to the IPv6 VIP address.

As illustrated in [Figure 5.9](#) on page 233, the VIP and virtual MAC addresses link to the HA port on each node. Select five IP addresses on the same network before you configure an HA pair, as follows:

- **VIP:** For core network services and for management purposes when the MGMT port is disabled. Both nodes share the same VIP.
- **Node 1 HA (active):** Source IP for the VIP and VRRP advertisements
- **Node 1 LAN1 (active):** For management through SSHv2 and listens for VRRP advertisements from the HA port
- **Node 2 HA (passive):** Listens for VRRP advertisements
- **Node 2 LAN1 (passive):** Source IP for SSL VPN to the VIP of the active node and receives bloxSYNC from the VIP

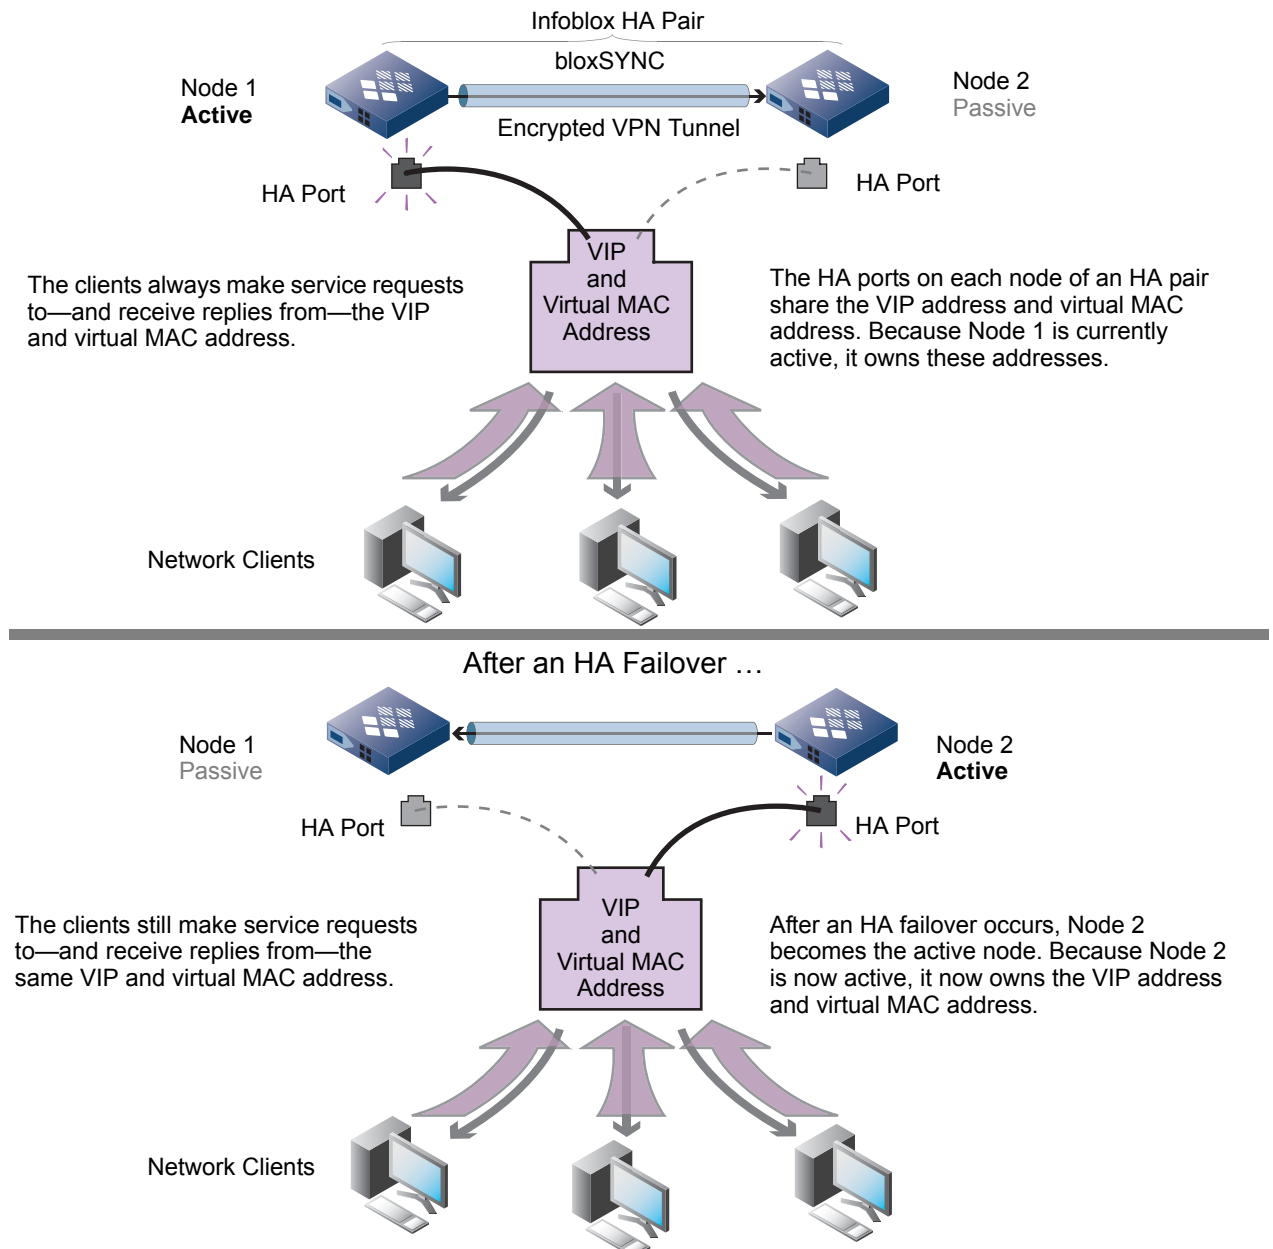
Figure 5.9 HA Pair



About HA Failover

The appliance supports HA through bloxHA™, which provides a robust failover mechanism. As described in [Planning for an HA Pair](#) on page 233, both nodes in an HA pair share a single VIP address and a virtual MAC address. The node that is currently active is the one whose HA port owns the VIP address and virtual MAC address. When a failover occurs, these addresses shift from the HA port of the previous active node to the HA port of the new active node, as illustrated in [Figure 5.10](#).

Figure 5.10 VIP Address and Virtual MAC Address and HA Failover



VRRP Advertisements

VRRP advertisements are periodic announcements of the availability of the HA node linked to the VIP. The two nodes in an HA pair include a VRID (virtual router ID) in all VRRP advertisements and use it to recognize VRRP advertisements intended for themselves. Only another appliance on the same subnet configured to use the same VRID responds to the announcements. The active node in an HA pair sends advertisements as multicast datagrams every second. It sends them from its HA port using the source IP address of the HA port (not from the VIP address) and the source MAC address 00:00:5e:00:01:vrp_id. The last two hexadecimal numbers in the source MAC address indicate the VRID number for this HA pair. For example, if the VRID number is 143, then the source MAC address is 00:00:5e:00:01:8f (8f in hexadecimal notation = 143 in decimal notation).

The destination MAC and IP addresses for all VRRP advertisements are 01:00:5e:00:00:12 and 224.0.0.18. Because a VRRP advertisement is a multicast datagram that can only be sent within the immediate logical broadcast domain, the nodes in an HA pair must be in the same subnet together.

As illustrated in [Figure 5.11](#), when you configure an HA pair, only the appliance configured to listen for VRRP advertisements with the same VRID number processes the datagrams, while all other appliances ignore them. The passive node in an Infoblox HA pair listens for these on its HA port and the active node listens on its LAN1 or LAN1 (VLAN) port. If the passive node does not receive three consecutive advertisements or if it receives an advertisement with the priority set to 0 (which occurs when you manually perform a forced failover or request the active node to restart, reboot, or shut down), it changes to the active state and assumes ownership of the VIP address and virtual MAC address.

If both nodes go offline, the one that comes online first becomes the active node. If they come online simultaneously, or if they enter a dual-active state—that is, a condition arises in which both appliances assume an active role and send VRRP advertisements, possibly because of network issues—then the appliance with the numerically higher VRRP priority becomes the active node. The priority is based on system status and events.

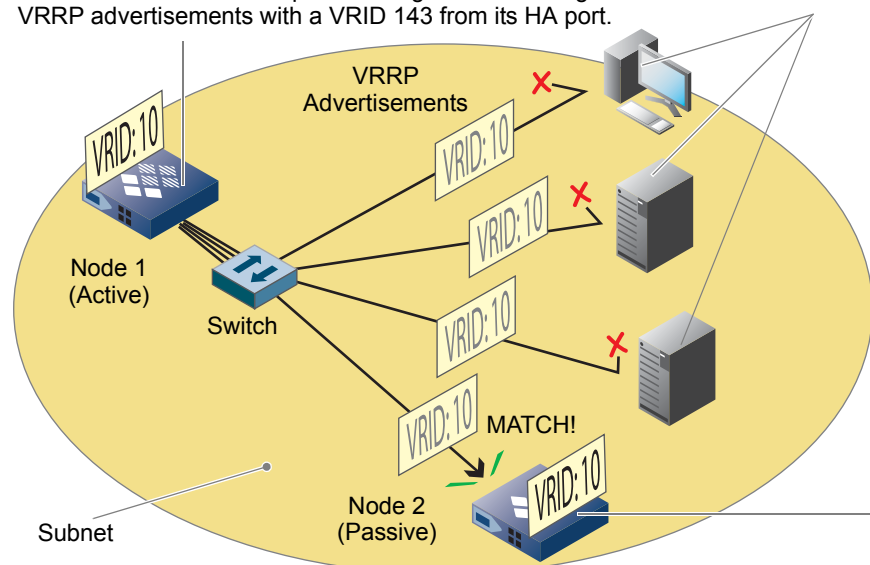
If both nodes have the same priority, then the appliance whose HA port has a numerically higher IP address becomes the active node. For example, if the IP address of the HA port on Node 1 is 10.1.1.80 and the IP address of the HA port on Node 2 is 10.1.1.20, then Node 1 becomes the active node.

For more information about VRRP, see *RFC 3768, Virtual Router Redundancy Protocol (VRRP)*.

Figure 5.11 VRRP Advertisements with a Unique VRID

After you finish configuring Node 1 of the HA pair to use VRID 143—a number that is unique for this subnet—it starts listening for VRRP advertisements with that VRID. When it does not receive any for three seconds, it becomes the active node in the HA pair and begins multicasting VRRP advertisements with a VRID 143 from its HA port.

Any device on that subnet that is not configured to listen for VRRP advertisements with VRID 143 drops the packet.

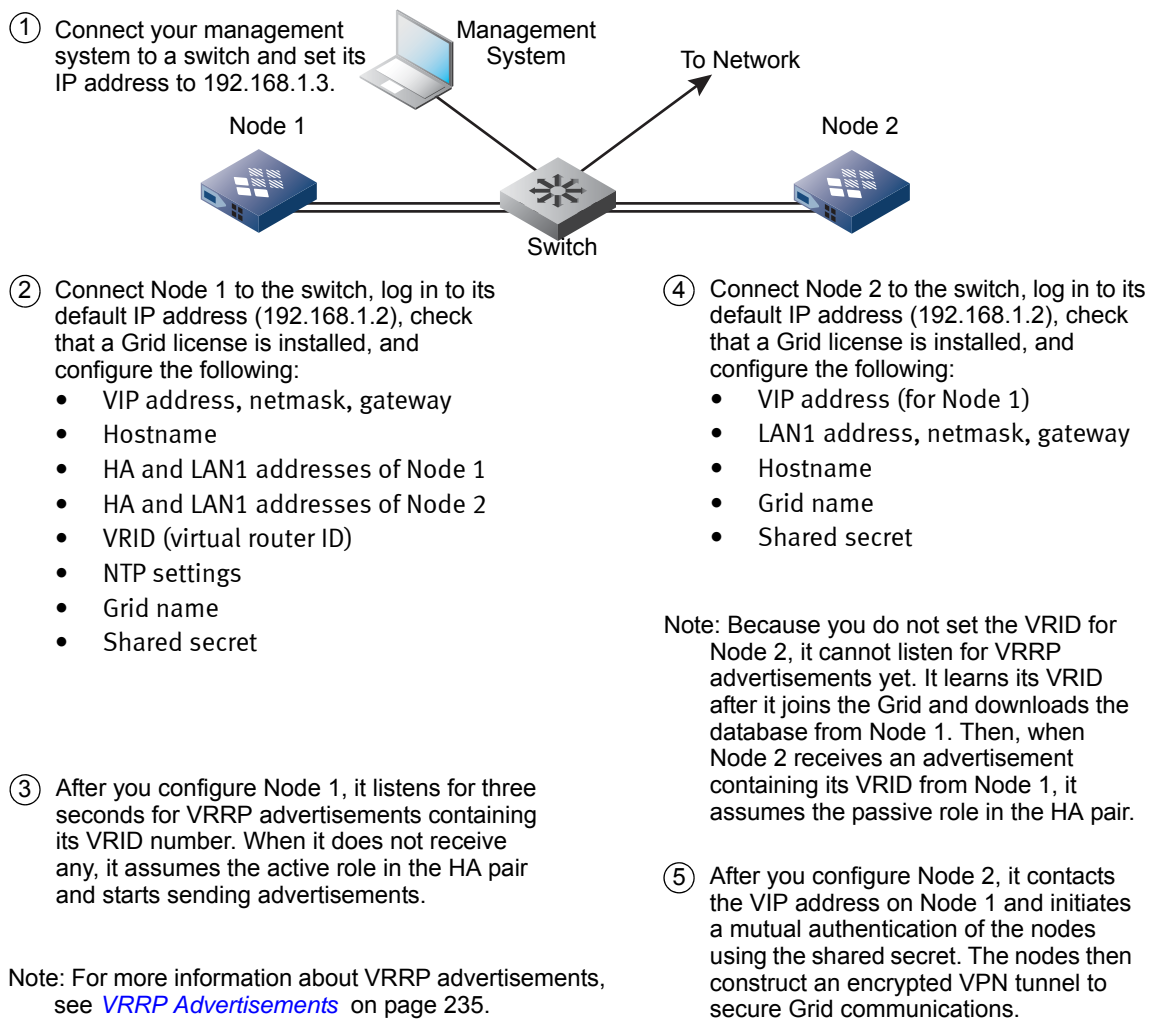


After you finish configuring Node 2 to join the HA pair, it initiates a connection with Node 1. The two appliances establish a VPN tunnel between themselves, using the HA connection name and shared secret to authenticate each other. Node 2 downloads the database from Node 1 and learns its VRID. Node 2 then begins listening for VRRP advertisements on its HA port. When it receives an advertisement from Node 1, Node 2 recognizes it and becomes the passive node.

Creating a Grid Master

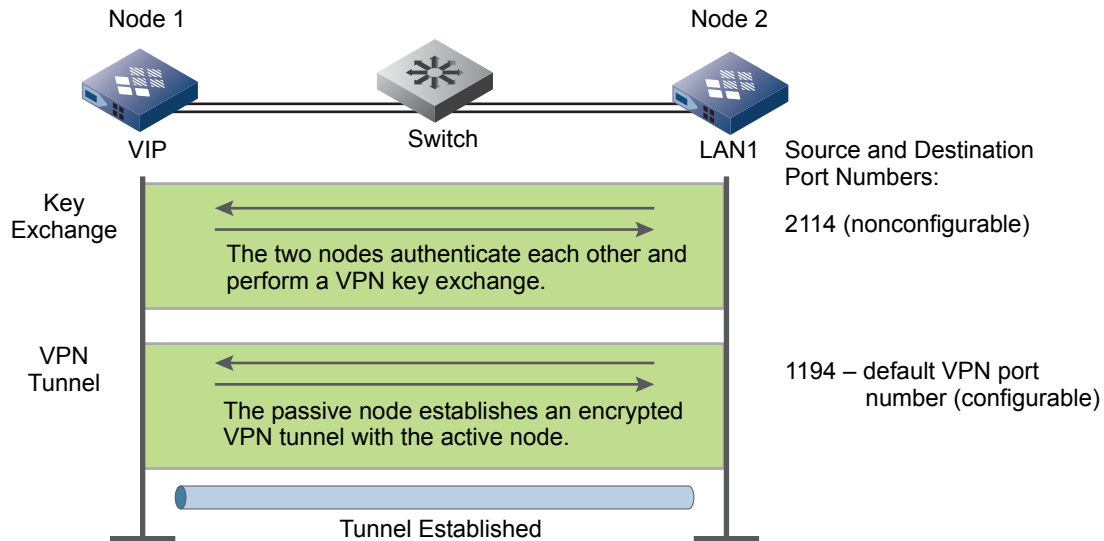
To create a Grid, you first create a Grid Master and then add members. Although the Grid Master can be a single appliance (a “single master”), a more resilient design is to use an HA pair (an “HA master”) to provide hardware redundancy. For information about HA pairs, see [About HA Pairs](#) on page 233. The basic procedure for forming two appliances into an HA master is shown in [Figure 5.12](#). All Infoblox hardware platforms, except for the Infoblox-250-A and appliances with a 50 GB disk, support configuration as a Grid Master or Grid Master candidate. For information about which vNIOS appliance supports configuration as a Grid Master, see [Supported vNIOS Appliance Configurations](#) on page 1305.

Figure 5.12 Initially Configuring a Pair of Appliances as a Grid Master



After the two nodes form an HA pair, Node 2 initiates a key exchange and creates an encrypted VPN tunnel with Node 1. The two nodes communicate between the VIP interface linked to the HA port on Node 1 and the LAN1 port on Node 2. The initialization of VPN communications between the two nodes is shown in [Figure 5.13](#) on page 237.

Figure 5.13 Establishing a VPN Tunnel for Grid Communications

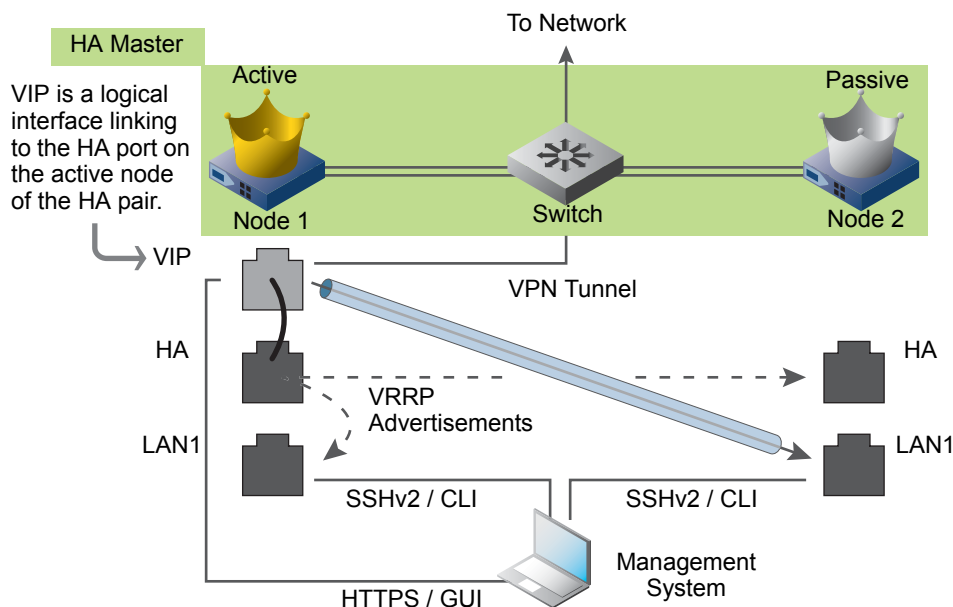


After the nodes establish a VPN tunnel between themselves, Node 1 sends Node 2 its entire database (its configuration settings and service data). Because the configuration contains the VRID (virtual router ID) for the HA pair, Node 2 starts listening for VRRP advertisements containing that VRID number. Because Node 1 is already sending such advertisements, Node 2 receives one and assumes the passive role in the HA pair.

After the initial transmission of its database, Node 1 continues to send Node 2 real-time database updates through the VPN tunnel.

Node 1 maintains the synchronization of the database throughout the Grid—which, at this point, has no other members—sends VRRP advertisements indicating its physical and network health, and—if configured to do so—provides network services. Node 2 maintains a state of readiness to assume mastership in the event of a failover. You can see the flow of HA- and Grid-related traffic from ports on the active node to ports on the passive node in [Figure 5.14](#). This illustration also shows the ports that you can use for management traffic and network service.

Figure 5.14 Traffic and Ports that an HA Grid Master Uses



Note: If you enable the MGMT port, you can only make an HTTPS connection to the IP address of the active node. If you try to connect to the IP address of the passive node, the appliance redirects your browser to the IP address of the active node.

SSHv2, however, behaves differently from HTTPS. If you enable the MGMT port and define its network settings for both nodes in the HA pair, you can make an SSHv2 connection to the IP addresses of the LAN1 and MGMT ports on both the active and passive nodes.

From the management system, you can manage the active node of the HA master by making an HTTPS connection to the VIP interface and using the GUI, and by making an SSHv2 connection to the LAN1 port (and MGMT port, if enabled) and using the CLI. If you enable the MGMT port on an HA pair, you can make an HTTPS connection through the MGMT port on the active node, and you can make an SSHv2 connection through the LAN1 or MGMT port on the active and passive nodes.

Note: For information about enabling and using the MGMT port, the Infoblox GUI, and SSH, see [Using the MGMT Port](#) on page 359, [Logging in to the GUI](#) on page 48, and [Enabling Remote Console Access](#) on page 343.

Port Numbers for Grid Communication

If connectivity between Grid members must pass through a firewall, the firewall policies must allow the initial key exchange and subsequent VPN traffic to pass. The key exchange uses UDP with a source and destination port of 2114. VPN traffic uses UDP with a default source and destination port of 1194. The VPN port number is configurable.

To configure the VPN port number:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and click **Grid Properties** -> **Edit**.
3. In the **General** tab of the *Grid Properties* editor, type a new port number in the **VPN Port** field.
4. Save the configuration.
5. When Grid Manager displays a warning indicating that a product restart is required, click **Yes** to continue.

The product automatically restarts.

A member and master first perform a handshake to authenticate each other and exchange encryption keys. Then they build an encrypted VPN tunnel between themselves. The member typically initiates both of these connections. The master only initiates a key exchange if you manually promote a member to the role of master (see [Promoting a Master Candidate](#) on page 270). [Figure 5.13](#) on page 237 shows the typical connection exchange and default port usage not only between the two nodes forming an HA pair but also between a member and master when the member joins a Grid.

The member and master key exchange occurs when an appliance joins a Grid, during master promotion, and when a member reconnects to a Grid after becoming disconnected. At all other times, Grid-related communications occur through encrypted VPN tunnels.

Grid Setup Wizard

The Grid Setup Wizard simplifies configuring a Grid. You can use it to configure an HA or single Grid Master and to join appliances to a Grid. The Grid Setup Wizard appears when you first log in to the appliance. After that, you can access it at anytime as follows:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Grid Properties** -> **Setup (Grid Setup Wizard)**.

Creating an HA Grid Master

To create a Grid, you first create a Grid Master and then add members. Although you can define a single appliance as a Grid Master, using an HA pair provides hardware redundancy for this vital component of a Grid. The following procedure explains how to put two NIOS appliances on the network and use the Grid Setup Wizard to configure them as Nodes 1 and 2 to form an HA Grid Master. For information about which vNIOS appliance supports configuration as an HA Grid Master, see [Supported vNIOS Appliance Configurations](#) on page 1305.

Configuring the Connecting Switch

To ensure that VRRP (Virtual Router Redundancy Protocol) works properly, configure the following settings at the port level for all the connecting switch ports (HA, LAN1, and LAN2):

- Spanning Tree Protocol: Disable. For vendor specific information, search for “HA” in the Infoblox Knowledge Base system at <https://support.infoblox.com>.
- Trunking: Disable
- EtherChannel: Disable
- IGMP Snooping: Disable
- DHCP Snooping: Disable or Enable Trust Interface

Note: You must disable DHCP Snooping to successfully run DHCP services on the Grid. For more information about DHCP services, see [About Infoblox DHCP Services](#) on page 774.

- Port Channeling: Disable
- Speed and Duplex settings: Match these settings on both the Infoblox appliance and switch
- Disable other dynamic and proprietary protocols that might interrupt the forwarding of packets

Note: By default, a NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between its LAN1 or LAN1 (VLAN), HA, and MGMT ports and the Ethernet ports on the connecting switch. If the two appliances fail to auto-negotiate the optimal settings, see [Modifying Ethernet Port Settings](#) on page 354 for steps you can take to resolve the problem.

Placing Both Appliances on the Network

1. Connect the power cable from each NIOS appliance to a power source and turn on the power. If possible, connect the appliances to separate power circuits. If one power circuit fails, the other might still be operative.
2. Connect Ethernet cables from the LAN1 port and the HA port on each appliance to a switch on the network.

Note: The Ethernet ports on the Infoblox-250-A, 550-A, 1050-A, 1550-A, 1552-A, 1852-A, 2000-A, Trinzic 810, Trinzic 820, Trinzic 1410, Trinzic 1420 and IB-4010 appliances are autosensing, so you can use either a straight-through or cross-over Ethernet cable for these connections.

3. Use the LCD on one appliance or make a console connection to it, and configure the network settings of its LAN1 port so that it is on the local subnet and you can reach it on the network. IPv4 addressing is supported on the LCD; ensure that you have the correct network address values before configuration of the appliance.

Note: For details about using the LCD and console, refer to the installation guide that shipped with your product.

4. Similarly, configure the LAN1 port on the other appliance so that it is in the same subnet as the first appliance.
5. Connect your management system to the network so that it can reach the IP addresses of the LAN1 ports on both appliances.

HA Master – Node 1

1. On your management system, open a browser window, and connect to https://ip_addr, where *ip_addr* is the address of the LAN1 port on Node 1. IPv4 and IPv6 values are valid, based on the LAN1 port configuration.
2. Log in using the default user name and password *admin* and *infoblox*. For detailed information about logging in to the GUI, see [Logging in to the GUI](#) on page 48.
3. Review the End-User License Agreement. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following:

- **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
- **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
- **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.

Click **I Accept**. The *Grid Setup* wizard appears.

4. On the first screen, select **Configure a Grid Master** and click **Next**.
5. On the next screen, specify the Grid properties and click **Next**:
 - **Grid Name:** Enter a text string that the two appliances use to authenticate each other when establishing a VPN tunnel between them. The default Grid name is **Infoblox**.
 - **Shared Secret:** Enter a text string that both appliances use as a shared secret to authenticate each other when establishing a VPN tunnel between them. The default shared secret is **test**.
 - **Show Password:** Select this to display the password. Clear the check box to conceal the password.
 - **Hostname:** Enter a valid domain name for the appliance.
 - **Is the Grid Master an HA pair?:** Select **Yes**.
6. On the next screen, specify the network properties and click **Next**:
 - **Virtual Router ID:** Enter the VRID (virtual router ID). This must be a unique VRID number—from 1 to 255—for this subnet.
 - **Required Ports and Addresses:** This table lists some of the network settings on the HA pair. Some fields are prepopulated by Grid Manager based on the existing configuration of the appliance. All fields are required. Enter correct information for the following by clicking the field:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IP address for the corresponding port.
 - **Subnet Mask:** Specify an appropriate subnet mask.
 - **Gateway:** Type the default gateway for the interface.
 - **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
 - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **IPv6 (VIP) Optional:** Enter the following information to add an IPv6 VIP:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IPv6 address for the Grid member on the interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef). For IB-4030 appliance, use a /128 CIDR for IPv6 while configuring multiple interfaces.
 - **Prefix Length:** Choose the CIDR netmask for the subnet to which the VIP address connects. CIDR is an alternative to classful subnet masking that organizes IP addresses into subnetworks. Also known as supernetting, CIDR allows multiple subnets to be grouped together for network routing. The prefix length ranges from 0 to 128, with common-sense values ranging from /48 to /128 due to the larger number of bits in the IPv6 address. Note that the IB-4030 supports the same netmask as the LAN1 interface or a /128 prefix.
 - **Gateway:** Do one of the following:

- Type the IPv6 address of the default gateway of the subnet to which the VIP address connects.
- Type **auto** to enable the appliance to acquire the IP address of the default gateway and the link MTU from router advertisements.

Note: You can now define a link-local address as the default IPv6 gateway and isolate the LAN segment so the local router can provide global addressing and access to the network and Internet. This is supported for both LAN1 and LAN2 interfaces as well as LAN1 and LAN2 in the failover mode.

- **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
7. Optionally, enter a new password and click **Next**. The password must be a single string (no spaces) that is at least four characters long.
 8. Select the time zone of the Grid Master and indicate whether the Grid Master synchronizes its time with an NTP (Network Time Protocol) server.
 - If you choose to enable NTP, click the Add icon and enter the IP address of an NTP server. You can enter IP addresses for multiple NTP servers.
 - If you choose to disable NTP, set the date and time for the appliance.
 - Click **Next**.
 9. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following and then click **Next**:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Email:** Enter an email address to which Infoblox sends a copy of the usage report. The email address you enter here is also displayed in the **Customer Improvement** tab in the Grid Properties editor. This is optional.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.
 10. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**. The application restarts after you click **Finish**.

Note: The Grid Setup wizard provides options such as not changing the default password and manually entering the time and date. However, changing the password and using an NTP server improve security and accuracy (respectively), and so these choices are presented here.

Record and retain this information in a safe place. If you forget the shared secret, you need to contact Infoblox Technical Support for help. When you add an appliance to the Grid, you must configure it with the same Grid name, shared secret, and VPN port number that you configure on the Grid Master.

11. Close the management window.
The configuration for Node 1 is complete.

HA Master – Node 2

1. On your management system, open a new browser window, and connect to `https://ip_addr`, where `ip_addr` is the address of the LAN1 port on Node 2. IPv4 or IPv6 values are valid.
When you enter an IPv6 address, enclose the address in square brackets (as in `https://[ip_addr]` or `https://[2001:db8::256:ABCD:EF12:34:1]`).
2. Log in using the default user name and password `admin` and `infoblox`.

3. Review the End-User License Agreement. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.

Click **I Accept**. The *Grid Setup* wizard appears.

4. On the first screen, select **Join Existing Grid** and click **Next**.
5. On the next screen, specify the Grid properties and click **Next**
 - **Grid Name:** Enter a text string that the two appliances use to authenticate each other when establishing a VPN tunnel between them. This must match the Grid name you entered for node 1.
 - **Grid Master's IP Address:** Enter the same VIP you entered for node 1.
 - **Shared Secret:** Enter a text string that both appliances use as a shared secret to authenticate each other when establishing a VPN tunnel between them. This must match your entry in node 1.
6. On the next screen verify the IP address settings of the member and click **Next**.
The last screen displays the settings you specified in the previous panels of the wizard.
7. Verify that the information is correct and click **Finish**.
The setup of the HA master is complete. From now on, when you make an HTTPS connection to the HA pair, use the VIP address.

Creating a Single Grid Master

Although using an HA master is ideal because of the hardware redundancy it provides, you can also use a single appliance as the Grid Master. Infoblox recommends frequent backups if the Grid Master is a single appliance, and there is no master candidate. For information about which vNIOs appliance supports configuration as a single Grid Master, see [Supported vNIOs Appliance Configurations](#) on page 1305.

Setting up an appliance as a single Grid Master is very easy. If the appliance has the DNSOne package with the Grid upgrade, it is already a Grid Master. You simply need to define the network settings for its LAN1 port. The various procedures for defining the network settings for the LAN1 port of a single independent appliance apply here as well; that is, you can use any of the following procedures to define the network settings for the LAN1 port of the appliance that you want to make a single Grid Master:

- LCD – See [Method 1 – Using the LCD](#) on page 277. (LCD configuration does not support IPv6 address entry.)
- Console port – [Method 2 – Using the CLI](#) on page 277.

You can also use the NIOS Grid Setup Wizard to create a single Grid Master. In addition to providing a simple method accompanied by helpful information, the setup wizard allows you to change the admin password and configure time settings for the appliance.

Using the Setup Wizard

To create a single Grid Master using the *Grid Setup* wizard:

1. Connect the power cable from the NIOS appliance to a power source and turn on the power.
2. Connect an Ethernet cable from the LAN1 port on the appliance to a switch on the network.

Note: The Ethernet ports on the Infoblox-250-A, 550-A, 1050-A, 1550-A, 1552-A, 1852-A, 2000-A, Trinzic 810, Trinzic 820, Trinzic 1410, Trinzic 1420 and IB-4010 appliances are autosensing, so you can use either a straight-through or cross-over Ethernet cable for these connections.

3. If you have not changed the default IP address (192.168.1.2/24) of the LAN1 port through the LCD or CLI—and the subnet to which you connect the appliance does not happen to be 192.168.1.0/24—put your management system in the 192.168.1.0/24 subnet and connect an Ethernet cable between your management system and the NIOS appliance.
4. Open a web browser and make an HTTPS connection to the IP address of the LAN1 port. To reach the default IP address, enter: **https://192.168.1.2**.

Several certificate warnings appear during the login process. This is normal because the preloaded certificate is self-signed (and, therefore, is not in the trusted certificate stores in your browser) and has the hostname `www.infoblox.com`, which does not match the destination IP address you entered in step 3. To stop the warning messages from occurring each time you log in to the GUI, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information about certificates, see [Creating a Login Banner](#) on page 49.

5. Log in using the default user name **admin** and password **infoblox**.
6. Review the End-User License Agreement. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.

Click **I Accept**. The *Grid Setup* wizard appears.

7. On the first screen, select **Configure a Grid Master** and click **Next**.
8. On the next screen, specify the Grid properties and click **Next**:
 - **Grid Name:** Enter a text string that the Grid Master and appliances joining the Grid use to authenticate each other when establishing a VPN tunnel between them. The default Grid name is **Infoblox**.
 - **Shared Secret:** Enter a text string that the Grid Master and appliances joining the Grid use as a shared secret to authenticate each other when establishing a VPN tunnel between them. The default shared secret is **test**.
 - **Show Password:** Select this to display the password. Clear the check box to conceal the password.
 - **Hostname:** Enter a valid domain name for the appliance.
 - **Is the Grid Master an HA pair?:** Select **No**.
9. On the next screen, configure the network settings and click **Next**:
 - **Host Name:** Enter a valid domain name for the appliance.

Note: IPv6 addressing and connectivity requires an IPv4 address configuration on the same interface.

- **IP Address:** Type the IP address for the corresponding port.
- **Subnet Mask:** Specify an appropriate subnet mask.
- **Gateway:** Type the default gateway for the interface.
- **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.

- **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
- **IPv6 (VIP) Optional:** Enter the following information to add an IPv6 VIP:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IPv6 address on the interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef). For IB-4030 appliance, use a /128 CIDR for IPv6 while configuring multiple interfaces.
 - **Prefix Length:** Choose the CIDR netmask for the subnet to which the VIP address connects. CIDR is an alternative to classful subnet masking that organizes IP addresses into subnetworks. Also known as supernetting, CIDR allows multiple subnets to be grouped together for network routing. The prefix length ranges from 0 to 128, with common-sense values ranging from /48 to /128 due to the larger number of bits in the IPv6 address. Note that the IB-4030 supports the same netmask as the LAN1 interface or a /128 prefix.
 - **Gateway:** Do one of the following:
 - Type the IPv6 address of the default gateway of the subnet to which the VIP address connects.
 - Type **auto** to enable the appliance to acquire the IP address of the default gateway and the link MTU from router advertisements.
 - **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
- 10. Optionally, enter a new password and click **Next**. The password must be a single hexadecimal string (no spaces) that is at least four characters long.
- 11. Select the time zone of the Grid Master and indicate whether the Grid Master synchronizes its time with an NTP (Network Time Protocol) server, and then click **Next**.
 - If you choose to enable NTP, click the Add icon and enter the IP address of an NTP server. You can enter IP addresses for multiple NTP servers.
 - If you choose to disable NTP, set the date and time for the appliance.
- 12. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following and then click **Next**:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Email:** Enter an email address to which Infoblox sends a copy of the usage report. The email address you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. This is optional.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.
- 13. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**. The application restarts after you click **Finish**.

Note: The *Grid Setup* wizard provides options such as not changing the default password and manually entering the time and date. However, changing the password and using an NTP server improve security and accuracy (respectively), and so these choices are presented here.

Record and retain this information in a safe place. If you forget the shared secret, you need to contact Infoblox Technical Support for help. When you add an appliance to the Grid, you must configure it with the same Grid name, shared secret, and VPN port number that you configure on the Grid Master.

The last screen of the setup wizard states that the changed settings require the appliance to restart. When you click **Finish**, the appliance restarts.

The setup of the single master is complete. From now on, when you make an HTTPS connection to the appliance, use its new IP address.

Adding Grid Members

You can add single appliances and HA pairs to a Grid, forming single members and HA members respectively. A single Grid member can be either an Infoblox appliance or a vNIOS appliance. For information about which vNIOS appliance supports configuration as an HA Grid member, see [Supported vNIOS Appliance Configurations](#) on page 1305.

You can also define an HA member on the Grid Master and then add two individual NIOS appliances to the Grid as Node 1 and Node 2 to complete the HA member you defined on the master.

New members inherit all settings that you create at the Grid level unless you override them at the member level.

The process for adding either a single appliance or HA pair to a Grid involves the following steps:

1. Adding and configuring Grid members on the Grid Master. In addition to defining the network and appliance settings for a member, you can also configure service settings before you join the member or HA pair to the Grid.
2. Joining the appliance or HA pair to the Grid. This includes defining the VIP or IP address of the Grid Master, the Grid name, and the shared secret on the single appliance or HA pair. If an appliance or HA pair cannot join the Grid because of MTU (maximum transmission unit) limitations on its network link, you can reduce the MTU that the master uses when communicating with it. See [Setting the MTU for VPN Tunnels](#) on page 270. If the Grid Master is behind a NAT device and there are members on both sides of that NAT device, you must create a NAT group, as described in [NAT Groups](#) on page 226.

In a large scale deployment of Grids across multiple sites, consider remotely provisioning your Grid members before joining them to the Grid. For more information about this feature, see [Auto-Provisioning NIOS Appliances](#) on page 250.

In situations where you want to define certain configurations on an offline Grid member and associate DNS and DHCP data to the member before deploying it, you can use the pre-provisioning feature to accomplish this. For more information, see [Pre-Provisioning NIOS Appliances](#) on page 252.

Adding a Single Member

The basic steps necessary to add a single member are as follows:

1. Define the network settings of the LAN1 port of the single appliance on the Grid Master.
2. Initiate the join Grid operation during which you specify the VIP or IP address of the Grid Master, the Grid name, and the shared secret on the single appliance. For information, see [Joining Appliances to the Grid](#) on page 249.

In addition, you can configure on the Grid Master the service settings such as DNS zones and records, DHCP networks and address ranges, and so on for a member before or after you join the appliance to the Grid. The basic steps for adding a single member are presented below.

For information on how to configure a vNIOS appliance as a Grid member, refer to the *Quick Start Guide for Installing vNIOS Software on Riverbed Services Platforms* and the *Quick Start Guide for Installing vNIOS Software on VMware Platforms*.

Configuring a Single Member on the Grid Master

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
 - **Member Type**: Specify the appliance type of the Grid member. If the member is an Infoblox appliance, select **Infoblox**, which is the default. For a vNIOS appliance, select **Riverbed**, or **Virtual NIOS** for VMware.
 - **Host Name**: Type the FQDN (fully qualified domain name) of the appliance that you are adding to the Grid.
 - **Time Zone**: If the Grid member is in a different time zone from the Grid, click **Override** and select a time zone.
 - **Comment**: Type a comment that provides some useful information about the appliance, such as its location.
 - **Master Candidate**: Select this option to designate this appliance as a master candidate. For supported vNIOS appliances, see [Supported vNIOS Appliance Configurations](#) on page 1305.
4. Enter the following information about the member that you are adding to the Grid and click **Next**:
 - **Standalone Member**: Select this option.
 - **Required Ports and Addresses**: This table lists some of the network settings on the appliance. Some fields are prepopulated by Grid Manager based on the existing configuration of the appliance. All fields are required. For a standalone member, enter information about the LAN1 port. Enter correct information for the following by clicking the field:
 - **Interface**: Displays the name of the interface. You cannot modify this.
 - **Address**: Type the IP address for the corresponding port.
 - **Subnet Mask**: Specify an appropriate subnet mask.
 - **Gateway**: Type the default gateway for the interface.
 - **VLAN Tag**: For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
 - **Port Settings**: From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **DSCP Value**: Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
 - **IPv6 LAN1 Optional**: Enter the following information to add an IPv6 LAN1:
 - **Interface**: Displays the name of the interface. You cannot modify this.
 - **Address**: Type the IPv6 address on the interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef). For IB-4030 appliance, use a /128 CIDR for IPv6 while configuring multiple interfaces.
 - **Prefix Length**: Choose the CIDR netmask for the subnet to which the VIP address connects. CIDR is an alternative to classful subnet masking that organizes IP addresses into subnetworks. Also known as supernetting, CIDR allows multiple subnets to be grouped together for network routing. The prefix length ranges from 0 to 128, with common-sense values ranging from /48 to /128 due to the larger number of bits in the IPv6 address. Note that the IB-4030 supports the same netmask as the LAN1 interface or a /128 prefix.
 - **Gateway**: Do one of the following:
 - Type the IPv6 address of the default gateway of the subnet to which the VIP address connects.
 - Type **auto** to enable the appliance to acquire the IP address of the default gateway and the link MTU from router advertisements.

- **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
 - **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
5. Optionally, define extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 6. Do one of the following:
 - Click **Save & Edit** to add the single member to the Grid and launch the editor. You can configure additional properties, such as the MTU size, or add the member to a NAT group.
 - Click **Save & New** to add the single member to the Grid and launch the wizard again to add another member.
 - Click **Save & Close** to add the single member to the Grid and close the wizard.

Adding an HA Member

The basic steps necessary to add an HA member are as follows:

1. Define the network settings of the HA pair on the Grid Master.
2. Initiate the join Grid operation, during which you specify the VIP or IP address of the Grid Master, the Grid name, and the shared secret on the HA pair. For information, see [Joining Appliances to the Grid](#) on page 249.

In addition, on the Grid Master you can configure the service settings such as DNS zones and records, DHCP networks and address ranges, and so on for a member before or after you join the HA pair to the Grid. The basic steps for adding an HA member are presented below.

Note: The procedure for adding an HA pair to a Grid when it uses the MGMT port of the active node for Grid communications differs slightly from that described below. See [Grid Communications](#) on page 362.

Configuring an HA Member on the Grid Master

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
 - **Member Type:** Specify the appliance type of the Grid member. If the member is an Infoblox appliance, select Infoblox, which is the default. For a vNIOS appliance on VMware, select **Virtual NIOS**.
 - **Host Name:** Type the FQDN (fully qualified domain name) for the HA member.
 - **Time Zone:** If you want the Grid member to have a different time zone, click **Override** and select a time zone.
 - **Comment:** Type a comment that provides some useful information about the appliance, such as its location.
 - **Master Candidate:** select this check box to designate this appliance as a master candidate. For supported vNIOS appliances, see [Supported vNIOS Appliance Configurations](#) on page 1305.
4. Enter the following information about the member that you are adding to the Grid and click **Next**:
 - **High Availability Pair:** Select this option.
 - **Virtual Router ID:** Enter a unique VRID number—from 1 to 255—for the local subnet.
 - **Required Ports and Addresses:** This table lists some of the network settings on the appliance. Some fields are prepopulated by Grid Manager based on the existing configuration of the appliance. All fields are required. For an HA pair, enter information about the following interfaces: VIP, Node 1 HA and LAN ports, Node 2 HA and LAN ports. The VIP address and the IP addresses for all the ports must be in the same subnet. Enter correct information for the following by clicking the field:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IP address for the corresponding port.

- **Subnet Mask:** Specify an appropriate subnet mask.
- **Gateway:** Type the default gateway for the interface.
- **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
- **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
- **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
- **IPv6 LAN1 Optional:** Enter the following information to add an IPv6 LAN1:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IPv6 address on the interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef). For IB-4030 appliance, use a /128 CIDR for IPv6 while configuring multiple interfaces.
 - **Prefix Length:** Choose the CIDR netmask for the subnet to which the VIP address connects. CIDR is an alternative to classful subnet masking that organizes IP addresses into subnetworks. Also known as supernetting, CIDR allows multiple subnets to be grouped together for network routing. The prefix length ranges from 0 to 128, with common-sense values ranging from /48 to /128 due to the larger number of bits in the IPv6 address. Note that the IB-4030 supports the same netmask as the LAN1 interface or a /128 prefix.
 - **Gateway:** Do one of the following:
 - Type the IPv6 address of the default gateway of the subnet to which the VIP address connects.
 - Type **auto** to enable the appliance to acquire the IP address of the default gateway and the link MTU from router advertisements.
 - **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
 - **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.

Note: When the system operates in HA mode, should the IPv6-addressed VIP value be deleted, the IPv6 address of the HA port will also be deleted.

5. Optionally, define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Do one of the following:
 - Click **Save & Edit** to add the HA member to the Grid and launch the editor. You can configure additional properties, such as the MTU size, or add the member to a NAT group.
 - Click **Save & New** to add the HA member to the Grid and launch the wizard again to add another member.
 - Click **Save & Close** to add the HA member to the Grid and close the wizard.

Changing the Member Type

When you change the **Member Type** from *Infoblox* to *Virtual NIOS*, Infoblox displays an error indicating that the network port of a vNIOS member must be set to Automatic. If you encounter this error, follow the steps mentioned below to change the **Member Type** to **Virtual NIOS**:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the **Toolbar** and click **Add** -> **Add Grid Member**.

3. In the *Add Grid Member* wizard, leave the **Member Type** as **Infoblox**, fill other details and click **Next**.
4. In the **Network** tab select **High Availability Pair**.
5. Change the port settings to **Automatic** for **Node1 HA**.
6. Select **Standalone Member**.
7. Click **Previous** and change the **Member Type** to **vNIOS**.

Joining Appliances to the Grid

You can use the Grid Setup Wizard or access the *Join Grid* dialog box to join appliances to a Grid. The Grid Setup Wizard launches when you first log in to an appliance. You can also launch it from the Toolbar as described in [Grid Setup Wizard](#) on page 238.

To join a single appliance and HA pair to a Grid using the Grid Manager GUI:

1. Log in to the appliance or HA pair that you want to add to the Grid. The appliance or HA pair must be online and able to reach the Grid Master.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
3. Expand the Toolbar and click **Join Grid**.
4. In the *Join Grid* dialog box, enter the following:
 - **Virtual IP of Grid Master:** Type the VIP address of the HA Grid Master or the LAN1 address of the single Grid Master for the Grid to which you want to add the appliance.
 - **Grid Name:** Type the name of the Grid.
 - **Grid Shared Secret:** Type the shared secret of the Grid.
 - **Use MGMT port to join Grid:** If you have already enabled the MGMT port (see [Grid Communications](#) on page 362), this option becomes available. Select it to connect to the Grid through the MGMT port.

5. Click **OK** to begin the join operation.

To confirm that the appliance has successfully joined the Grid, log in to the Grid Master and navigate to the **Grid** tab, select the **Grid Manager** -> **Members** tab. This panel lists the Grid members. Check the icon in the Status column of the newly added member. (green = the appliance has joined the Grid and is functioning properly; yellow = the appliance is in the process of joining the Grid; red = the appliance has not joined the Grid). You can also use the CLI command `set network` to join an appliance to a Grid.

To join a single appliance and HA pair to a Grid using the Grid Setup Wizard:

1. Log in to the appliance or HA pair that you want to add to the Grid. The appliance or HA pair must be online and able to reach the Grid Master.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
3. Expand the Toolbar and click **Grid Properties** -> **Setup (Grid Setup Wizard)**.
4. On the next screen, specify the Grid properties and click **Next**
 - **Grid Name:** Enter a text string that the two appliances use to authenticate each other when establishing a VPN tunnel between them. This must match the Grid name you entered for node 1.
 - **Grid Master's IP Address:** Enter the same VIP you entered for node 1.
 - **Shared Secret:** Enter a text string that both appliances use as a shared secret to authenticate each other when establishing a VPN tunnel between them. This must match your entry in node 1.
5. On the next screen verify the IP address settings of the member and click **Next**
6. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**.

To confirm that the appliance has successfully joined the Grid, log in to the Grid Master and navigate to the **Grid** tab, select the **Grid Manager** -> **Members** tab. This panel lists the Grid members. Check the icon in the Status column of the newly added member. (Green = The appliance has joined the Grid and is functioning properly; Yellow = The appliance is in the process of joining the Grid; Red = The appliance has not joined the Grid). You can also use the CLI command `set network` to join an appliance to a Grid.

Auto-Provisioning NIOS Appliances

In addition to using the Grid Setup Wizard or access the *Join Grid* dialog box to join appliances to a Grid, you can set up an appliance using the auto-provisioning feature, which allows a DHCP server to automatically assign an IP address to the appliance. You can then join the auto-provisioned appliance to the Grid.

Auto-provisioning is enabled by default for physical appliances, but it is not supported for vNIOS appliances. When you connect the appliance to the network, a lease request is automatically sent to the DHCP server. The DHCP server fingerprints the client as “Infoblox Appliance,” as the DHCP client provides the unique option sequence (1,28,2,2,3,3,15,6,12) and vendor ID (INFOBLOX). The DHCP server assigns a DHCP lease and a dynamic IP address to the appliance. If the DHCP lease request fails, the default IP address (192.168.1.2) is assigned to the appliance. The DHCP client tries to send the lease request for a duration of one minute when the appliance is either in the factory default state or in the auto-configured default IP address state after a reboot. If you do not use auto-provisioning to set up the appliance, then you can wait one minute before connecting the appliance to the network. Otherwise, the DHCP server will assign a dynamic IP address to the appliance. Note that if you have already set the IP address for the appliance through the Infoblox CLI, GUI, or API, then auto-provisioning is disabled for the appliance and the lease address is not requested. When auto-provisioning is enabled for an appliance, the DNS, DHCP, FTP, TFTP, HTTP, NTP, bloxTools, Captive Portal, Reporting services, as well as backup and restore are disabled for the member until a static IP address is set for the appliance. You can join a single appliance or HA pair to the Grid. After the appliance joins the Grid, the static IP address is set for the appliance.

Complete the following to set up an appliance using auto-provisioning and to join the auto-provisioned appliance to the Grid Master:

1. Connect the appliance to a network using an Ethernet cable, connect it to a power source, and then turn on the power. For information about cabling the appliance to a network and powering the appliance, refer to the user guide or installation guide that ships with the product.

A lease request is automatically sent to the DHCP server that assigns a DHCP lease and a dynamic IP address to the appliance. The DHCP client tries to send the lease request for a duration of one minute and if the request fails, the default IP address (192.168.1.2) is assigned to the appliance.

2. Join the appliance to the Grid Master. You can join the auto-provisioned appliance to the Grid Master using the *Connect* dialog box. For more information, see [Joining Auto-Provisioned Appliances to the Grid](#) on page 250. You can also join an appliance to the Grid using the *Join Grid* dialog box. For more information, see [Joining Appliances to the Grid](#) on page 249.

A static IP address is set and auto-provisioning is automatically disabled for the appliance after it joins the Grid.

Note: When auto-provisioning is disabled for an appliance and the network address is not preserved, auto-provisioning will be re-enabled and a DHCP lease request is sent to the DHCP server if you reset the appliance using the CLI command `reset all` or reset the database using the CLI command `reset database`. However, if the static IP address for an appliance is set and network settings are preserved, auto-provisioning will be re-enabled for the appliance but the lease address will not be requested if you reset the database using the CLI command `reset database`.

Joining Auto-Provisioned Appliances to the Grid

You can join a predefined appliance with a DHCP lease to the Grid Master using the *Connect* dialog box. You can join a single appliance or an HA pair to the Grid Master. For an HA pair, the member which is offline will join the Grid Master and it will become the active node. When both the members of an HA pair are offline, Node 1 of an HA pair is joined to the Grid Master.

Only superusers can join a Grid member to the Grid Master. If the Grid member fails to join the Grid, then the remote console is enabled for the appliance and you can join the appliance to a Grid through the remote console. You can log in to the remote console using the user name, **admin** and the Grid shared secret as the password.

To join a single appliance or an HA pair to a Grid Master, complete the following:

1. Log in to the Grid Master. Note that the single appliance or the HA pair must be online and the Grid Master must be able to reach the appliance.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
3. Add the appliance as a Grid member. For information about adding Grid members to the Grid, see [Adding Grid Members](#) on page 245.
4. Select the Grid member that you want to join to the Grid Master, expand the Toolbar and click **Connect**.
5. The following fields are displayed in the *Connect* dialog box:
 - **Host Name:** The name of the member.
 - **Configured IPv4 Address:** The IPv4 address of the member.
 - **Site:** The site to which the IP address belongs. This is one of the predefined extensible attributes.
 - **Temporary IPv4 Address:** Enter the IPv4 address of the DHCP lease or click **Select** to select the DHCP lease. Grid Manager displays the *Lease Selector* dialog box from which you can select the DHCP lease. Note that the *Lease Selector* displays the active DHCP leases which are fingerprinted as “Infoblox Appliance”.
6. Click **Next** to retrieve the appliance information.
 The Grid Master uses SSL to connect to the appliance and it gets the appliance information. Grid Manager displays the following information for the appliance:
 - **Remote Appliance Type:** The appliance type.
 - **Remote Appliance Serial Number:** The serial number of the appliance.
 - **Licenses:** Grid Manager displays the Grid license and the licenses that are pre-provisioned on the member. It displays the following information:
 - **Type:** The license type.
 - **String:** The license string. If the license string is not displayed, you can enter or paste it here.
7. Click **Connect** to join the appliance to the Grid Master.

To confirm that the appliance has successfully joined the Grid, check the status in the **Status** column of the newly added member. (Green = The appliance has joined the Grid Master and is functioning properly; Yellow = The appliance is in the process of joining the Grid Master; Red = The appliance has not joined the Grid Master).

Pre-Provisioning NIOS Appliances

Before joining a member to the Grid, you can first enable provisional licenses and make necessary configurations on the offline member, which allows DNS and DHCP data to be associated with the member prior to its deployment. Note that pre-provisioned members are treated as offline members. There are a few guidelines to consider before you pre-provision a member. For more information about the guidelines, see [Guidelines for Pre-provisioning Offline Grid Members](#) on page 252.

When you add a new member to the Grid, the **Pre-Provisioning** tab is displayed in the *Grid Member Properties* editor. You can pre-provision the member by defining its hardware model and enable certain provisional licenses through the **Pre-Provisioning** tab. This tab is not displayed after the member successfully joins the Grid. NIOS supports the following provisional licenses: DHCP, DNS, Microsoft Management, FireEye, and RPZ (Response Policy Zone). You must enable provisional licenses before you can make supported configurations on the pre-provisioned member. For more information about these licenses, see [About Provisional Licenses](#) on page 253.

To pre-provision an offline Grid member and join it to the Grid at a later time, complete the following:

1. Add a new single member or HA member to the Grid, as described in [Adding a Single Member](#) on page 245 or [Adding an HA Member](#) on page 247.
2. Pre-provision the offline member, as described in [Configuring Pre-Provisioned Members](#) on page 253.
3. Configure services to use the pre-provisioned member.
4. Obtain permanent licenses you have specified for pre-provisioning and use the `set license` CLI command to install the licenses on the member. For more information about CLI commands, refer to the *Infoblox CLI Guide*.
5. Join the pre-provisioned member to the Grid, as described in [Joining Appliances to the Grid](#) on page 249. For guidelines about joining pre-provisioned members, see [Joining Pre-Provisioned Members to the Grid](#) on page 254.

Guidelines for Pre-provisioning Offline Grid Members

Before you pre-provision a Grid member, consider the following:

- A pre-provisioned Grid member is an offline member. When you upgrade a Grid that has a pre-provisioned member, the upgrade behaves the same way as it does when you upgrade the Grid that has an offline member. Note that you cannot pre-provision a member or update its settings during a scheduled upgrade. For more information about upgrades, see [About Upgrades](#) on page 406.
- You cannot change the pre-provisioned member configuration after you save it. To change the configuration, you must first delete the member and pre-provision it again. If you want to delete certain provisional licenses or change the hardware model for the pre-provisioned member, you must also first delete the existing member and define a new one. For information about deleting a member, see [Removing a Grid Member](#) on page 270.
- When you assign a network, zone, or IPv4 DHCP failover association to a pre-provisioned member, the **Restart Service** button is not displayed. If you restart any service on a pre-provisioned member, no action is actually taken even though you may receive a message indicating that the operation may take a few minutes. When you join the member to the Grid, NIOS will run respective member services on the joined member. For more information about service restarts, see [Restarting Services](#) on page 386.
- NIOS allows you to backup information about the pre-provisioned member. When you perform a forced restore however, NIOS does not restore the pre-provisioned licenses if you have already installed permanent NIOS licenses on the corresponding member. For more information about backup and restore, see [Backing Up and Restoring Configuration Files](#) on page 423.
- You can use **Manage Member Services** to manage the pre-provisioned member services. For more information, see [Monitoring Member Services](#) on page 1012.

Configuring Pre-Provisioned Members

The pre-provisioning feature is disabled by default. You must select a supported hardware model for the member to enable this feature.

To pre-provision an offline member, login to the Grid Master and complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Add** -> **Add Grid Member** from the Toolbar.
2. In the *Add Grid Member* wizard, add a new member as described in [Adding a Single Member](#) on page 245.
3. After you add the member to the Grid, select the member in the **Members** tab and click the Edit icon.
4. In the *Grid Member Properties* editor, select the **Pre-Provisioning** tab, and complete the following:
 - **Member Type:** Displays the member type that you have selected in the **General** tab. The pre-provisioning feature is supported only for **Infoblox**, **Virtual NIOS**, and **Riverbed** member types. Note that you must select a hardware model for the member in order to enable the pre-provisioning feature.
 - **Hardware Model:** Select the hardware model from the drop-down list. Grid Manager displays only the supported hardware models for the specified member type. Once you select the hardware model, the pre-provisioning feature is enabled for the member. NIOS allows you to pre-provision HA members that have the same or different hardware models for Node 1 and Node 2. A few hardware specific features, such as DSCP, VLAN, LAN2, and LOM (Light Out Management), are enabled based on the pre-provisioned hardware model you specify here.
 - **Provisional Licenses:** Select the licenses that you want to enable for the pre-provisioned member. You can select the licenses only after you have specified the hardware model for the member. Once you select and enable a license, you can no longer modify the hardware model for the member. Note that the permanent licenses that you later add to the member must include the ones that are specified for pre-provisioning.
5. Save the configuration.

Note: After you save the configuration, you can no longer modify the hardware model for the member. You also cannot disable any provisional licenses, though you can add new ones. To disable provisional licenses, you must first remove the pre-provisioned member and then configure a new one.

About Provisional Licenses

If a member has never joined a Grid, you can pre-provision this member provided that you define the hardware model for the member and assign provisional licenses to it. Provisional licenses are not permanent NIOS licenses. Though they do not have expiration dates or validity periods, you must replace these licenses with corresponding permanent licenses before you join the member to the Grid.

Note: Before you join the member to the Grid, use the CLI command `set license` to add corresponding permanent licenses that you have specified for pre-provisioning. For information about CLI commands, refer to the *Infoblox CLI Guide*.

NIOS supports the following provisional licenses: DHCP (dhcp), DNS (dns), Microsoft Management (ms_management), FireEye (fireeye), and RPZ (rpz).

After you configure the offline member, you can select the pre-provisioned member from the corresponding wizards and editors based on the required license(s). The following table lists the wizards and editors from which you can select a pre-provisioned member when required pre-provisioned licenses are enabled:

Wizards and editors from which you can select a pre-provisioned member	Required license(s)
DNS Zones and Name Server Groups	dns
DHCP IPv4 and IPv6 networks	dhcp

Wizards and editors from which you can select a pre-provisioned member	Required license(s)
IPv4 DHCP Failover Association	dhcp
Microsoft servers Note that the initial synchronization with Microsoft servers is read-only. When you join the appliance to the Grid, the appliance removes all Microsoft management objects that you have configured on the Microsoft servers after the synchronization. The configuration on the Microsoft servers will replace the configuration on the NIOS appliance.	ms_management
FireEye integrated zones	fireeye, dns
RPZs (Response Policy Zones)	rpz, dns

Note: If you configure a DHCP Failover using an online member and a pre-provisioned member, assign it to a range, and start DHCP service, no addresses will be served because the initial synchronization does not happen due to the pre-provisioned offline member. NIOS logs the following message in the syslog:

```
2013-12-24T08:37:23+00:00 daemon (none) dhcpd[8790]: info DHCPDISCOVER from
cb:86:a8:45:6c:5c via 10.120.21.236: not responding (recovering)
```

Joining Pre-Provisioned Members to the Grid

Before you join a pre-provisioned member to the Grid, ensure that you verify the member type, hardware model, and provisional licenses for the member. For information about how to join a member to the Grid, see [Joining Appliances to the Grid](#) on page 249.

Note the following about joining a pre-provisioned member to the Grid:

- If you install fewer permanent licenses than the specified provisional licenses, you cannot join the member to the Grid.
- If the pre-provisioned member does not have any provisional licenses enabled, you can join the member to the Grid provided that you install a permanent Grid license on the member.
- You must install at least the set of permanent licenses that were specified for pre-provisioning along with any other needed licenses, except for the following:
 - You can join the member to the Grid if the pre-provisioned member is a vNIOS virtual appliance and has only the DNS license enabled, and you install both the vNIOS and DNS licenses on the member.
 - Similarly, you can join the member to the Grid if the pre-provisioned member is a vNIOS virtual appliance and has both DNS and DHCP licenses enabled, and you install the vNIOS, DNS, and DHCP licenses on the member.
- After you successfully join the pre-provisioned member to the Grid, provisional licenses are removed and permanent licenses take effect.
- After the member joins the Grid successfully, the **Pre-Provisioning** tab is not displayed in the *Grid Member Properties* editor.

Configuration Example: Configuring a Grid

In this example, you configure seven NIOS appliances in a Grid serving internal DHCP and DNS for an enterprise with the domain name corp100.com. There are four sites: HQ and three branch offices. A hub-and-spoke VPN tunnel system connects the sites, with HQ at the hub. The distribution and roles of the NIOS appliances at the four sites are as follows:

- HQ site (four appliances in two HA pairs):
 - HA Grid Master – hidden primary DNS server
 - HA member – secondary DNS server and DHCP server for HQ
- Site 1 (two appliances in an HA pair): HA member – secondary DNS server and DHCP server for Site 1
- Site 2 (one appliance): single member – secondary DNS server and DHCP server for Site 2

Note: When adding an Infoblox appliance to an existing Grid, you must first check whether the Grid is running the minimum required software release of the appliance. For information, refer to the document, *Minimum Required Release Software for Hardware Platforms*, that was shipped with your product.

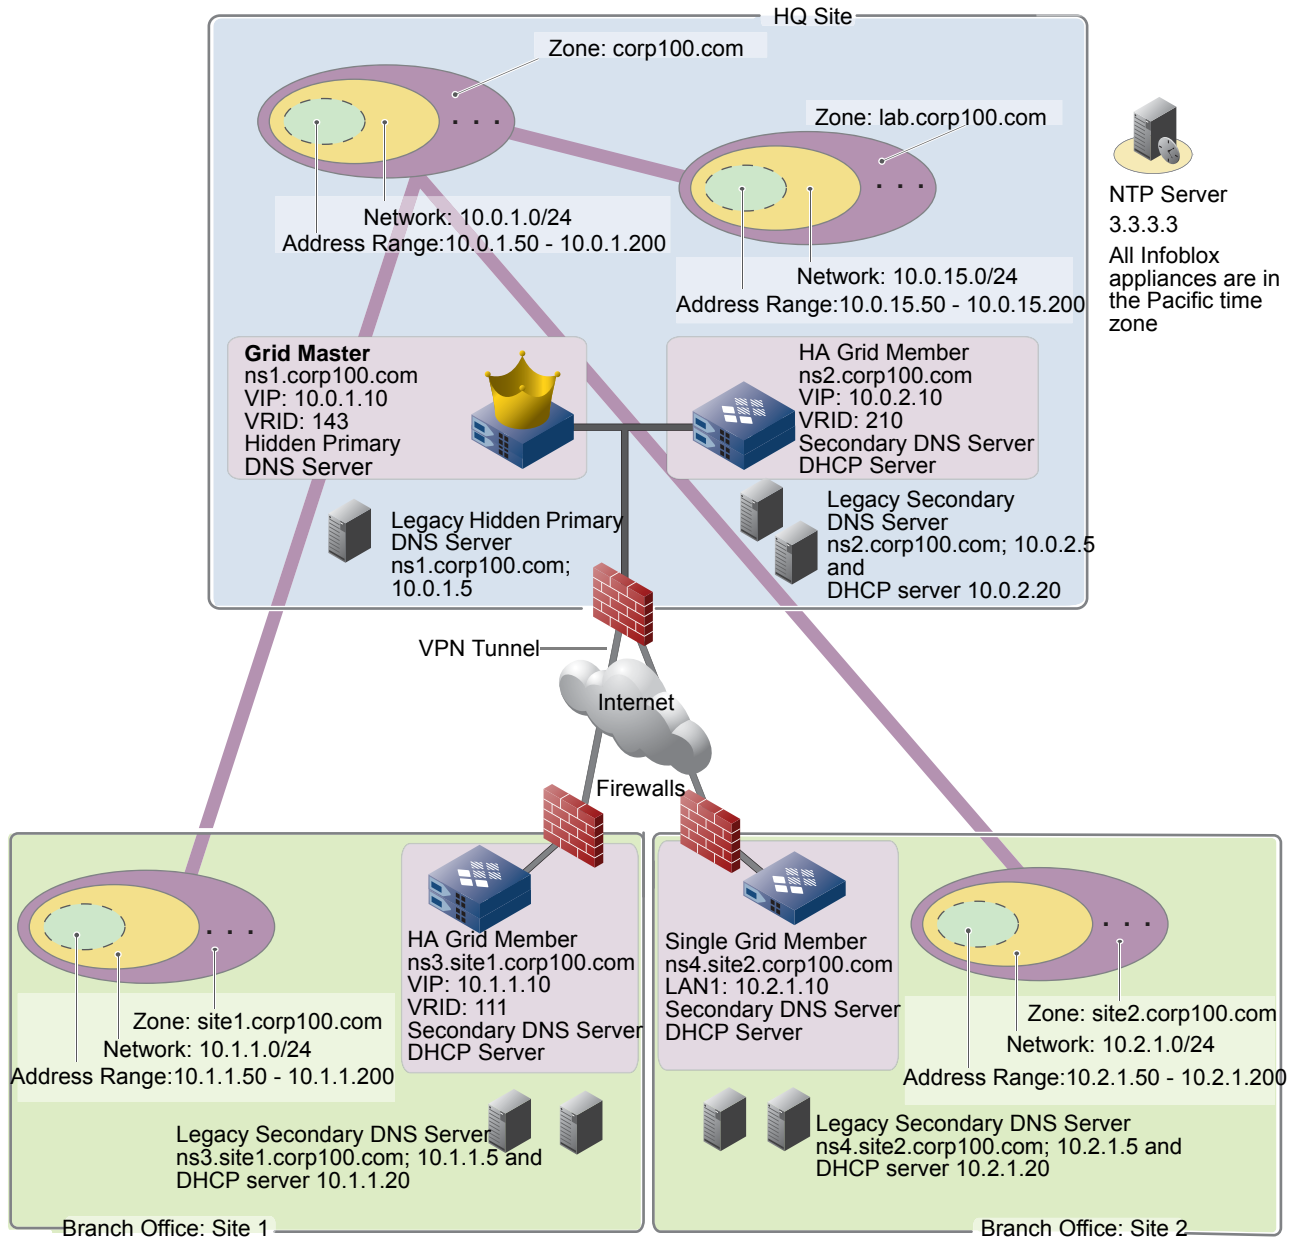
To create a Grid, you first create a Grid Master and then add members. The process involves these three steps:

1. Configuring two appliances at HQ as the Grid Master. See [Create the Grid Master](#) on page 257.
2. Logging in to the Grid Master and defining the members that you want to add to the Grid; that is, you configure Grid member settings on the Grid Master in anticipation of later joining those appliances to the Grid. See [Define Members on the Grid Master](#) on page 259.
3. Logging in to the individual appliances and configuring them so that they can reach the Grid Master over the network and join the Grid. See [Join Appliances to the Grid](#) on page 260.

After creating the Grid and adding members, you use the Data Import Wizard to import DHCP and DNS data from legacy servers. See [Import DHCP Data](#) on page 262 and [Import DNS Data](#) on page 263.

Finally, you transition DHCP and DNS service from the legacy servers to the Infoblox Grid members. See [Enable DHCP and Switch Service to the Grid](#) on page 267.

Figure 5.15 Network Diagram



Cable All Appliances to the Network and Turn On Power

Cable the NIOS appliances to network switches. After cabling each appliance to a switch and connecting it to a power source, turn on the power. For information about installing and cabling the appliance, refer to the user guide or installation guide that ships with the product.

1. At HQ and Site 1, connect Ethernet cables from the LAN1 and HA ports on the appliances in each HA pair to a switch, connect the appliances to power sources, and turn on the power for each appliance.

Note: When connecting the nodes of an HA pair to a power source, connect each node to a different power source if possible. If one power source fails, the other might still be operative.

2. At Site 2, connect an Ethernet cable from the LAN1 port on the single appliance to a switch, connect the appliance to a power source, and turn on the power for that appliance.

Create the Grid Master

Note: IPv6 addressing is fully supported on Infoblox Grid Masters, HA pairs and standalone HA pairs and appliances. Examples in the sections of this chapter use IPv4.

Configure two appliances at HQ to be the two nodes that make up the HA pair forming the Grid Master.

Grid Master – Node 1

1. By using the LCD or by making a console connection to the appliance that you want to make Node 1 of the HA pair for the Grid Master, change the default network settings of its LAN1 port to the following:
 - IP Address: 10.0.1.6
 - Netmask: 255.255.255.0
 - Gateway: 10.0.1.1
2. Connect your management system to the HQ network, open a browser window, and connect to <https://10.0.1.6>.
3. Log in using the default user name and password admin and infoblox.
4. Review the End-User License Agreement and click **I Accept**.
The Grid Setup Wizard appears.
5. On the first screen, select **Configure a Grid Master** and click **Next**.
6. Specify the Grid properties:
 - **Grid Name:** Enter corp100.
 - **Shared Secret:** Enter Mg1kW17d.
 - **Show Password:** Clear the check box to conceal the password.
 - **Hostname:** Enter ns1.corp100.com.
 - **Is the Grid Master an HA pair?:** Select **Yes**.
7. Specify the network properties and click **Next**:
 - **Virtual Router ID:** Enter 143.
 - **Required Ports and Addresses:** Enter the following to set up the HA pair:

Interface	Address	Subnet Mask	Gateway	Port Settings
VIP	10.0.1.10	255.255.255.0	10.0.1.1	Automatic
Node 1 HA	10.0.1.7	255.255.255.0	10.0.1.1	Automatic
Node 2 HA	10.0.1.9	255.255.255.0	10.0.1.1	Automatic

Interface	Address	Subnet Mask	Gateway	Port Settings
Node 1 LAN1	10.0.1.6	255.255.255.0	10.0.1.1	Automatic
Node 2 LAN1	10.0.1.8	255.255.255.0	10.0.1.1	Automatic

8. Enter a new password: 1n85w2lF. Retype it and click **Next**.
9. Complete the following:
 - **Time zone:** Select (UTC – 8:00 Pacific Time (US and Canada), Tijuana
 - Enable NTP, click the Add icon and enter the IP address of the NTP server: 3.3.3.3
10. Click **Finish**.
When you click Finish, the Infoblox GUI application restarts.

Grid Master – Node 2

1. By using the LCD or by making a console connection to the appliance that you want to make Node 2 of the HA pair for the Grid Master, change the default network settings of its LAN1 port to the following:
 - IP Address: 10.0.1.8
 - Netmask: 255.255.255.0
 - Gateway: 10.0.1.1
2. In the login window, type 10.0.1.8 in the Hostname field.
3. Log in using the default user name and password admin and infoblox.
4. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box.
5. Expand the Toolbar and click **Join Grid** and specify the following:
 - **Virtual IP of Grid Master:** 10.0.1.10.
 - **Grid Name:** Enter corp100.
 - **Grid Shared Secret:** Enter Mg1kW17d.
6. Confirm the configuration, and then on the last screen of the wizard, click Finish.
The HTTPS session terminates, but the login window remains open.
7. In the login window, type 10.0.1.10 (the VIP address for the Grid Master) in the Hostname field.
8. Log in using the default user name admin and the password 1n85w2lF.
9. To check the status of the two nodes of the HA Grid Master, navigate to the **Grid** tab, select the **Grid Manager** -> **Members** tab. This panel lists the Grid members. Check the icon in the Status column of the Grid Master. (green = the appliance has joined the Grid and is functioning properly; yellow = the appliance is in the process of joining the Grid; red = the appliance has not joined the Grid). You can also use the CLI command `set network` to join an appliance to a Grid. Check that the status indicators are all green in the Detailed Status panel.

During the joining process, an appliance passes through the following four phases:

1. Offline – the state when a Grid member—in this case, the second node of the HA pair composing the Grid Master—is not in contact with the active node of the master
2. Connecting – the state when an appliance matching a member configuration contacts the master to join the Grid and negotiates secure communications and Grid membership
3. Synchronizing – the master transmits its entire database to the member
4. Running – the state when a member is in contact with the master and is functioning properly

Note: Depending on the network connection speed and the amount of data that the master needs to synchronize with the member, the process can take from several seconds to several minutes to complete.

Define Members on the Grid Master

Before logging in to and configuring the individual appliances that you want to add to the Grid, define them first on the Grid Master.

HQ Site – HA Member

1. From the **Grid** tab, select the **Grid Manager** -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, complete the following and click **Next**:
 - **Host Name**: Enter **ns2.corp100.com**.
 - **Comment**: Enter **HQ Site - ns2.corp100.com**.
4. Enter the following information about the member that you are adding to the Grid and click **Save & Close**:
 - **High Availability Pair**: Select this option.
 - **Virtual Router ID**: 210
 - **Required Ports and Addresses**:

Interface	Address	Subnet Mask	Gateway	Port Settings
VIP	10.0.2.10	255.255.255.0	10.0.2.1	Automatic
Node 1 HA	10.0.2.7	255.255.255.0	10.0.2.1	Automatic
Node 2 HA	10.0.2.9	255.255.255.0	10.0.2.1	Automatic
Node 1 LAN1	10.0.2.6	255.255.255.0	10.0.2.1	Automatic
Node 2 LAN1	10.0.2.8	255.255.255.0	10.0.2.1	Automatic

Site 1 – HA Member

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
 - **Host Name**: Enter **ns3.site1.corp100.com**
 - **Comment**: Enter **Site 1 - ns3.site1.corp100.com**
4. Specify the following information about the member that you are adding to the Grid and click **Save & Close**:
 - **High Availability Pair**: Select this option.
 - **Virtual Router ID**: Enter **111**.
 - **Required Ports and Addresses**:

Interface	Address	Subnet Mask	Gateway	Port Settings
VIP	10.1.1.10	255.255.255.0	10.1.1.1	Automatic
Node 1 HA	10.1.1.7	255.255.255.0	10.1.1.1	Automatic
Node 2 HA	10.1.1.9	255.255.255.0	10.1.1.1	Automatic
Node 1 LAN1	10.1.1.6	255.255.255.0	10.1.1.1	Automatic
Node 2 LAN1	10.1.1.8	255.255.255.0	10.1.1.1	Automatic

Site 2 – Single Member

1. From the **Grid** tab, select the **Grid Manager -> Members** tab.
2. Expand the Toolbar and click **Add -> Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
 - **Host Name:** ns4.site2.corp100.com
 - **Comment:** Site 2- ns4.site2.corp100.com
4. Specify the following information about the member that you are adding to the Grid and click **Next**:
 - **Standalone Member:** Select this option.
 - **Address:** Enter **10.2.1.10**.
 - **Subnet Mask:** Enter **255.255.255.0**.
 - **Gateway:** Enter **10.2.1.1**.
 - **Port Settings:** Select **AUTOMATIC**.
5. Save the configuration and click **Restart** if it appears at the top of the screen.
6. Log out from the Grid Master.

Join Appliances to the Grid

To complete the process of adding appliances to the Grid, log in to and configure each individual appliance so that it can contact the Grid Master.

HQ Site – HA Grid Member (Node 1)

Make a console connection to the appliance that you want to make Node 1 in the HA pair, and enter the following:

```
Infoblox > set network
NOTICE: All HA configuration is performed from the GUI. This interface is used only
to configure a standalone node or to join a Grid.
Enter IP address: 10.0.2.6
Enter netmask [Default: 255.255.255.0]:
Enter gateway address [Default: 10.0.2.1]:
Become Grid member? (y or n): y
Enter Grid Master VIP: 10.0.1.10
Enter Grid Name: corp100
Enter Grid Shared Secret: MglkW17d
New Network Settings:
IP address: 10.0.2.6
Netmask: 255.255.255.0
Gateway address: 10.0.2.1
Join Grid as member with attributes:
  Grid Master VIP: 10.0.1.10
  Grid Name: corp100
  Grid Shared Secret: MglkW17d
WARNING: Joining a Grid will replace all the data on this node!
  Is this correct? (y or n): y
  Are you sure? (y or n): y
```

The Infoblox application restarts. After restarting, the appliance contacts the Grid Master and joins the Grid as Node 1.

HQ Site – HA Member (Node 2)

Make a console connection to the appliance that you want to make Node 2 in the HA pair, and enter exactly the same data you entered for Node 1 except that the IP address is 10.0.2.8.

After the application restarts, the appliance contacts the Grid Master and joins the Grid as Node 2, completing the HA member configuration for the HQ site.

Site 1 – HA Grid Member (Node 1)

Make a console connection to the appliance that you want to make Node 1 in the HA pair at Site 1, and use the `set network` command to configure its basic network and Grid settings. Use the following data:

- IP Address: 10.1.1.6
- Netmask: 255.255.255.0
- Gateway: 10.1.1.1
- Grid Master VIP: 10.0.1.10
- Grid Name: corp100
- Grid shared secret: Mg1kW17d

The Infoblox application restarts. After restarting, the appliance contacts the Grid Master and joins the Grid as Node 1.

Site 1 – HA Grid Member (Node 2)

Make a console connection to the appliance that you want to make Node 2 in the HA pair at Site 1, and enter exactly the same data you entered for Node 1 except that the IP address is 10.1.1.8.

After the application restarts, the appliance contacts the Grid Master and joins the Grid as Node 2, completing the HA member configuration for Site 1.

Site 2– Single Grid Member

Make a console connection to the appliance that you want to make Node 1 in the HA pair at Site 1, and use the `set network` command to configure its basic network and Grid settings. Use the following data:

- IP Address: 10.2.1.10
- Netmask: 255.255.255.0
- Gateway: 10.2.1.1
- Grid Master VIP: 10.0.1.10
- Grid name: corp100
- Grid shared secret: Mg1kW17d

The Infoblox application restarts. After restarting, the appliance contacts the Grid Master and joins the Grid.

To check the status of all the Grid members, log in to the Grid Master at 10.0.1.10, and from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, select 10.0.1.10 and click the Detailed Status icon. Check that the status indicators are all green in the Detailed Status panel. As an appliance joins a Grid, it passes through the following phases: Offline, Connecting, (Downloading Release from Master), Synchronizing, and Running.)

Note: Depending on the network connection speed and the amount of data that the master needs to synchronize with the member, the process of joining a Grid can take from several seconds to several minutes to complete.

The Grid setup is complete.

Import DHCP Data

The Data Import Wizard is a software tool that you can download from the Infoblox Support site to your management system. With it, you can import data from legacy DHCP and DNS servers to NIOS appliances. In this example, you use it to import both DHCP and DNS data to the Grid Master at 10.0.1.10, which then uses the database replication mechanism to send the imported data to other Grid members. In the wizard, you also specify which Grid members serve the imported data. The wizard supports various types of DHCP formats, such as the following:

- ISC DHCP
- Lucent VitalQIP
- Microsoft
- Nortel NetID
- CSV (comma-separated values); you can also import IPAM data in CSV format

In this example, all the DHCP data is in standard ISC DHCP format.

Importing DHCP Data for HQ and Site 2

1. Save the DHCP configuration file from your legacy DHCP server at 10.0.2.20 to a local directory.
2. Visit www.infoblox.com/support, log in with your support account, and download the Data Import Wizard. The Data Import Wizard application downloads to a container within a Java sandbox on your management system and immediately launches, displaying the Welcome page.
3. After reading the information in the left panel, click Next.
4. Select Import to Infoblox Appliance, enter the following, and then click Next:
 - Hostname or IP address: 10.0.1.10
 - Username: admin
 - Password: 1n85w2IF
5. Select the following, and then click Next:
 - What kind of data would you like to import? DHCP/IPAM
 - Which legacy system are you importing from? ISC DHCP
 - Which appliance will be serving this data? 10.0.2.10
6. Type the path and file name of the DHCP configuration file saved from the legacy server, and then click Next.
or
Click Browse, navigate to the file, select it, click Open, and then click Next.
7. In the Global DHCP Configuration table, double-click the Value cell for the domain-name-servers row, and change the IP addresses to 10.0.2.10.
8. When satisfied with the data, click Import.
You can view the status of the importation process and a summary report in the Data Import Wizard Log.
9. To enable DDNS updates, log in to the Grid Master, from the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Configuration**.
10. In the **DDNS** -> **Basic** tab of the Grid DHCP Properties editor, select **Enable DDNS Updates**.
11. Save the configuration and click **Restart** if it appears at the top of the screen.
12. To check the imported DHCP configuration file, from the **Data Management** tab, select the **DHCP** tab, -> **Members** tab -> 10.0.2.10 -check box. Expand the Toolbar and click View DHCP Configuration.
13. In the DHCP configuration file, check that all the imported subnets are present, and navigate to the beginning of the file and check that you see the **ddns-updates on** statement. (If you see **ddns-updates off**, enable DDNS updates for the Grid as explained in steps 9–12.)

Importing DHCP Data for Site 1

1. Repeat the steps in [Importing DHCP Data for HQ and Site 2](#), saving the DHCP configuration file from your legacy DHCP server at 10.1.1.20, and importing it to the Grid Master at 10.0.1.10 for the member with IP address 10.1.1.10 to serve.
2. Check the imported DHCP configuration file by logging in to the Grid Master and from the **Data Management** tab, select the **DHCP** tab -> Members tab -> 10.1.1.10 -check box. Expand the Toolbar and click View DHCP Configuration.

Importing DHCP Data for Site 3

1. Repeat the steps in [Importing DHCP Data for HQ and Site 2](#), saving the DHCP configuration file from your legacy DHCP server at 10.1.1.20, and importing it to the Grid Master at 10.0.1.10 for the member with IP address 10.3.1.10 to serve.
2. After the importation process completes, check the imported DHCP configuration file by logging in to the Grid Master and from the **Data Management** tab, select the **DHCP** tab -> Members tab -> 10.3.1.10 -check box. Expand the Toolbar and click View DHCP Configuration.

Import DNS Data

Using the Infoblox Data Import Wizard, import DNS data from the legacy hidden primary server at 10.0.1.5 to the new hidden primary server at 10.0.1.10 (the Grid Master). There are three phases to this task:

- [Before Using the Wizard](#):
 - Save the named.conf file from the legacy server to a file in a local directory on your management system.
 - Enable the legacy server to perform zone transfers to the NIOS appliance.
 - Configure three name server groups for the Grid, and allow the Grid Master/hidden primary DNS server at 10.0.1.10 to receive DDNS updates from the Grid members at 10.0.2.10, 10.1.1.10, and 10.3.1.10. These members act as secondary DNS servers and DHCP servers.
- [Using the Wizard](#) on page 264: Define the source, destination, and type of DNS data in the DNS configuration file (named.conf) that you want to import.
- [After Using the Wizard](#) on page 266: Check the imported DNS configuration file.

In this example, all the DNS data is in BIND 9 format. The Data Import Wizard supports various types of DNS formats, such as the following:

- BIND 4, 8, and 9
- Microsoft
- Lucent VitalQIP
- Nortel NetID

Before Using the Wizard

You must set up the legacy server and Grid Master before using the Data Import Wizard.

Legacy Server

1. Log in to the legacy name server at 10.0.1.5 and save the named.conf file, which contains all the DNS settings that you want to import into the Infoblox name server, to a local directory on your management system.
2. On the legacy server, enable zone transfers to the NIOS appliance.

Infoblox Grid Master – DDNS Updates

1. Log in to the Grid Master at 10.0.1.10, and from the **Data Management** tab, select the **DNS** tab -> **Members** tab -> 10.0.1.10 check box and select the Edit icon.
2. In the Member DNS Configuration editor, select the Updates -> **Basic** tab and enter the following:
 - Select **Override**.
 - Allow updates from: Click the Add icon and select IPv4 Address. Enter 10.0.2.10 in the **Name** field of the new row.
3. Click the Add icon again and add 10.1.1.10 and 10.2.1.10 as IP addresses from which you allow DDNS updates.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: When all DNS servers are members in the same Grid, the members use database replication to synchronize all their data—including DNS zone data. You can change the default behavior so that Grid members use zone transfers instead. In this example, Grid members use database replication.

Infoblox Grid Master – Name Server Groups

1. From the **Data Management** tab, select the **DNS** tab -> **Name Server Groups** tab.
2. Click the Add icon to open the Add Name Server Group wizard.
3. Enter the following:
 - **Name Server Group Name:** HQ-Group
4. Click the Add icon and add the following:
 - **Grid Primary:** ns1.corp100.com; Stealth: Select this check box.
 - **Grid Secondary:** ns2.corp100.com; Grid replication (recommended): Select this check box.
5. Click Save & New.
6. Repeat steps 2 to 4 to create another group. Name it Site1-Group, and use ns1.corp100.com as the hidden primary server, ns3.site1.corp100.com as a secondary server, and Grid replication for zone updates.
7. Repeat steps 2 to 4 to create another group. Name it Site2-Group, and use ns1.corp100.com as the hidden primary server, ns4.site2.corp100.com as a secondary server, and Grid replication for zone updates.

Using the Wizard

While progressing through the Data Import Wizard, you must define the source, destination, and type of DNS data that you want to import. You then make some simple modifications to the data and import it.

Defining the Source, Destination, and Type of DNS Data

1. Launch the Data Import Wizard.
2. After reading the information in the left panel of the welcome page, click Next.
3. Select Import to Infoblox Appliance, enter the following, and then click Next:
 - Hostname or IP address: 10.0.1.10
 - Username: admin
 - Password: 1n85w2IF

The Data Import Wizard Log opens in a separate window behind the wizard. Leave it open while you continue.

4. Select the following, and then click Next:
 - What kind of data would you like to import? DNS
 - Which legacy system are you importing from? BIND 9
 - Which appliance will be serving this data? 10.0.1.10

5. Select the following, and then click Next:
 - What BIND 9 DNS configuration file would you like to use? Click Browse, navigate to the named.conf file you saved from the legacy server, select it, and then click Open.
 - What type of BIND 9 DNS data do you want to import? DNS zone information and DNS record data
 - Where is the BIND 9 DNS record data? Zone transfer(s) from a DNS server; 10.0.1.5

The wizard displays two tables of data. The upper table contains global DNS server configuration parameters. The lower table contains zone configurations.

The Data Import Wizard Log presents a summary listing the number of views, zones, and DNS records in the configuration file.

Modifying DNS Data

While importing data from the legacy DNS server, you cancel the importation of global configuration settings, and apply the name server groups you created in [Before Using the Wizard](#) on page 263 to the zones you want to import.

1. In the Global DNS Configuration table, select all rows by clicking the top row and then SHIFT+clicking the bottom row.
2. Right-click the selected rows to display the Set Import Options dialog box, select Do not import, and then click Apply.
3. In the DNS Zones table, clear the Import check box for the default view.
4. Select corp100.com, lab.corp100.com and all the corresponding reverse-mapping zones.

Tip: You can use SHIFT+click to select multiple contiguous rows and CTRL+click to select multiple noncontiguous rows.

5. Right-click the selected rows, and then select Set Import Options.
6. In the Set Import Options dialog box, enter the following, and then click Apply:
 - Set Zone Type: No change
 - Set Import Option: No change
 - Set View: default
 - Set Member: HQ-Group master
7. Select site1.corp100.com and all the reverse-mapping zones with 1 in the second octet in the zone name (1.1.10.in-addr.arpa, 2.1.10.in-addr.arpa, 3.1.10.in-addr.arpa, and so on).
8. Right-click the selected rows, and select Set Import Options.
9. In the Set Import Options dialog box, make the same selections as in [Step 6](#), but choose Site1-Group master from the Set Member drop-down list.
10. Similarly, select site2.corp100.com and all the reverse-mapping zones with 2 in the second octet in the zone name.
11. Right-click the selected rows, and select Set Import Options.
12. In the Set Import Options dialog box, make the same selections as in [Step 6](#), but choose Site2-Group master from the Set Member drop-down list.

Importing DNS Data

1. Click Import.
The wizard imports the global DNS parameters and zone-specific configuration settings from the named.conf file and performs a zone transfer of the data from the legacy server.
2. Use the Data Import Wizard Log to monitor progress and review results afterward.
The log lists all the zones that the wizard imports and concludes with a total of all the successfully and unsuccessfully imported zones.

Note: If the wizard is unable to import a zone, an error message with an explanation appears in the log.

3. To close the Data Import Wizard, click Exit. This closes the Data Import Wizard Log as well.

After Using the Wizard

After you import data, you must restart services on the Grid Master and delete the A records for the legacy servers from the corp100.com zone. You can also confirm that the imported data is correct and complete by checking the DNS configuration and the forward- and reverse-mapping zones.

1. Log in to the Grid Master (10.0.1.10), select the **Grid** tab, expand the Toolbar, and then click the Restart Services icon.

Note: When importing data through the wizard rather than entering it through the GUI, the Restart Services icon does not change to indicate you must restart service for the appliance to apply the new data. Still, restarting service on the Grid Master is necessary for the imported configuration and data to take effect.

2. To remove A records for the legacy servers, from the **Data Management** tab, select **DNS** tab -> **Zones** tab -> corp100.com.
3. Expand the Records section, select the following A records in the corp100.com zone, and then click the Delete icon:
 - ns1 (for 10.0.1.5)
 - ns2 (for 10.0.2.5)
 - ns3.site1.corp100 (for 10.1.1.5)
 - ns4.site3.corp100 (for 10.2.1.5)
4. Remove the respective A records for legacy servers from the site1.corp100 and site3.corp100 subzones.
5. To check the imported DNS configuration file, from the **Data Management** tab, select **DNS** tab -> **Members** tab -> 10.0.1.10 check box. Expand the Toolbar and click View -> View DNS Configuration.

Note: If you do not see the imported DNS configuration file, make sure you enabled DNS and restarted services.

6. Scroll through the DNS configuration log to check that each imported zone has an allow-update statement like the following one for the 10.1.10.in-addr.arpa reverse-mapping zone:

```
zone "10.1.10.in-addr.arpa" in {
    ...
    allow-update { key DHCP_UPDATER; 10.0.2.10; 10.1.1.10; 10.2.1.10; };
    ...
};
```

Enable DHCP and Switch Service to the Grid

Finally, you must enable DHCP service on the three Grid members at 10.0.2.10, 10.1.1.10, and 10.2.1.10, and switch DNS and DHCP service from the legacy DNS and DHCP servers to them.

1. Log in to the Grid Master (10.0.1.10) and from the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> 10.0.2.10 check box. Expand the Toolbar and click **Start**.
2. Repeat step 1 to enable DHCP on 10.1.1.10 and 10.3.1.10.

Note: Start the DNS service, as described in [Starting and Stopping the DNS Service](#) on page 565.

The Grid members are ready to serve DHCP and DNS, and send DDNS updates.

3. Take the legacy DHCP and DNS servers offline.

Managing a Grid

After you configure a Grid Master and add members, you might need to perform the following tasks:

- [Changing Grid Properties](#) on page 267
- [Configuring Security Level Banner](#) on page 268
- [Configuring Notice and Consent Banner](#) on page 268
- [Configuring Informational Level Banner](#) on page 269
- [Setting the MTU for VPN Tunnels](#) on page 270
- [Removing a Grid Member](#) on page 270
- [Promoting a Master Candidate](#) on page 270

Changing Grid Properties

You can change a Grid name, its shared secret, and the port number of the VPN tunnels that the Grid uses for communications. Note that changing the VPN port number, time zone, date or time requires a product restart.

To modify the properties of a Grid:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> click the **Basic** tab, and then modify any of the following:
 - **Grid Name:** Type the name of a Grid. The default name is *Infoblox*.
 - **Shared Secret:** Type a shared secret that all Grid members use to authenticate themselves when joining the Grid. The default shared secret is *test*.
 - **Shared Secret Retype:** Type the shared secret again to confirm its accuracy.
 - **Time Zone:** Choose the applicable time zone from the drop-down list.
 - **Date:** Click the calendar icon to select a date or enter the date in YYYY/MM/DD format.
 - **Time:** Click the clock icon to select a time or enter the time in HH:MM:SS format. For afternoon and evening hours, use the integers 13-24.
 - **VPN Port:** Type the port number that the Grid members use when communicating with the Grid Master through encrypted VPN tunnels. The default port number is 1194. For more information, see [Port Numbers for Grid Communication](#) on page 238.
 - **Enable Recycle Bin:** Select the check box to enable the Recycle Bin. The Recycle Bin stores deleted items when the user deletes Grid, DNS, or DHCP configuration items. Enabling the Recycle Bin allows you to undo deletions and to restore the items on the appliance at a later time. If you do not enable this feature, deleted items from the GUI are permanently removed from the database.

- **Audit Logging:** Select one of the following:
 - **Detailed:** This is the default type. It is automatically selected. It provides detailed information on all administrative changes such as the date and time stamp of the change, administrator name, changed object name, and the new values of all properties.
 - **Brief:** Provides information on administrative changes such as the date and time stamp of the change, administrator name, and the changed object name. It does not show the new value of the object.
- In the *Grid Properties* editor, select the **General** tab -> click the **Advanced** tab (or click Toggle Advanced Mode) and select the **Enable GUI Redirect from Member** check box to allow the appliance to redirect the Infoblox GUI from a Grid member to the Grid Master.

4. Save the configuration.

If you changed the VPN port number, time zone, date or time, Grid Manager displays a warning indicating that a product restart is required. Click **Yes** to continue, and then log back in to Grid Manager after the application restarts.

Configuring Security Level Banner

You can publish a security banner that indicates the security level of the Infoblox Grid. It appears on the header and footer of all pages of Grid Manager. The security level can be Top Secret, Secret, Confidential, Restricted, and Unclassified. Each message type is associated with a predefined security level color. You can modify this color at any point of time. Grid Manager automatically uses an appropriate contrasting text font color that goes with the banner color. Only superusers can configure and enable this feature.

To configure the advanced security level banner for a Grid:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **Security** tab -> **Advanced** tab.
4. Complete the following:
 - **Enable Security Banner:** Select this to enable the display of the security banner.
 - **Security Level:** From the drop-down list, select the security level for the banner.
 - **Security Level Color:** The default color is displayed in the drop-down. If necessary using the drop-down list, select the required color for the security level banner. When you change the security level, Grid Manager resets default color for that level.
 - **Classification Message:** Enter the message you want to display in the security banner. You can enter up to 190 characters.
5. Save the configuration.

Security banner appears on the header and footer of the Grid Manager screen including the Login screen.

Configuring Notice and Consent Banner

You can configure and publish a notice and consent banner as the first login screen that includes specific terms and conditions you want end users to accept before they log in to the Infoblox Grid. When an end user tries to access Grid Manager, this banner is displayed as the first screen. The user must accept the terms and conditions displayed on the consent screen before accessing the login screen of Grid Manager. Only superusers can configure and enable this feature.

To configure the notice and consent banner:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the **Toolbar** and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **Security** tab -> **Advanced** tab, and then complete the following:
 - **Enable Notice and Consent Banner:** Select the check box to enable the display of the notice and consent banner. In the text field, enter the message that you want to be included in the banner. The message cannot exceed 10,000 characters.

4. Save the configuration.

This banner appears as the first screen when users access Grid Manager. Users must read the terms and conditions and then click **Accept** on the consent screen before they can access the login screen of Grid Manager.

Configuring Informational Level Banner

You can publish the informational banner for multiple uses, such as to indicate whether the Infoblox Grid is in production or a lab system. The banner can also be used for issuing messages of the day. The informational level banner appears on the header of the Grid Manager screen. You can publish the banner information you want and set the banner color. Grid Manager automatically uses an appropriate contrasting text font color that goes with the banner color. Only superusers can configure and enable this feature.

To configure the advanced informational banner for a Grid:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab
4. Complete the following:
 - **Enable informational GUI Banner:** Select the check box to enable the display of the informational banner message.
 - **Banner Color:** The default color is displayed in the drop-down. If necessary using the drop-down list, select the required color for the informational level banner.
 - **Message:** Enter the message you want to display in the informational banner. You can enter up to 190 characters.
5. Save the configuration.

Informational banner appears on the header of the Grid Manager screen.

Configuring Recursive Deletions of Networks and Zones

Through Grid Manager, you can configure the group of users that are allowed to delete or schedule the deletion of a network container and its child objects as well as a zone and its child objects. For information about how to delete a network container or zone, see [Deleting Network Containers](#) on page 465 and [Removing Zones](#) on page 634.

When you select **All Users** or **Superusers**, these users can choose to delete a parent object and re-parent its child objects, or they can choose to delete a parent object and all its child objects. These options appear only if a network container or a zone has child objects. For information about scheduling recursive deletion of network containers and zones, see [Scheduling Recursive Deletions of Network Containers and Zones](#) on page 76.

When you select **Nobody**, all the users can delete the parent object only. All the child objects, if any, are re-parented. For information about scheduling deletions, see [Scheduling Deletions](#) on page 76. Note that you can restrict specific users to perform recursive deletions of network containers and zones only through Grid Manager. These settings do not prevent other users from performing recursive deletions through the API.

Note: You must have Read/Write permission to all the child objects in order to delete a parent object. Recursive deletion is applicable to all zone types except stub and forward-mapping zones.

The appliance puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary. When you restore a parent object from the Recycle Bin, all its contents, if any, are re-parented to the restored parent object. For information about the Recycle Bin, see [Using the Recycle Bin](#) on page 64.

To configure the group of users to perform recursive deletions:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab.

4. Under **Present the option of recursive deletion of networks or zones to**, select one of the following:
 - **All Users:** Select this to allow all users, including superusers and limited-access users, to choose whether they want to delete the parent object and its contents or the parent object only when they delete a network container/network or a zone. This is selected by default.
 - **Superuser:** Select this to allow only superusers to choose whether they want to delete the parent object and its contents or the parent object only when they delete a network container/network or a zone.
 - **Nobody:** When you select this, users can only delete the parent object (network container or zone). All child objects, if any, are re-parented.
5. Save the configuration.

Setting the MTU for VPN Tunnels

You can configure the VPN MTU (maximum transmission unit) for any appliance with a network link that does not support the default MTU size (1500 bytes) and that cannot join a Grid because of this limitation. If an appliance on such a link attempts to establish a VPN tunnel with a Grid Master to join a Grid, the appliance receives a PATH-MTU error, indicating that the path MTU discovery process has failed. For information about the MTU discovery process, see *RFC 1191, Path MTU Discovery*.

To avoid this problem, you can set a VPN MTU value on the Grid Master for any appliance that cannot link to it using a 1500-byte MTU. When the appliance contacts the master during the key exchange handshake that occurs during the Grid-joining operation, the master sends the appliance the MTU setting to use.

To set the VPN MTU for a Grid member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box -> Edit icon.
2. Select the **Network** -> **Advanced** tab of the *Grid Member Properties* editor.
3. In the **VPN MTU** field, enter a value between 600 and 1500.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Removing a Grid Member

You might want or need to remove a member from a Grid, perhaps to disable it or to make it an independent appliance or an independent HA pair. Before you remove a member, make sure that it is not assigned to serve any zones or networks.

To remove a Grid member, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and click the Delete icon.

Promoting a Master Candidate

To promote a master candidate to a Grid Master, you must have previously designated a Grid member as a master candidate. Select the **Master Candidate** option in the **General** tab of the *Grid Member Properties* editor to designate the member as a master candidate. Before promoting a master candidate, check your firewall rules to ensure that the master candidate can communicate with all the Grid members. For information, see [Grid Communications](#) on page 225.

To promote a master candidate, you can make a direct serial connection to the console port on the active node of an HA candidate or to the console port on a single candidate. You can also make a remote serial connection (using SSH v2) to the candidate. Enter the following Infoblox CLI command to promote a master candidate:

```
set promote_master.
```

You can do one of the following to promote a master candidate:

- Immediately notify all Grid members about the promotion.
- Set a sequential notification to provide wait time for Grid members to join the new Grid Master. Staggering the restarts of Grid members can minimize DNS outages. The sequential order for Grid members to join the new Grid Master begins with the old Grid Master and then the Grid members in FQDN order. The default delay time is 120 seconds. You can configure the delay time from a minimum of 30 seconds up to 600 seconds.

Note: During a Grid Master promotion, ensure that you do not designate a Grid member as a Grid Master candidate or promote a master candidate.

To promote a master candidate, do the following:

1. Establish a serial connection (through a serial console or remote access using SSH) to the master candidate. For information about making a serial connection, see [Method 2 – Using the CLI](#) on page 277.
2. At the CLI prompt, use the command `set promote_master` to promote the master candidate and send notifications to all Grid members immediately, or promote the master candidate to the Grid Master immediately and specify the delay time for the Grid members to join the new Grid Master. For more information about the command, refer to the *Infoblox CLI Guide*.
3. To verify the new master is operating properly, log in to the Infoblox Grid Manager on the new master using the VIP address for an HA master or the IP address of the LAN1 port for a single master.
4. Check the icons in the **Status** column. Also, select the master, and then click the Detailed Status icon in the table toolbar. You can also check the status icons of the Grid members to verify that all Grid members have connected to the new master. If you have configured delay time for Grid member notification, it will take some time for some members to connect to the new master. You can also check your firewall rules and log in to the CLI to investigate those members.

Note: Note that when you promote the master candidate to a Grid Master, the IP address will change accordingly. If you have configured a FireEye appliance, then any changes in the Grid Master IP address, FireEye zone name, associated network view or the DNS view will affect the **Server URL** that is generated for a FireEye appliance. The FireEye appliance will not be able to send alerts to the updated URL when there is a change in the IP address. You must update the URL in the FireEye appliance to send alerts to the NIOS appliance. For more information, see [Configuring FireEye RPZs](#) on page 1254.

About the Master Grid

A Master Grid provides centralized management of multiple Grids. When a Grid is managed by a Master Grid, the Master Grid icon appears on the left side of the top panel of Multi-Grid Manager. Assuming you have permission, you can click this icon to access Multi-Grid Manager. In addition, the Toolbar provides several functions for joining the Master Grid, editing its properties and leaving the Master Grid. For more information about the Master Grid and these functions, refer to the *Multi-Grid Manager Administrator Guide*.



Chapter 6 Deploying Independent Appliances

This chapter explains how to deploy single independent appliances and independent HA pairs. Independent appliances run NIOS without the Grid upgrade and are deployed independently from a Grid. This chapter includes the following sections:

- [Independent Deployment Overview](#) on page 275
 - [System Manager GUI](#) on page 276
- [Deploying a Single Independent Appliance](#) on page 276
 - [Method 1 – Using the LCD](#) on page 277
 - [Method 2 – Using the CLI](#) on page 277
 - [Method 3 – Using the Infoblox NIOS Startup Wizard](#) on page 279
- [Configuration Example: Deploying a NIOS Appliance as a Primary DNS Server](#) on page 281
 - [Cabling the Appliance to the Network and Turning On Power](#) on page 282
 - [Specifying Initial Network Settings](#) on page 282
 - [Specifying Appliance Settings](#) on page 282
 - [Enabling Zone Transfers on the Legacy Name Server](#) on page 283
 - [Importing Zone Data on an Independent Appliance](#) on page 284
 - [Designating the New Primary on the Secondary Name Server \(at the ISP Site\)](#) on page 286
 - [Configuring NAT and Policies on the Firewall](#) on page 287
- [Deploying an Independent HA Pair](#) on page 287
 - [Using the Infoblox NIOS Startup Wizard to Configure an HA Pair](#) on page 287
- [Configuration Example: Configuring an HA Pair for Internal DNS and DHCP Services](#) on page 290
 - [Cabling Appliances to the Network and Turning On Power](#) on page 291
 - [Specifying Initial Network Settings](#) on page 292
 - [Specifying Appliance Settings](#) on page 292
 - [Enabling Zone Transfers](#) on page 294
 - [Importing Zone Data](#) on page 294
 - [Defining Networks, Reverse-Mapping Zones, DHCP Ranges, and Infoblox Hosts](#) on page 294
 - [Defining Multiple Forwarders](#) on page 297
 - [Enabling Recursion on External DNS Servers](#) on page 297
 - [Modifying the Firewall and Router Configurations](#) on page 298

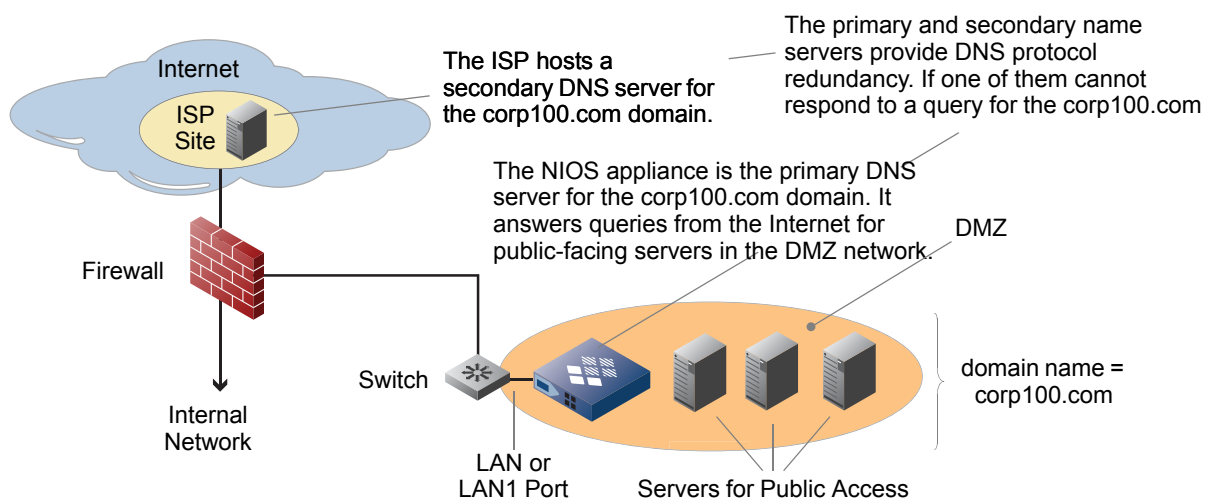
- [*Enabling DHCP and Switching Service to the NIOS Appliance*](#) on page 299
 - [*Managing and Monitoring*](#) on page 299
- [*Verifying the Deployment*](#) on page 300
 - [*Single Independent Appliance*](#) on page 300
 - [*Independent HA Pair*](#) on page 300
- [*Infoblox Tools for Migrating Bulk Data*](#) on page 301

INDEPENDENT DEPLOYMENT OVERVIEW

Note: Infoblox appliances support IPv4 and IPv6 networking configurations in most deployments cited in this chapter. You can set the LAN1 port to an IPv6 address and use that address to access the NIOS UI and the NIOS Setup Wizard. All HA operations can be applied across IPv6. Topics in this and following chapters generally use IPv4 examples. Also note that LAN2 and the MGMT port also support IPv6. DNS and DHCP services are fully supported in IPv6 for the LAN2 port. Example networks throughout this chapter use IPv4 addressing.

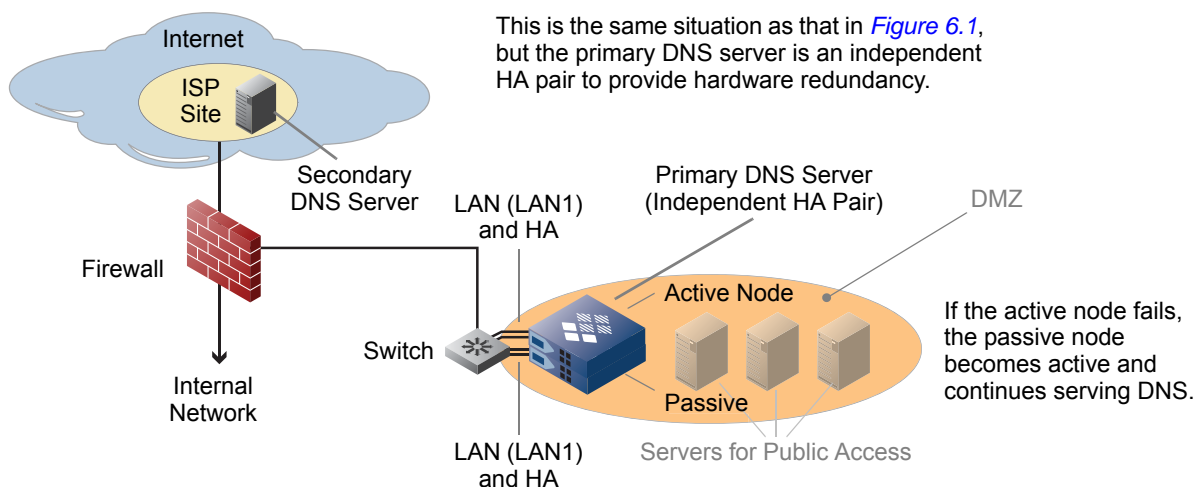
You can deploy the NIOS appliance as a Grid member in an Infoblox Grid or independently as a standalone deployment. Grids offer many advantages for large organizations while independent deployments can be sufficient for smaller sites. For example, if your ISP hosts one name server to respond to external DNS queries, you can deploy a single independent NIOS appliance as the other name server, as shown in [Figure 6.1](#).

Figure 6.1 Single Independent Appliance as a DNS Server



Using primary and secondary name servers provides DNS protocol redundancy, and configuring two DHCP servers as DHCP failover peers provides DHCP protocol redundancy. However, you can only have hardware redundancy if you deploy appliances in an HA (high availability) pair. Should the active node in an HA pair fail, the passive node becomes active and begins serving data, as shown in [Figure 6.2](#). For more information about HA pairs, see [About HA Pairs](#) on page 233.

Figure 6.2 Independent HA Pair



System Manager GUI

When you deploy an independent appliance, you use System Manager to manage the appliance. Though other chapters in this guide contain information that assumes a Grid deployment and describes the Grid Manager GUI, most of the configuration procedures are applicable to an independent appliance, with the following differences:

- In the Dashboard, there is no *Grid Status* widget, and the *Members Status* widget in Grid Manager is the *System Status* widget in System Manager.
- Functions related to a Grid, such as joining a Grid and managing Grid licenses, do not exist in System Manager.
- The Grid related tabs and functions in Grid Manager are the system related tabs and functions in System Manager.
- Functions related to the **Members** tab in Grid Manager appear in the **Nodes** tab or the Toolbar of another subtab in System Manager.

For example, the following navigation path for a Grid:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click **HTTPS Cert** -> **Download Certificate** from the Toolbar.

is the following for an independent appliance:

- From the **System** tab, select the **System Manager** tab -> **Nodes** tab, and then click **HTTPS Cert** -> **Download Certificate** from the Toolbar.

DEPLOYING A SINGLE INDEPENDENT APPLIANCE

To deploy a single independent NIOS appliance, you cable its LAN1 port to the network and change its default IP settings so that it can connect to its surrounding IP address space. The default LAN settings are as follows:

- IP address: 192.168.1.2
- Netmask: 255.255.255.0
- Gateway: 192.168.1.1

When deploying a single independent appliance, you can use one of the following methods to set up the initial configuration:

- *Method 1 – Using the LCD*
 - Requirements: Physical access to a powered up NIOS appliance.
 - Advantage: You do not need any other equipment.
- *Method 2 – Using the CLI*
 - Requirements: A serial connection from your management system to the console port on the NIOS appliance. You can also enable remote console access so that you can use the CLI over a network connection. For information, see [Enabling Remote Console Access](#) on page 343.
 - Advantage: You do not need to change the IP address of the management system to connect to the NIOS appliance.
- *Method 3 – Using the Infoblox NIOS Startup Wizard*
 - Requirements: An HTTPS connection from your management system to the LAN1 port on the NIOS appliance.
 - Advantage: The wizard provides step-by-step guidance for changing not only the IP settings for the LAN1 port, but also changing the appliance host name and admin password, setting the system clock, and—if using NTP (Network Time Protocol)—enabling the NIOS appliance to be an NTP client.


Note that you can configure network settings using the Startup wizard any time after you have configured the appliance. To start the wizard, from System Manager, select the **System** tab, and then click **System Properties** -> **Startup Wizard** from the Toolbar.

After you configure the network settings on a single independent appliance, you can migrate data from legacy DNS and DHCP servers to the NIOS appliance. Several tools and methods are available for migrating data and configuration settings. For a list of the available options, see [Infoblox Tools for Migrating Bulk Data](#) on page 301.

Method 1 – Using the LCD

Some of the NIOS appliances have an LCD and navigation buttons on the front panel that allow you to view system status and license information as well as configure network settings for the LAN1 port.

Figure 6.3 Infoblox LCD and Navigation Buttons



You can deploy a single independent NIOS appliance by setting its LAN1 port IP address, netmask, and gateway through the LCD. This is the simplest method because you do not need anything other than a physical access to the appliance to complete the initial configuration.


1. Connect the power cable from the NIOS appliance to a power source and turn on the power.
At startup, the Infoblox logo appears in the LCD on the front panel of the appliance. Then the LCD scrolls repeatedly through a series of display screens.
2. To change the network settings for the LAN1 port, press one of the navigation buttons.
The LCD immediately goes into the input mode, in which you can enter the IP address, netmask, and gateway for the LAN1 port.
3. Use the navigation buttons to enter an IP address, netmask, and gateway address for the LAN1 port.
4. Cable the LAN1 port of the NIOS appliance to a network as described in the installation guide that shipped with your product.

Method 2 – Using the CLI

You can use the Infoblox CLI to make an initial network configuration through the `set network` command. To access the CLI, make a direct serial connection from your management system.

1. Connect a console cable from the console port on your workstation to the male DB-9 console port on the NIOS appliance.
The DB-9 pin assignments follow the EIA232 standard. You can use one RJ-45 rollover cable and two female RJ-45-to-female DB-9 adapters, or a female DB-9-to-female DB-9 null modem cable.

Figure 6.4 Console Connection



2. Use a serial terminal emulation program, such as Hilgraeve Hyperterminal[®] (provided with Windows[®] operating systems), to launch a session. The connection settings are:
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: Xon/Xoff
3. Log in to the appliance using the default username and password (admin and infoblox).
4. At the Infoblox command prompt, enter **set network** to change the network settings, such as the IP address, netmask, and gateway for the LAN1 port.

Note: In the following example, the variable *ip_addr1* is the IP address of the LAN1 port and *ip_addr2* is the IP address of the gateway for the subnet on which you set the *ip_addr1* address. Infoblox appliances support IPv4 and IPv6 networking configurations in all deployments cited in this chapter. You can set the LAN1 port to an IPv6 address and use that address to access the NIOS UI. IPv6 transport is not allowed unless an IPv4 configuration is defined first.

```
Infoblox > set network
```

```
NOTICE: All HA configuration is performed from the GUI. This interface is used only to configure a standalone node or to join a Grid.
```

```
Enter IP address: ip_addr1
```

```
Enter netmask: [Default: 255.255.255.0]: netmask
```

```
Enter gateway address [Default: n.n.n.1]: ip_addr2
```

```
Configure IPv6 network settings? (y or n): y
```

```
Enter IPv6 address []: 2001:db8:a22:a00::29
```

```
Enter IPv6 cidr []: 64
```

```
Enter IPv6 gateway [Default: ::1] 2001:db8:a22:a00::1
```

```
Become Grid member? (y or n): n
```

You can press N to avoid configuring IPv6 on the command line. After you confirm your network settings, the Infoblox application automatically restarts.

5. Cable the LAN1 port to a network. For information about installing and cabling the appliance, refer to the user guide or installation guide that was shipped with the product.

Method 3 – Using the Infoblox NIOS Startup Wizard

When you first make an HTTPS connection to a NIOS appliance, the Infoblox NIOS Startup Wizard guides you through the deployment options and basic network settings. You can also change the password of the superuser (admin) and set up the system clock.

Note that you can configure network settings using the Startup wizard any time after you have configured the appliance. To start the wizard, from Grid Manager, select the **System** tab, and then click **System Properties -> Startup Wizard** from the Toolbar.

To make an HTTPS session to the appliance, you must be able to reach its IP address from the management system.

Note: If you have already set the IP address of the LAN1 port through the LCD or CLI so that you can reach it over the network—and you have already cabled the appliance to the network—you can skip the first step.

1. If you have not changed the default IP address (192.168.1.2/24) of the LAN1 port through the LCD or CLI—and the subnet to which you connect the appliance is not 192.168.1.0/24—put your management system in the 192.168.1.0/24 subnet and connect an Ethernet cable between the management system and the appliance.
2. Open an Internet browser window and enter `https:// <IP address of the appliance>` to make an HTTPS connection. For information about supported browsers, see [Supported Browsers](#) on page 46.

Several certificate warnings may appear during the login process because the preloaded certificate is self-signed and has the hostname `www.infoblox.com`, which may not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information, see [Managing Certificates](#) on page 53.

3. Enter the default username and password (admin and infoblox) on the Grid Manager login page, and then click **Login** or press Enter. For information, see [Logging in to the GUI](#) on page 48.
4. Read the Infoblox End-User License Agreement. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.

Click **I Accept**. The *NIOS Setup* wizard appears.

5. In the *NIOS Setup* wizard, select **Configuring a standalone appliance**. To configure an independent HA pair, see [Deploying an Independent HA Pair](#) on page 287.
6. Click **Next** and complete the following to configure network settings:
 - **Host Name:** Enter a valid domain name for the appliance.
 - **IP Address:** Displays the IP address of the LAN1 port, which in factory default is the 192.168.1.2/24 address. You can apply an IPv4 and IPv6 address for the network settings. (IPv4 configuration is required before configuring IPv6.) Ensure that you know the gateway IP address for either protocol type. As an example, once you define the initial IPv4 value for the IP Address, Subnet Mask and Gateway settings, you can add desired IPv6 settings after the appliance is up and running. NIOS Appliance ports may run IPv4 and IPv6 protocols simultaneously.
 - **Subnet Mask:** Displays the subnet mask of the LAN1 port.
 - **Gateway:** Displays the IP address of the gateway of the subnet on which the LAN1 port is set.
 - **Port Settings:** Select the port settings from the drop-down list. The list contains all the settings supported by the hardware model. The default is **Automatic**. The appliance automatically detects the port settings.

7. Click **Next** and complete the following to set admin password:
 - **Yes:** To change the default password.
 - **No:** To keep the default password. Infoblox recommends that you change the default password.
When you select **Yes**, complete the following:
 - **Password:** Enter a password for the superuser admin account. The password must be a single alphanumeric string without spaces and at least four characters long. The password is case-sensitive.
 - **Retype Password:** Enter the same password.
8. Click **Next** and complete the following to configure time settings:
 - **Time Zone:** Select the applicable time zone from the drop-down list. The default is **(UTC) Coordinated Universal Time**.
 - **Would you like to enable NTP?:**
 - Select **Yes** to synchronize the time with external NTP servers, and then click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter either the IP address or the resolvable host name of an NTP server. You can view a list of public NTP servers at ntp.isc.org.
 - Select **No** to specify the time settings for the appliance.
 - **Date:** Enter the date in YYYY-MM-DD format. You can also click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter the time in HH:MM:SS AM/PM format. You can also click the clock icon to select a time from the drop-down list.
9. If you want to participate in the Infoblox Customer Experience Improvement Program, complete the following and then click **Next**:
 - **Participate in the Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality. For more information about the program, see [Participating in the Customer Experience Improvement Program](#) on page 1023.
 - **Support ID (optional):** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. The value you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. Infoblox includes this ID in the data report.
 - **Email:** Enter an email address to which Infoblox sends a copy of the usage report. The email address you enter here is also displayed in the **Customer Improvement** tab in the *Grid Properties* editor. This is optional.
 - **Infoblox Privacy Policy:** Click here to view the Infoblox privacy policy. The appliance displays the policy in a new browser tab.
10. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
11. Click **Finish**.
The appliance restarts and disconnects Grid Manager.

CONFIGURATION EXAMPLE: DEPLOYING A NIOS APPLIANCE AS A PRIMARY DNS SERVER

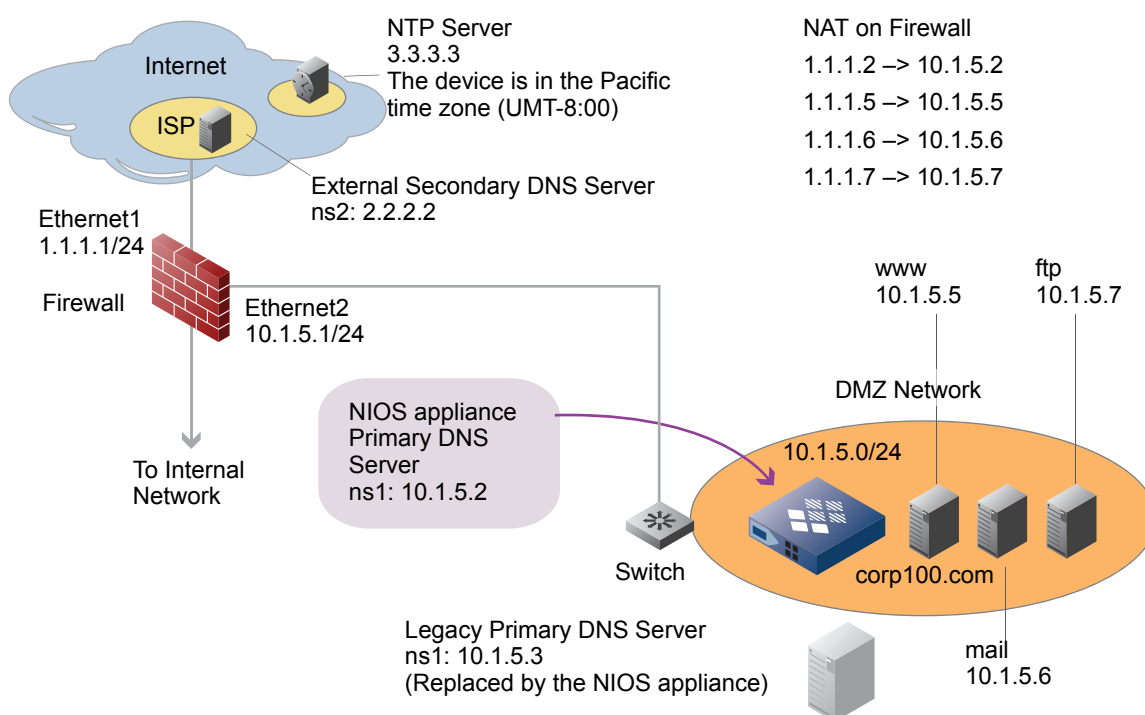
In this example, you configure the NIOS appliance as a primary DNS server for corp100.com. Its FQDN (fully-qualified domain name) is ns1.corp100.com. The interface IP address of the LAN1 port is 10.1.5.2/24. Because this is a private IP address, you must also configure the firewall to perform NAT (network address translation), mapping the public IP address 1.1.1.2 to 10.1.5.2. Using its public IP address, ns1 can communicate with appliances on the public network. The FQDN and IP address of the external secondary DNS server are ns2.corp100.com and 2.2.2.2. The ISP hosts this server.

The primary and secondary servers answer queries for the following public-facing servers in the DMZ:

- www.corp100.com
- mail.corp100.com
- ftp.corp100.com

When you create the corp100.com zone on the NIOS appliance, you import zone data from the legacy DNS server at 10.1.5.3.

Figure 6.5 Example 1 Network Diagram



The NIOS appliance is the primary DNS server for the corp100.com domain. It answers queries from the Internet for the three public-facing servers in the DMZ network:

- www.corp100.com
- mail.corp100.com
- ftp.corp100.com

Cabling the Appliance to the Network and Turning On Power

Connect an Ethernet cable from the LAN1 port of the NIOS appliance to a switch in the DMZ network and turn on the power. For information about installing and cabling the appliance, refer to the user guide or installation guide that ships with the product.

Specifying Initial Network Settings

Before you can configure the NIOS appliance through Grid Manager, you must be able to make a network connection to it. The default network settings of the LAN1 port are 192.168.1.2/24 with a gateway at 192.168.1.1 (the HA and MGMT ports do not have default network settings). To change these settings to suit your network, use either the LCD or the console port.

In this example, you change the IP address/netmask of the LAN1 port to 10.1.5.2/24, and the gateway to 10.1.5.1.

LCD

The NIOS appliance has an LCD and navigation buttons on its front panel.

At startup, the Infoblox logo appears in the LCD on the front panel of the appliance. Then the LCD scrolls repeatedly through a series of display screens.

1. To change the network settings from the default, press one of the navigation buttons.
The LCD immediately goes into input mode, in which you can enter the IP address, netmask, and gateway for the LAN1 port.
2. Use the navigation buttons to enter the following information:
 - IP Address: 10.1.5.2
 - Netmask: 255.255.255.0
 - Gateway: 10.1.5.1

Specifying Appliance Settings

When you make the initial HTTPS connection to the NIOS appliance, the NIOS Startup Wizard guides you through the basic deployment of the appliance on your network. Use the wizard to enter the following information:

- Deployment: single independent appliance
 - Host name: ns1.corp100.com
 - Password: SnD34n534
 - NTP (Network Time Protocol) server: 3.3.3.3; time zone: (UMT – 8:00 Pacific Time (US and Canada), Tijuana)
1. Open an Internet browser window and enter **https://10.1.5.2**.
 2. Accept the certificate when prompted.
Several certificate warnings may appear during the login process. This is normal because the preloaded certificate is self-signed and has the hostname www.infoblox.com, which does not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully-qualified domain name) of the appliance. This is a very simple process. For information about certificates, see [Creating a Login Banner](#) on page 49.
 3. Enter the default username and password (admin and infoblox) on the Grid Manager login page, and then click **Login** or press Enter. For information, see [Logging in to the GUI](#) on page 48.
 4. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed. Grid Manager may take a few seconds to load your user profile.
 5. In the *NIOS Startup* wizard, select **Configuring a standalone appliance**.

6. Click **Next** and complete the following to configure network settings:
 - **Host Name:** Enter `ns1.corp100.com`.
 - **IP Address:** Enter `10.1.5.2` as the IP address for the LAN1 port.
 - **Subnet Mask:** Enter `255.255.255.0` as the subnet mask for the LAN1 port.
 - **Gateway:** Enter `10.1.5.1` as the gateway of the subnet on which the LAN1 port is set.
 - **Port Settings:** Use the default value **Automatic**.
7. Click **Next** and complete the following to set admin password:
 - **Would you like to set admin password?:** Click **Yes**.
 - **Password:** Enter `sND34n534`.
 - **Retype Password:** Enter `sND34n534` again.
8. Click **Next** and complete the following to configure the time settings:
 - **Time Zone:** Select **UMT – 8:00 Pacific Time (US and Canada), Tijuana** from the drop-down list.
 - **Would you like to enable NTP?:** Select **Yes** to synchronize the time with external NTP servers, and then click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter `3.3.3.3` in the **NTP Server** field.
9. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
10. Click **Finish**.

Enabling Zone Transfers on the Legacy Name Server

To allow the appliance to import zone data from the legacy server 10.1.5.3, you must configure the legacy server to allow zone transfers to the appliance at 10.1.5.2.

Legacy BIND Server

1. Open the `named.conf` file using a text editor and change the `allow-transfer` statement as shown below:
For All Zones — To set the `allow-transfer` statement as a global statement in the `named.conf` file for all zones:

```
options {
    zone-statistics yes;
    directory "/var/named/named_conf";
    version "";
    recursion yes;
    listen-on { 127.0.0.1; 10.1.5.3; };
    ...
    allow-transfer {10.1.5.2; };
    transfer-format many-answers;
};
```

- For a Single Zone** — To set the `allow-transfer` statement in the `named.conf` file for the `corp100.com` zone:

```
zone "corp100.com" in {
    type master;
    allow-transfer {10.1.5.2;};
    notify yes;
};
```

2. After editing the `named.conf` file, restart DNS service on the appliance for the change to take effect.

Legacy Windows 2000/2003 Server

1. Click **Start** -> **All Programs** -> **Administrative Tools** -> **DNS**.
2. Click **+** (for `ns1`) -> **+** (for Forward Lookup Zones) -> **corp100.com**.
3. Right-click **corp100.com**, and then select **Properties** -> **Zone Transfers**.
4. On the *Zone Transfers* page in the *corp100.com Properties* dialog box, enter the following:

- **Allow zone transfers:** Select this.
 - **Only to the following servers:** Select this.
 - **IP address:** Enter **10.1.5.2**, and then click **Add**.
5. To save the configuration and close the *corp100.com Properties* dialog box, click **OK**.

Importing Zone Data on an Independent Appliance

You can import zone data from a legacy server or manually enter it. When you import both forward-mapping and reverse-mapping zone data, the NIOS appliance automatically creates Infoblox host records if corresponding A and PTR records are present. You can then modify the host records to add MAC addresses. However, if you only import forward-mapping zone data, the NIOS appliance cannot create host records from just the A records. In that case, because you cannot later convert A records to host records, it is more efficient to create the corp100.com zone, and define host records manually.

Infoblox host records are data models that represent IP devices within the Infoblox semantic database. The NIOS appliance uses a host object to define A, PTR, and CNAME resource records in a single object as well as a DHCP fixed address if you include a MAC address in the host object definition. The host object prevents costly errors because you only maintain a single object for multiple DNS records and a DHCP fixed address. Therefore, it is advantageous to use host records instead of separate A, PTR, and CNAME records.

Note: If you only have forward-mapping zones on your legacy servers and you want to add reverse-mapping zones and automatically convert A records to host records in the imported forward-mapping zones and create reverse host records in corresponding reverse-mapping zones, create the reverse-mapping zones on the NIOS appliance and then import the forward-mapping zones data. The NIOS appliance automatically converts the imported A records to host records in the forward-mapping zones and creates reverse host records in the reverse-mapping zones.

You also have the option of using the Data Import Wizard for loading DNS and DHCP data. For large data sets, this option is an efficient approach. To download the Data Import Wizard, visit www.infoblox.com/import/.

In this example, when you create the corp100.com forward-mapping zone, you import zone data for the existing corp100.com zone from the legacy server at 10.1.5.3. When you create the 1.1.1.0/24 reverse-mapping zone, you also import the reverse-mapping zone records from the legacy server. After the appliance has both the forward- and reverse-mapping zone data, it converts the A and PTR records to Infoblox host records.

Creating a Name Server Group

1. Open an Internet browser window, enter **https://10.1.5.2**, and then log in to Grid Manager using the username **admin** and password **SnD34n534**.
2. From the **Data Management** tab, select the **DNS** tab → **Name Server Groups** tab, and then click the Add icon to create a name server group.
3. In the *Name Server Group* wizard, complete the following:
 - **Name:** Enter **corp100** as the group name.
 - **Name Servers:** Click the Add icon → **Primary**.
 - In the *Add Primary* section, Grid Manager displays the host name of the independent appliance. Click **Add**. Grid Manager adds the independent system as the primary server.
 - Click the Add icon → **External Secondary**.
 - In the *Add External Secondary* section, complete the following:
 - **Name:** Enter **ns2.corp100.com**.
 - **Address:** Enter **2.2.2.2**.
 - **Stealth:** Clear this check box.
 - Click **Add**. Grid Manager adds the external secondary to the name server group.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Creating a Forward-Mapping Zone

Note: To import zone data, you must first create a zone and save it.

1. To create an authoritative zone, from the **Data Management** tab, select the **DNS** tab -> **Zones** tab, and then click the Add icon -> **Authoritative Zone**.
2. In the *Add Authoritative Zone* wizard, select **Add an authoritative forward-mapping zone**.
3. Click **Next** and complete the following:
 - **Name:** Enter `corp100.com`.
 - **Comment:** Enter `DNS zone`.
4. Click **Next** to assign a name server group to the zone.
5. Click the **Zones** tab, select the **corp100.com** check box, and then click the Edit icon.
6. In the *Authoritative Zone* editor, select the **Name Servers** tab, and then complete the following:
 - **Use this name server group:** Select this, and then select **Corp100** from the drop-down list.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

Importing Zone Data

1. To import zone data to the corp100.com zone that you created earlier, click the **Zones** tab, select the **corp100.com** check box, and then click **Import Zone** from the Toolbar.
2. In the *Import Zone* editor, complete the following:
 - **Address:** Enter the IP address `10.1.5.3` from which you want to import zone data.
 - **Automatically create hosts from A records:** Select this to enable the appliance to create host records from the imported A records.
3. Click **Import**.
4. After successfully importing the zone data, click **corp100.com** in the **Zones** tab.
You can see all the imported forward-mapping zone data in the *Records* panel. Because you have not yet imported the reverse-mapping zone data, most of the records appear as A records.
5. To import the reverse-mapping zone data, from the **Zones** tab, click the Add icon -> **Authoritative Zone**.
6. In the *Add Authoritative Zone* wizard, select **Add an authoritative IPv4 reverse-mapping zone**.
7. Click **Next** and complete the following:
 - **IPv4 Network:** Enter `1.1.1.0`.
 - **Netmask:** Select **24** from the drop-down list.
 - **Comment:** Enter `Reverse-mapping zone`.
8. Click **Save & Close**.
9. To assign a name server group to the reverse-mapping zone, click the **Zones** tab, select the **1.1.1.in-addr.arpa** check box, and then click the Edit icon.
10. In the *Authoritative Zone* editor, select the **Name Servers** tab, and then complete the following:
 - **Use this name server group:** Select this, and then select **Corp100** from the drop-down list.
11. Click **Save & Close**.
12. To import reverse-mapping zone data, click the **Zones** tab, select the **corp100.com** check box, and then click **Import Zone** from the Toolbar.
13. In the *Import Zone* editor, complete the following:
 - **Address:** Enter `10.1.5.3` from which you want to import zone data.
 - **Automatically create hosts from A records:** Select this to enable the appliance to create host records from the imported A records.

14. Click **Import**.
15. After successfully importing the zone data, click **1.1.1.in-addr.arpa** in the **Zones** tab.
You can see all the imported reverse-mapping zone data in the *Records* panel.
16. Click **corp100.com** in the Forward Mapping Zones list.
Because you have now imported both the forward- and reverse-mapping zone data, most of the records appear as host records.
17. Finally, you must remove the ns1 host record for the legacy server (value 1.1.1.3). To remove it, select the **ns1** check box (the host record for 1.1.1.3), and then click the Delete icon.

Designating the New Primary on the Secondary Name Server (at the ISP Site)

In this example, the external secondary name server is maintained by an ISP, so you must contact your ISP administrator to change the IP address of the primary (or *master*) name server. (If you have administrative access to the secondary name server, you can make this change yourself.)

Because a firewall performing NAT exists between the secondary and primary name servers, specify the NAT address 1.1.1.2 for the primary name server instead of 10.1.5.2.

Secondary BIND Server

1. Open the named.conf file using a text editor and set ns1 (with NAT address 1.1.1.2) as the primary (or *master*) from which ns2 receives zone transfers in the named.conf file for the corp100.com zone:

```
zone "corp100.com" in {
    type slave;
    masters {1.1.1.2;};
    notify yes;
    file "/var/named/db.corp100.com";
};
```
2. After editing the named.conf file, restart DNS service for the change to take effect.

Secondary Windows 2000/2003 Server

1. Click **Start -> All Programs -> Administrative Tools -> DNS**.
2. Click **+** (for ns2) -> **+** (for Forward Lookup Zones) -> **corp100.com**.
3. Right-click **corp100.com**, and then select **Properties -> General**.
4. On the *General* page in the *corp100.com Properties* dialog box, enter the following:
 - **Zone file name:** corp100.com.dns
 - **IP address:** Enter **1.1.1.2**, and then click **Add**.
 - In the IP Address field, select **1.1.1.3** (the NAT IP address of the legacy DNS server), and then click **Remove**.
5. To save the configuration and close the *corp100.com Properties* dialog box, click **OK**.

Configuring NAT and Policies on the Firewall

Change the NAT and policy settings on the firewall to allow bidirectional DNS traffic to and from ns1.corp100.com and NTP traffic from ns1.corp100.com to the NTP server at 3.3.3.3.

For example, enter the following commands on a Juniper firewall running ScreenOS 4.x or later:

```
set address dmz ns1 10.1.5.2/32
set address untrust ntp_server 3.3.3.3/32
set interface ethernet1 mip 1.1.1.2 host 10.1.5.2
set policy from dmz to untrust ns1 any dns permit
set policy from untrust to dmz any mip(1.1.1.2) dns permit
set policy from dmz to untrust ns1 ntp_server ntp permit
```

At this point, the new DNS server can take over DNS service from the legacy server. You can remove the legacy server and unset any firewall policies permitting traffic to and from 10.1.5.3.

DEPLOYING AN INDEPENDENT HA PAIR

To deploy an independent HA pair, you cable the HA and LAN1, LAN1 (VLAN), or LAN2, LAN2 (VLAN) ports to the network and configure the IP settings for these ports and the VIP address within the same subnet. For more information about HA pairs, see [About HA Pairs](#) on page 233.

The default LAN1 or LAN2 settings are as follows:

- IP address: 192.168.1.2
- Netmask: 255.255.255.0
- Gateway: 192.168.1.1.

You can configure an HA pair using the Infoblox NIOS Startup Wizard. IPv4 and IPv6 network values are supported for the NIOS Startup Wizard and for HA Pair configuration. The NIOS appliance MGMT port also can be configured to support an IPv6 address.

- Requirements: HTTPS connections from your management system to the Ethernet ports on the two appliances.
- Advantage: The startup wizard provides step-by-step guidance for configuring the network settings of the VIP address and HA and LAN1 or LAN1 (VLAN) ports on both nodes, for setting the host name, admin password, and system clock, and—if using NTP (Network Time Protocol)—for enabling the HA pair as an NTP client.

Using the Infoblox NIOS Startup Wizard to Configure an HA Pair

When you first make an HTTPS connection to the NIOS appliance, the Infoblox NIOS Startup Wizard guides you through various deployment options, basic network settings, and opportunities for changing the password of the superuser *admin* and for setting the system clock.

Configuring the Connecting Switch

To ensure that VRRP (Virtual Router Redundancy Protocol) works properly, configure the following settings at the port level for all the connecting switch ports (HA, LAN1, LAN1 (VLAN), LAN2, and LAN2 (VLAN)):

- Spanning Tree Protocol: Disable. For vendor specific information, search for “HA” in the Infoblox Knowledge Base system at <http://www.infoblox.com/en/support/support-center-login.html>.
- Trunking: Disable
- EtherChannel: Disable
- IGMP Snooping: Disable
- Port Channeling: Disable
- Speed and Duplex settings: Match these settings on both the Infoblox appliance and switch
- Disable other dynamic and proprietary protocols that might interrupt the forwarding of packets

Note: By default, a NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between its LAN1, HA, and MGMT ports and the Ethernet ports on the connecting switch. If the two appliances fail to auto-negotiate the optimal settings, see [Modifying Ethernet Port Settings](#) on page 354 for steps you can take to resolve the problem.

Putting Both Nodes on the Network

1. Use one of the methods described in [Deploying a Single Independent Appliance](#) on page 276 to configure the network settings of the LAN1 port of each node so that they are on the same subnet and you can reach them across the network.
2. Cable the LAN1 port and the HA port on each node to the network switch.

Note: The Ethernet ports on the Infoblox-250, -250-A, -550, -550-A, -1050, -1050-A, -1550, -1550-A, -1552, -1552-A, 1852-A and -2000 appliances are autosensing, so you can use either a straight-through or cross-over Ethernet cable for these connections.

3. Cable your management system to the network switch.

Configuring Node 1

1. Open an Internet browser window and enter `https:// <the IP address of the appliance>` to make an HTTPS connection to the first node. For information about supported browsers, see [Supported Browsers](#) on page 46.

Several certificate warnings may appear during the login process because the preloaded certificate is self-signed and has the hostname `www.infoblox.com`, which may not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information, see [Creating a Login Banner](#) on page 49.

2. Enter the default username and password (admin and infoblox) on the Grid Manager login page, and then click **Login** or press Enter. For information, see [Logging in to the GUI](#) on page 48.
3. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed. Grid Manager may take a few seconds to load your user profile.
4. In the *NIOS Startup* wizard, select **Configuring an HA pair**. Click **Yes** for the first appliance of the HA pair.
5. Click **Next** and complete the following to configure network settings:
 - **Host Name:** Enter a valid domain name for the node.
 - **HA Pair Name:** Enter a name for the HA pair. The default name is **Infoblox**.
 - **Shared Secret:** Enter the shared secret that both nodes use to authenticate each other when establishing a VPN tunnel for ensuing bloxSYNC traffic. The default shared secret is **test**.
 - **Show Password:** Select this to display the shared secret. Clear it to conceal the shared secret.
6. Click **Next** and complete the following to set properties for the first node:
 - **Virtual Router ID:** Enter the VRID (virtual router ID). This must be a unique VRID number—from 1 to 255—for this subnet.
 - **Required Ports and Addresses:** Enter information for the interfaces VIP, Node 1 HA, Node 2 HA, Node 1 LAN, and Node 2 LAN. Some fields are prepopulated by Grid Manager based on the existing configuration of the appliance. All fields are required. Click the empty fields and complete the following information:
 - **Address:** IP address of the interface.
 - **Subnet Mask:** The subnet mask of the interface.
 - **Gateway:** The IP address of the gateway for the subnet on which the interfaces are set. This is the same for all interfaces.

- **Port Settings:** Select the port settings from the drop-down list. The list contains all settings supported by the hardware model. The default is **Automatic**. The appliance automatically detects the port settings.
7. Click **Next** and complete the following to set admin password:
 - **Yes:** To change the default password.
 - **No:** To keep the default password.
 If you select **Yes**, complete the following:
 - **Password:** Enter a password for the superuser admin account. The password cannot contain spaces and it must be at least four characters long. The password is case-sensitive.
 - **Retype Password:** Enter the same password.
 8. Click **Next** and complete the following to configure time settings:
 - **Time Zone:** Select the applicable time zone from the drop-down list. The default is **(UTC) Coordinated Universal Time**.
 - **Would you like to enable NTP?:**
 - Select **Yes** to synchronize the time with external NTP servers. Click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter either the IP address or the resolvable host name of an NTP server. You can view a list of public NTP servers at ntp.isc.org.
 - Select **No** to specify a date and time.
 - **Date:** Enter the data in YYYY-MM-DD format. You can also click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter the time in HH:MM:SS AM/PM format. You can also click the clock icon to select a time from the drop-down list.
 9. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
 10. Click **Finish**.

Configuring Node 2

1. Open an Internet browser window and enter `https:// <the IP address of the appliance>` to make an HTTPS connection to the second node. For information about supported browsers, see [Supported Browsers](#) on page 46.
 Several certificate warnings may appear during the login process because the preloaded certificate is self-signed and has the hostname `www.infoblox.com`, which may not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information, see [Creating a Login Banner](#) on page 49.
2. Enter the default username and password (admin and infoblox) on the Grid Manager login screen, and then click **Login** or press Enter. For information, see [Logging in to the GUI](#) on page 48.
3. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed. Grid Manager may take a few seconds to load your user profile.
4. In the *NIOS Startup* wizard, select **Configuring an HA pair** to configure an independent HA pair. Click **No** to configure the second node of the HA pair.
5. Click **Next** and complete the following to configure network settings:
 - **HA Virtual IP address:** Enter the VIP (virtual IP) address and its netmask.
 - **HA Pair Name:** Enter a name for the HA pair. The default name is **Infoblox**. Ensure that you use the same name as the first node.
 - **Shared Secret:** Enter a text string that both nodes use as a shared secret to authenticate each other when establishing a VPN tunnel. The default shared secret is `test`. This must be the same shared secret that you entered on the first appliance.
 - **Show Password:** Click this to display the shared secret. Clear it to conceal the shared secret.

6. Click **Next**, and then complete the following to set properties for the second appliance:
 - **IP Address:** Enter the IP address of the appliance.
 - **Subnet Mask:** Enter the subnet mask of the appliance.
 - **Gateway:** Enter the IP address of the gateway of the subnet of the interface.
7. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
8. Click **Finish**.

The setup of the HA pair is complete. When you next make an HTTPS connection to the HA pair, use the VIP address.

CONFIGURATION EXAMPLE: CONFIGURING AN HA PAIR FOR INTERNAL DNS AND DHCP SERVICES

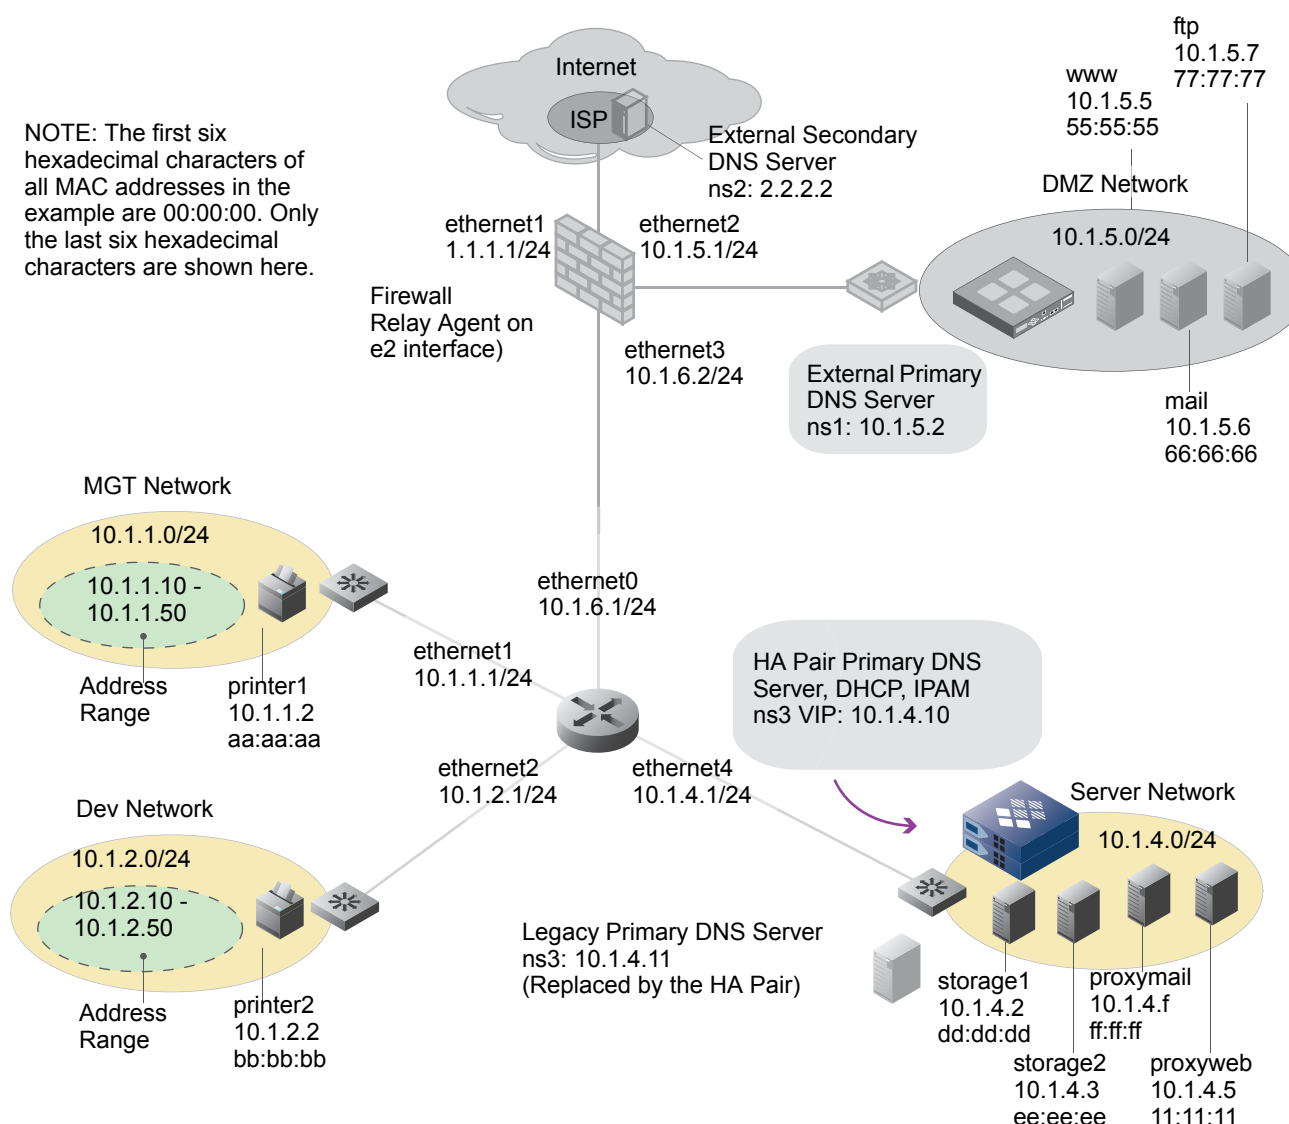
In this example, you set up an HA pair of NIOS appliances to provide internal DNS and DHCP services. The HA pair answers internal queries for all hosts in its domain (corp100.com). It forwards internal queries for external sites to ns1.corp100.com at 10.1.5.2 and ns2.corp100.com at 2.2.2.2. It also uses DHCP to provide dynamic and fixed addresses.

The HA pair consists of two appliances (nodes). The IP addresses of the VIP (virtual IP) address of the HA pair and the HA and LAN1 ports on each node, are as follows:

HA Pair IP Addresses	
VIP 10.1.4.10 (the address that the active node of the HA pair uses)	
Node 1	Node 2
<ul style="list-style-type: none">• LAN1 10.1.4.6• HA 10.1.4.7	<ul style="list-style-type: none">• LAN1 10.1.4.8• HA 10.1.4.9

The virtual router ID number for the HA pair is 150. The ID number must be unique for this network segment. When you create the corp100.com zone on the HA pair, you import DNS data from the legacy server at 10.1.4.11.

Figure 6.6 Example 2 Network Diagram



An HA pair of NIOS appliances provides internal DNS services. It answers internal queries for all hosts in its domain. It forwards internal queries for external sites to ns1 and ns2. It also serves DHCP, providing both dynamic and fixed addresses. For information on configuring the NIOS appliance as a primary DNS server, see [Configuration Example: Deploying a NIOS Appliance as a Primary DNS Server](#) on page 281.

Cabling Appliances to the Network and Turning On Power

Connect Ethernet cables from the LAN1 and HA ports on both NIOS appliances to a switch in the server network and turn on the power for both appliances. For information about installing and cabling the appliance, refer to the user guide or installation guide that ships with the product.

Specifying Initial Network Settings

Before you can configure the appliances through Grid Manager, you must be able to make a network connection to them. The default network settings of the LAN1 port are 192.168.1.2/24 with a gateway at 192.168.1.1 (the HA and MGMT ports do not have default network settings). To change these settings, you can use the LCD or make a console connection to each appliance.

Node 1

Using the LCD or console port on one of the appliances, enter the following information:

- IP Address: 10.1.4.6 (for the LAN1 port)
- Netmask: 255.255.255.0
- Gateway: 10.1.4.1

Node 2

Using the LCD or console port on the other appliance, enter the following information:

- IP Address: 10.1.4.8 (for the LAN1 port)
- Netmask: 255.255.255.0
- Gateway: 10.1.4.1

After you confirm your network settings, the Infoblox GUI application automatically restarts.

Specifying Appliance Settings

When you make the initial HTTPS connection to a NIOS appliance, the Infoblox NIOS Startup Wizard guides you through the basic deployment of the appliance on your network. To set up an HA pair, you must connect to and configure each appliance individually.

Node 1

1. Open an Internet browser window and enter **https://10.1.4.6**.
2. Accept the certificate when prompted.
Several certificate warnings may appear during the login process. This is normal because the preloaded certificate is self-signed and has the hostname `www.infoblox.com`, which does not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully-qualified domain name) of the appliance. This is a very simple process. For information about certificates, see [Creating a Login Banner](#) on page 49.
3. Enter the default username and password (admin and infoblox) on the Grid Manager login page, and then click **Login** or press Enter. For information, see [Logging in to the GUI](#) on page 48.
4. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed. Grid Manager may take a few seconds to load your user profile.
5. In the *NIOS Startup* wizard, select **Configuring an HA pair**. Click **Yes** to configure the first appliance.
6. Click **Next** and complete the following to configure network settings:
 - **Host Name:** Enter `ns3.corp100.com`.
 - **HA Pair Name:** Use the default name `Infoblox`.
 - **Shared Secret:** Enter `37eeT1d`.

7. Click **Next** and complete the following to set properties for the first node:
 - **Virtual Router ID:** Enter **150**.
 - **Required Ports and Addresses:** In the table, click the empty fields and enter the following information for each corresponding interface:
 - **VIP:** **10.1.4.10**
 - **Node 1 HA:** **10.1.4.7**
 - **Node 2 HA:** **10.1.4.9**
 - **Node 1 LAN:** **10.1.4.6**
 - **Node 2 LAN:** **10.1.4.8**
 - **Subnet Mask:** **255.255.255.0**
 - **Gateway:** **10.1.4.1**

Note: Some fields are prepopulated by Grid Manager based on the existing configuration of the appliance. All fields are required.

8. Click **Next** and complete the following to set admin password:
 - **Would you like to set admin password?:** Click **No**.
9. Click **Next** and complete the following to configure time settings:
 - **Time Zone:** Select **UMT – 8:00 Pacific Time (US and Canada), Tijuana** from the drop-down list.
 - **Would you like to enable NTP?:** Select **Yes** to synchronize the time with external NTP servers, and then click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter **3.3.3.3** in the **NTP Server** field.
10. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
11. Click **Finish**.

Node 2

1. From the **System** tab, select the **System Manager** tab, and then click **System Properties** -> **Setup Wizard** from the Toolbar.
2. In the *NIOS Startup* wizard, select **Configuring an HA pair** to configure an independent HA pair. Click **No** for configuring node 2 of the HA pair.
3. Click **Next**, and then complete the following to configure network settings:
 - **HA Virtual IP address:** Enter **10.1.4.10**.
 - **HA Pair Name:** Use the default name **Infoblox**.
 - **Shared Secret:** Enter **37eeT1d**.
 - **Show Password:** Click this to display the shared secret.
4. Click **Next**, and then complete the following to set properties for the second appliance:
 - **IP Address:** Enter **10.1.4.8**.
 - **Subnet Mask:** Enter **255.255.255.0**.
 - **Gateway:** Enter **10.1.4.1**.
5. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
6. Click **Finish**.

The setup of the HA pair is complete. From now on, when you make an HTTPS connection to the HA pair, use the VIP address 10.1.4.10.

Enabling Zone Transfers

To allow the NIOS appliance to import zone data from the legacy server at 10.1.4.11, you must configure the legacy server to allow zone transfers to the appliance at 10.1.4.10.

Legacy BIND Server

1. Open the `named.conf` file using a text editor and change the `allow-transfer` statement to allow zone transfers to the appliance at 10.1.4.10. For a sample of the required changes to the `named.conf` file, see [Legacy BIND Server](#) on page 283.
2. After editing the `named.conf` file, restart DNS service for the change to take effect.

Legacy Windows 2000/2003 Server

Navigate to the *corp100.com Properties* dialog box, and add 10.1.4.10 to the list of IP addresses to which you want to allow zone transfers. For more detailed navigation and configuration instructions, see [Legacy Windows 2000/2003 Server](#) on page 283.

Importing Zone Data

You can import zone data from a legacy server to an independent HA pair, as described in [Importing Zone Data on an Independent Appliance](#) on page 284. Use the following information:

- Forward-mapping zone: `corp100.com`
- Import zone from: `10.1.4.11`
- Reverse-mapping zone: `1.1.1.0`

Defining Networks, Reverse-Mapping Zones, DHCP Ranges, and Infoblox Hosts

In this task, you enter data manually. For large data sets, you have the option of using the Data Import Wizard for loading DNS and DHCP configurations and data to make the process more efficient. To download the Data Import Wizard, visit www.infoblox.com/import/.

Networks

You can create all the subnetworks individually (which in this example are 10.1.1.0/24, 10.1.2.0/24, 10.1.4.0/24, and 10.1.5.0/24), or you can create a parent network (10.1.0.0/16) that encompasses all the subnetworks and then use the Infoblox split network feature to create the individual subnetworks automatically. The split network feature accomplishes this by using the IP addresses that exist in the forward-mapping zones to determine which subnets it needs to create. This example uses the split network feature. For information about creating networks, see [Configuring IPv4 Networks](#) on page 845.

1. From the **Data Management** tab, select the **IPAM** tab, and then click **Add -> Add IPv4 Network** from the Toolbar.
2. In the *Add Network* wizard, complete the following:
 - **Address:** 10.1.0.0
 - **Netmask:** Use the netmask slider to select the /16 (255.255.0.0) netmask.
3. Click **Next** to select a server. Click the Add icon. Grid Manager displays **ns3.corp100.com** in the table.
4. Click **Save & Close**.
5. In the IPAM tab, select the **10.1.0.0/16** check box, and then select **Split** from the Toolbar.

6. In the *Split Network* dialog box, complete the following:
 - **Subnetworks:** Move the slider to 24.
 - **Immediately Add:** Select **Only networks with ranges and fixed addresses.**
 - **Automatically create reverse-mapping zones:** Select this check box.
7. Click **OK**.
 The appliance creates the following 24-bit subnets for the imported Infoblox hosts:
 - 10.1.1.0/24
 - 10.1.2.0/24
 - 10.1.4.0/24
 - 10.1.5.0/24
8. From the **IPAM** tab, click the **10.1.1.0/24** check box, and then click the Edit icon.
9. In the *DHCP Network* editor, enter information in the following tabs:
 - General*
 - **Comment:** MGT
 - Server Assignment*
 - Add **ns3.corp100.com** as a server.
10. Click **Save & Close**.
11. To modify the other networks, repeat steps #8 – 10 for each network and use the following information:
 - 10.1.2.0/24 Network:
 - **Comment:** Dev
 - **Server Assignment:** ns3.corp100.com
 - 10.1.4.0/24 Network:
 - **Comment:** Server
 - **Server Assignment:** ns3.corp100.com
 - 10.1.5.0/24 Network:
 - **Comment:** DMZ
 - **Server Assignment:** ns3.corp100.com

DHCP Ranges

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *10.1.1.0/24*, and then click **Add -> DHCP Range** from the Toolbar.
2. In the *Add Range* wizard, complete the following:
 - Start:** 10.1.1.10
 - End:** 10.1.1.50
3. Click **Next**, and then select **Server**, Grid Manager displays **ns3.corp100.com** as the assigned member.
4. Click **Save & Close**.
5. In the **Networks** tab, click *10.1.2.0/24*, and then click **Add -> DHCP Range** from the Toolbar.
6. In the *Add Range* wizard, complete the following:
 - Start:** 10.1.2.10
 - End:** 10.1.2.50
7. Click **Next**, and then select **Server**, Grid Manager displays **ns3.corp100.com** as the assigned member.
8. Click **Save & Close**.

Infoblox Hosts

Defining both a MAC and IP address for an Infoblox host definition creates a DHCP host entry—like a fixed address—that you can manage through the host object. To add a MAC address to each host record that the appliance created when you imported forward- and reverse-mapping zone records:

1. From the **Data Management** tab, select the **IPAM** tab -> *10.1.1.0/24* -> *10.1.1.2*.
2. In the **Related Objects** tab, select the check box of the host record, and then click the Edit icon.
3. In the *Host Record* editor, click the MAC Address field, and then enter the following:
 - **MAC Address:** 00:00:00:aa:aa:aa
4. Click **Save & Close**.
5. Follow steps 1 – 4 to modify hosts with the following information:

printer2

- IP Address: 10.1.2.2
- MAC Address: 00:00:00:bb:bb:bb

storage1

- IP Address: 10.1.4.2
- MAC Address: 00:00:00:dd:dd:dd

storage2

- IP Address: 10.1.4.3
- MAC Address: 00:00:00:ee:ee:ee

proxymail

- IP Address: 10.1.4.4
- MAC Address: 00:00:00:ff:ff:ff

proxyweb

- IP Address: 10.1.4.5
- MAC Address: 00:00:00:11:11:11

www

- IP Address: 10.1.5.5
- MAC Address: 00:00:00:55:55:55

mail

- IP Address: 10.1.5.6
- MAC Address: 00:00:00:66:66:66

ftp

- IP Address: 10.1.5.7
- MAC Address: 00:00:00:77:77:77

Defining Multiple Forwarders

Because ns3.corp100.com is an internal DNS server, you configure it to forward DNS queries for external DNS name resolution to the primary and secondary DNS servers—ns1.corp100.com at 10.1.5.2 and ns2.corp100.com at 2.2.2.2.

1. From the **Data Management** tab, select the **DNS** tab, and then select **System DNS Properties** from the Toolbar.
2. In the *System DNS Properties* editor, click the Add icon in the **Forwarders** tab. Grid Manager adds a row to the table. Complete the following:
 - **Address:** Type **2.2.2.2**. Click Add again to add another forwarder.
 - **Address:** Type **10.1.5.2**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Each of the forwarders is assigned a random response time. The appliance sends the initial outbound query to the forwarder that has the lowest response time. If the first forwarder does not reply, the appliance tries the one with the next lowest random response time. The appliance adjusts and keeps track of the response times of the forwarders, and uses the quicker one for future queries. If the quicker forwarder does not respond, the appliance then uses another one.

Enabling Recursion on External DNS Servers

Because the HA pair forwards outbound queries to the two external DNS servers ns1.corp100.com (10.1.5.2) and ns2.corp100.com (2.2.2.2) for resolution, you must enable recursion on those servers. When a DNS server employs recursion, it queries other DNS servers for a domain name until it either receives the requested data or an error that the requested data cannot be found. It then reports the result back to the server that queried—in this case, the internal DNS server ns3.corp100.com (10.1.4.10), which in turn reports back to the DNS client.

Infoblox Server in the DMZ Network (ns1.corp100.com, 10.1.5.2)

1. From the **Data Management** tab, select the **DNS** tab, and then click **System DNS Properties** from the Toolbar.
2. In the *System DNS Properties* editor, select the **Allow Recursion** check box from the **Queries** tab, and then click the Add icon → **IPv4 Address**. Grid Manager adds a row to the **Allow recursive queries from** table. Complete the following:
 - **Permission:** Select **Allow** from the drop-down list.
 - **Name:** Enter **10.1.1.52**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

BIND Server at ISP Site (ns2.corp100.com, 2.2.2.2)

1. Open the named.conf file using a text editor and change the recursion and allow-recursion statements to allow recursive queries from 1.1.1.8 (the NAT address of ns3).

```
options {
zone-statistics yes;
directory "/var/named/named_conf";
version "";
recursion yes;
listen-on { 127.0.0.1; 2.2.2.2; };
...
allow-recursion {1.1.1.8;};
transfer-format many-answers;
};
```

2. After editing the named.conf file, restart DNS service for the change to take effect.

Windows 2000/2003 Server at ISP Site (ns2.corp100.com, 2.2.2.2)

1. Click **Start -> All Programs -> Administrative Tools -> DNS**.
2. Right-click **ns3**, and then select **Properties -> Advanced**.
3. On the *Advanced* page in the *ns3 Properties* dialog box, clear the **Disable recursion** check box.
4. To save the configuration change and close the *ns3 Properties* dialog box, click **OK**.

Modifying the Firewall and Router Configurations

Configure the firewall and router in your internal network to allow the following DHCP, DNS, and NTP traffic:

- To allow messages to pass from the DHCP clients in the DMZ—the web, mail, and FTP servers—to ns3 in the Server network, configure policies and DHCP relay agent settings on the firewall.
- To forward DHCP messages from DHCP clients in the MGT and Dev networks to ns3 in the Server network, configure relay agent settings on the router.
- To translate the private IP address of ns3 (10.1.4.10) to the public IP address (1.1.1.8) when forwarding DNS queries from ns3 to ns2, set a MIP (mapped IP) address on the firewall.
- To allow DNS queries from ns3 to ns1 and ns2 and NTP traffic from ns3 to the NTP server, configure firewall policies.

Firewall

For example, enter the following commands on a Juniper firewall running ScreenOS 4.x or later:

DHCP Relay Configuration

```
set address trust ns3 10.1.4.10/32
set interface ethernet2 dhcp relay server-name 10.1.4.10
set policy from dmz to trust ns1 ns3 DHCP-Relay permit
```

DNS Forwarding

```
set interface ethernet1 mip 1.1.1.8 host 10.1.4.10
set policy from trust to untrust ns3 ns2 dns permit
set policy from trust to dmz ns3 ns1 dns permit
```

NTP

```
set policy from dmz to untrust ns1 ntp_server ntp permit
```

Router

For example, enter the following commands on a Cisco router running IOS for release 12.x or later:

DHCP Relay Configuration

```
interface ethernet1
  ip helper-address 10.1.4.10
interface ethernet2
  ip helper-address 10.1.4.10
```

Enabling DHCP and Switching Service to the NIOS Appliance

With the Infoblox in place and the firewall and router configured for relaying DHCP messages, you can switch DHCP service from the legacy DHCP server at 10.1.4.11 to the HA pair at 10.1.4.10 (VIP address).

Tip: To minimize the chance of duplicate IP address assignments during the transition from the legacy DHCP server to the appliance, shorten all lease times to a one-hour length in advance of the DHCP server switch. Then, when you take the legacy DHCP server offline, the DHCP clients quickly move to the new server when their lease renewal efforts fail and they broadcast DHCPDISCOVER messages. To determine how far in advance you need to shorten the lease length, find the longest lease time (for example, it might be two days). Then change the lease length to one hour at a slightly greater interval of time before you plan to switch DNS service to the appliance (for example, three days before the switch over). By changing the lease length this far in advance, you can be sure that all DHCP leases will be one-hour leases at the time of the switch-over. If the longest lease length is longer—such as five days—and you want to avoid the increased amount of traffic caused by more frequent lease renewals over a six-day period, you can also employ a stepped approach: Six days before the switch-over, change the lease lengths to one-day leases. Then two days before the switch-over, change them to one-hour leases.

1. Open an Internet browser window, enter **https://10.1.4.10**, and then log in to the appliance using the username **admin** and password **SnD34n534**.
2. From the **Data Management** tab, select the **DHCP** tab, and then click **Start** from the Toolbar.
3. In the *Start Member DHCP Service* dialog box, click **Yes**.
The HA pair is ready to provide DHCP service to the network.
4. Take the legacy DHCP server at 10.1.4.11 offline.
When the DHCP clients are unable to renew their leases from the legacy DHCP server, they broadcast DHCPDISCOVER messages to which the new DHCP server responds.

Managing and Monitoring

Infoblox provides tools for managing IP address usage and several types of logs to view events of interest and DHCP and DNS data. After configuring the appliance, you can use the following resources to manage and monitor IP address usage, DNS and DHCP data, and administrator and appliance activity.

IPAM (IP Address Management)

IPAM offers the following services:

- Simple IP address modification – Within a single IP address-centric data set, you can modify the Infoblox host, DHCP, and DNS settings associated with that IP address.
- Address type conversion – Through IPAM functionality, you can make the following conversions:
 - Currently active dynamic addresses to fixed addresses, reserved addresses, or Infoblox hosts.
 - Fixed addresses to reservations or hosts.
 - Reservations to hosts.
- Device classification – You can make detailed descriptions of appliances in DHCP ranges and appliances defined as Infoblox hosts and as fixed addresses.
- Three distinct views of IP address usage – To monitor the usage of IP addresses on your network, you can see the following different views:
 - High-level overall network view: From the **Data Management** tab, select the **IPAM** tab -> *member*. You can view the network usage in the Net Map or List view. You can also drill down to specific IP address to get detailed information.
 - DHCP lease history records: From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Lease History**.

Logs

The following are some useful information:

- Logs, as described in [Monitoring the Appliance](#) on page 1003.
 - Audit Log – Contains administrator-initiated events.
 - System Log – Contains events related to hardware and software operations.
- DNS statistics, as described in [Configuring DNS Services](#) on page 555.
 - DNS Configuration – Contains DNS server settings for the Infoblox DNS server.
 - Zone Statistics – Contains the results of all DNS queries per zone.
- DHCP information, as described in [Configuring DHCP Properties](#) on page 791.
 - DHCP Configuration – Contains DHCP server settings and network, DHCP range, and host settings for the Infoblox DHCP server.
 - DHCP Leases – Contains a real-time record of DHCP leases.
 - DHCP Lease History – Contains an historical record of DHCP leases.
 - DHCP Statistics – Contains the number of currently assigned static and dynamic addresses, and the high and low watermarks per network.
 - Network Statistics – Contains the number of static hosts, dynamic hosts, and available hosts per network.

VERIFYING THE DEPLOYMENT

After you deploy a single independent appliance or HA pair, you can make an HTTPS connection to it, log in, and check its status.

Single Independent Appliance

From the Dashboard, check the appliance status in the *System Status* widget. For information, see [Member Status \(System Status\)](#) on page 123.

- If the Status icon is green, the appliance has a network connection and is operating properly.
- If the Status icon is red, there is a problem. To determine what it is, look at the system log file for this appliance by selecting the **Administration** tab -> **Logs** tab -> **Syslog**.

Independent HA Pair

1. Make an HTTPS connection to the VIP address of the HA pair, log in, and check the status of both nodes.
2. From the Dashboard, check the appliance status in the *System Status* widget. For information, see [Member Status \(System Status\)](#) on page 123.
 - If the Status icon is green, both nodes have connectivity with each other and are operating properly.
 - If the Status icon is yellow, the two nodes are in the process of forming an HA pair.
 - If the Status icon is red, the passive node is offline or there is a problem. To determine what it is, look at the system log file by selecting the **Administration** tab -> **Logs** tab -> **Syslog**. You can also gather information from the **System** tab -> **System Manager** tab. For information, refer to the online Help.

INFOBLOX TOOLS FOR MIGRATING BULK DATA

Typically, the next step after cabling a single independent appliance to a network and configuring its network settings—or cabling two independent appliances to a network and configuring them as an HA pair—is to import data from legacy DNS, DHCP, and TFTP servers. Infoblox provides several tools to accomplish this:

- The CSV import feature allows you to import DNS, DHCP, and IPAM data through Grid Manager. You can add, overwrite, or merge data using this feature. The appliance updates the database based on import settings and the data you specify in the data files. From the **Data Management** tab of Grid Manager, you can access the *Import Manager* editor from which you start a data import. You can also export existing data to a CSV file. You can use this file to modify data, and then re-import the data into the database using the CSV import feature. For information, see [About CSV Import](#) on page 86.
- The Infoblox Data Import Wizard is a useful tool that simplifies the importation of DNS, DHCP and IPAM, and TFTP settings and data into a NIOS appliance. For large data sets, this option is an efficient approach. To download the Data Import Wizard, visit www.infoblox.com/import/.
- For smaller DNS data sets, you can use the zone import feature, which allows you to import data on a per-zone basis (see [Importing Zone Data](#) on page 631).



Chapter 7 Managing Appliance Operations

Managing the operations of a NIOS appliance involves defining system parameters such as time, security, and port settings. It also includes configuring operations such as scheduling tasks, defining approval workflows, managing extensible attributes, and configuring access control for supported operations.

The tasks covered in this chapter include:

- [*Configuring Access Control*](#) on page 306
 - [*Administrative Permissions*](#) on page 306
 - [*Operations that Support Access Control*](#) on page 306
 - [*Defining Named ACLs*](#) on page 307
 - [*Managing Named ACLs*](#) on page 309
 - [*Applying Access Control to Operations*](#) on page 311
- [*Managing Time Settings*](#) on page 312
 - [*Changing Time and Date Settings*](#) on page 312
 - [*Changing Time Zone Settings*](#) on page 312
 - [*Monitoring Time Services*](#) on page 313
- [*Using NTP for Time Settings*](#) on page 313
 - [*Authenticating NTP*](#) on page 314
 - [*NIOS Appliance as NTP Client*](#) on page 315
 - [*Configuring a Grid to Use NTP*](#) on page 316
 - [*Configuring Grid Members to Use NTP*](#) on page 318
 - [*NIOS Appliances as NTP Servers*](#) on page 319
 - [*Configuring a NIOS Appliance as an NTP Server*](#) on page 320
 - [*Monitoring NTP*](#) on page 322
- [*About Extensible Attributes*](#) on page 322
 - [*Adding Extensible Attributes*](#) on page 324
 - [*Configuring Inheritable Extensible Attributes*](#) on page 326
 - [*Using Extensible Attributes*](#) on page 332
 - [*Viewing Extensible Attributes*](#) on page 330
 - [*Modifying Extensible Attributes*](#) on page 330
 - [*Deleting Extensible Attributes*](#) on page 331
 - [*Configuration Examples for Inheritable Extensible Attributes*](#) on page 334

- [*Managing Security Operations*](#) on page 343
 - [*Enabling Support Access*](#) on page 343
 - [*Enabling Remote Console Access*](#) on page 343
 - [*Permanently Disabling Remote Console and Support Access*](#) on page 344
 - [*Restricting GUI/API Access*](#) on page 344
 - [*Enabling HTTP Redirection*](#) on page 344
 - [*Modifying the Session Timeout Setting*](#) on page 344
 - [*Disabling the LCD Input Buttons*](#) on page 344
 - [*Configuring Security Features*](#) on page 344
- [*Configuring Ethernet Ports*](#) on page 346
 - [*About Virtual LANs*](#) on page 346
 - [*Implementing Quality of Service Using DSCP*](#) on page 348
 - [*Ethernet Port Usage*](#) on page 349
 - [*Modifying Ethernet Port Settings*](#) on page 354
 - [*Using the LAN2 Port*](#) on page 355
 - [*About Port Redundancy*](#) on page 356
 - [*Configuring the LAN2 Port*](#) on page 357
 - [*Enabling DHCP on LAN2*](#) on page 358
 - [*Enabling DNS on LAN2*](#) on page 358
- [*Using the MGMT Port*](#) on page 359
 - [*Appliance Management*](#) on page 360
 - [*Grid Communications*](#) on page 362
 - [*DNS Services*](#) on page 364
- [*About Lights Out Management*](#) on page 366
 - [*Enabling LOM*](#) on page 367
 - [*Modifying LOM Settings*](#) on page 369
 - [*Viewing LOM Users*](#) on page 369
 - [*IPMI Commands and Examples*](#) on page 369
- [*Setting Static Routes*](#) on page 372
- [*Enabling DNS Resolution*](#) on page 376
- [*Managing Licenses*](#) on page 377
 - [*Obtaining and Adding Licenses*](#) on page 377
 - [*Obtaining Temporary Licenses*](#) on page 378
 - [*Viewing Licenses*](#) on page 378
 - [*Backing Up Licenses*](#) on page 379
 - [*Removing Licenses*](#) on page 379
- [*Managing the Order of Match Lists*](#) on page 380
- [*Shutting Down, Rebooting, and Resetting a NIOS Appliance*](#) on page 380
 - [*Rebooting a NIOS Appliance*](#) on page 380
 - [*Shutting Down a NIOS Appliance*](#) on page 380
 - [*Resetting a NIOS Appliance*](#) on page 381
- [*Managing the Disk Subsystem on the Infoblox-2000-A and -4010*](#) on page 382
 - [*About RAID 10*](#) on page 382
 - [*Evaluating the Status of the Disk Subsystem*](#) on page 383
 - [*Disk Drive Front Panel LEDs*](#) on page 383

-
- [*Replacing a Failed Disk Drive*](#) on page 384
 - [*Disk Array Guidelines*](#) on page 385
 - [*Restarting Services*](#) on page 386
 - [*Canceling a Scheduled Restart*](#) on page 388

CONFIGURING ACCESS CONTROL

To effectively manage your core network services, you can grant legitimate hosts access to specific tasks and operations using an access control list (ACL) or anonymous access control entries (ACEs). Depending on your admin permissions, you can configure a named ACL, and then apply it to multiple operations, such as file distribution and DNS zone transfers. For information about admin permissions, see [About Administrative Permissions](#) on page 160.

When you define a named ACL, you add access control types such as IPv4 and IPv6 addresses, IPv4 and IPv6 networks, nested named ACLs, and TSIG key based ACEs to a list, and then grant each entry in the list the Allow or Deny permission. For information about named ACLs and how to configure them, see [Defining Named ACLs](#) on page 307. Note that each operation supports specific access control types. You cannot apply a named ACL to an operation that does not support the access control types contained in the named ACL. For more information about which NIOS operations support access control and which access control types each operation supports, see [Operations that Support Access Control](#) on page 306.

When you add or modify a named ACL, or when you import named ACLs and ACEs to an existing named ACL through CSV import, the appliance does not automatically validate the ACEs in the list. For more information about how to import named ACLs and ACEs, refer to the *Infoblox CSV Import Reference*. To avoid conflicts and unexpected results, you must perform validations for all named ACLs before you use them for access control. When the appliance detects a conflict or an optimized issue about a specific ACE during the validation process, it displays detailed information in a CSV file. For more information about ACL validation, see [Validating Named ACLs](#) on page 309.

Administrative Permissions

You can configure a named ACL at the Grid level and override it at the member and object level. Superusers and limited-access users with Read/Write permission to **All Named ACLs** can create, modify, and delete named ACLs. Users with Read-only permission to **All Named ACLs** can apply a named ACL to a supported object if they have Read/Write permission to the respective object. Other users can only view named ACLs and their entries. For information about admin permissions, see [About Administrative Permissions](#) on page 160.

Operations that Support Access Control

On the appliance, only certain operations support access control. You can apply one named ACL or multiple anonymous ACEs to each operation. However, you cannot apply multiple named ACLs or use a combination of named ACLs and ACEs. Note that each operation supports different access control types. For example, DNS zone transfers support IPv4 and IPv6 addresses and networks as well as TSIG key based ACEs, while AAAA filtering supports only IPv4 addresses and networks.

When you apply a named ACL to an operation, the appliance validates to ensure that the named ACL contains ACEs that are supported by the operation. The appliance also validates any new ACEs that you add to an existing named ACL. If a named ACL contains access control types that an operation does not support, the appliance displays an error message and you cannot apply that named ACL to the operation. Thus when defining a named ACL for a specific operation or applying an existing named ACL, ensure that it contains access control types that the operation supports. [Table 7.1](#) lists access control types for NIOS operations that support access control.

Table 7.1 Operations that Support Access Control

Operation Type	Supported Access Control Types				
	IPv4 Addresses and Networks	IPv6 Addresses and Networks	TSIG Key Based ACEs	DNSone 2.x TSIG Key	Any Address and Network
GUI and API Access	Yes	Yes	No	No	No
NTP Service and NTP Queries	Yes	No	No	No	Yes
File Distribution Services	Yes	No	No	No	No
Syslog Proxy Access Control	Yes	Yes	No	No	No
DNS Zone Transfers (excludes zone transfers for Microsoft servers)*	Yes	Yes	Yes	Yes	Yes
Dynamic DNS Updates	Yes	Yes	Yes	No	Yes
DNS Queries	Yes	Yes	Yes	No	Yes
Recursive Queries	Yes	Yes	Yes	No	Yes
Blackhole Lists	Yes	Yes	No	No	Yes
AAAA Filtering	Yes	No	No	No	Yes
Forward DNS Updates	Yes	Yes	Yes	No	Yes
Match Clients for DNS Views	Yes	Yes	Yes	Yes	Yes
Match Destinations for DNS Views	Yes	Yes	Yes	Yes	Yes
DNS64 Clients	Yes	Yes	No	No	Yes
DNS64 Mapped	Yes	No	No	No	Yes
DNS64 Exclude IPv6	No	Yes	No	No	Yes

Note: * Zone transfers for Microsoft servers do not support named ACLs. However, you can still use individual ACEs to configure access control. For more information about how to configure zone transfer settings for Microsoft servers, see [Setting Zone Properties](#) on page 971. In addition, the DNSone 2.x TSIG key supports only the “Allow” permission. You cannot change “Allow” to “Deny.”

Complete the following tasks to use a named ACL:

1. Define a named ACL, as described in [Defining Named ACLs](#) on page 307.
2. Validate the named ACL, as described in [Validating Named ACLs](#) on page 309.
3. Apply the named ACL to specific operations, as described in [Applying Access Control to Operations](#) on page 311.

Defining Named ACLs

Depending on how you plan to use a named ACL and which access control types an operation supports, you can add one or all of the following when you define a named ACL: IPv4 and IPv6 addresses, IPv4 and IPv6 networks, TSIG key based ACEs, DNSone 2.x TSIG keys. You can also add an existing named ACL as a nested ACL to a new or existing named ACL.

When configuring a named ACL, ensure that you define it correctly for the intended operations using the supported access control types. For example, if you want to apply a named ACL to AAAA filtering, do not include IPv6 addresses or networks in the named ACL because AAAA filtering does not support IPv6 addresses and networks. For information about supported access control types, see [Table 7.1](#).

To define a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab, and then click the Add icon.
2. In the *Add Named ACL* wizard, complete the following:
 - **Name:** Enter a name for the named ACL. You can enter up to 64 characters.
 - **Comment:** Enter additional information about the named ACL.
3. Click **Next**. Complete the following to add ACEs to the named ACL:
 - Click the Add icon and select one of the following access control types from the drop-down list. Depending on your selection, Grid Manager adds a row to the table directly or expands the panel before adding a row.
 - **IPv4 Address:** Select this to add an IPv4 address. Click the **Entry** field and enter the IPv4 address. The **Operation** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** When you select this, enter the network address in the **Address** field, select the netmask using the slider, and then select **Allow** or **Deny** from the **Permission** drop-down list. Click **Add** and Grid Manager adds the entry to the table.
 - **IPv6 Address:** Select this to add an IPv6 address. Click the **Entry** field and enter the IPv6 address. The **Operation** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv6 Network:** When you select this, enter the network address and its netmask in the **Address** field, and then select **Allow** or **Deny** from the **Permission** drop-down list. Click **Add** and Grid Manager adds the entry to the table.
 - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server. This name must match the name of the same TSIG key on other name servers.
 - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
 - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
 - **DNSone 2.x TSIG Key:** Select this when the client is a NIOS appliance running DNS One 2.x code. The appliance automatically populates the value of the key in the **Entry** field. The **Operation** column displays **Allow** by default. You cannot change the default permission.
 - **Any Address/Network:** Select this to allow or deny permission for any addresses and networks.
 - **Named ACL:** When you select this, Grid Manager displays the *Named ACLs* Selector. Select the named ACLs you want to add to the new ACL. If you have only one existing named ACL, Grid Manager automatically adds the named ACL to the list. The selected named ACL becomes a nested ACL in the newly created named ACL.

Note: The **Order** field in the table displays the position of each entry based on the order it is placed in the list. You can modify this number to change the order of an ACE. You can also select the ACE check box and use the up and down arrows next to the table to place the entry in the desired position.

4. Click **Next** to enter extensible attributes for the named ACL. For information, see [About Extensible Attributes](#) on page 322.
5. Save the configuration.

Managing Named ACLs

You can do the following after you have configured named ACLs for access control:

- Preview the list of ACEs in a named ACL, as described in [Previewing ACEs in Named ACLs](#) on page 309
- Validate ACEs in a named ACL, as described in [Validating Named ACLs](#) on page 309.
- View a complete list of configured named ACLs, as described in [Viewing Named ACLs](#) on page 310.
- Modify information in a named ACL, as described in [Modifying Named ACLs](#) on page 311.
- Apply a named ACL to supported operations, as described in [Applying Access Control to Operations](#) on page 311.
- Delete a named ACL, as described in [Deleting Named ACLs](#) on page 311.
- Export and print the list of named ACLs.

Previewing ACEs in Named ACLs

You can preview the list of ACEs in a named ACL when you add or modify it. When you click the Preview icon in the *Add Named ACL* wizard or *Named ACL* editor, the appliance lists all the entries in the named ACL, even if you have selected only one or a few entries in the wizard or editor.

To preview a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named_acl* check box, and then click the Preview icon.
2. In a separate browser window, Grid Manager displays the following information for each ACE in the named ACL:
 - **Entry:** Displays one of the following: IPv4 or IPv6 address, IPv4 or IPv6 network, or TSIG key. Note that if the named ACL contains nested ACLs, all entries in the nested ACLs are displayed in a flat view. Grid Manager does not display the name of the nested ACL.
 - **Type:** The access control type of the entry. This can be **IPv4 Address**, **IPv6 Address**, **IPv4 Network**, **IPv6 Network**, **TSIG Key**, or **DNSone 2.x TSIG Key**.
 - **Operation:** Displays the access permission for the entry. This can be **Allow** or **Deny**.

Validating Named ACLs

When you add or modify a named ACL, the appliance does not automatically validate the ACEs in the list. In addition, when you import named ACLs or ACEs to a named ACL, no automatic validation is performed. To avoid unintended consequences, ensure that you validate your named ACLs before you save them or use them for access control.

Note: When you click the Validate icon in the *Add Named ACL* wizard or *Named ACL* editor, the appliance validates all the entries in the named ACL, even if you have selected only one or a few entries in the wizard or editor.

The following examples demonstrate the importance of validating named ACLs:

Example 1

You configure a named ACL “foo” that includes a Deny permission to 10.0.0.0/16. You then assign “foo” to DNS zone transfers. You later import an “Allow/10.0.0.0/24” entry to “foo” through CSV import. The appliance appends the entry to the end of “foo.” When you perform an ACL validation on “foo” after a DNS service restart, the appliance displays a warning message indicating that the new “Allow/10.0.0.0/24” entry is now included in the previously configured “Deny/10.0.0.0/16” entry. Since DNS service works on a first-match access control basis, zone transfers will not be allowed for the 10.0.0.0/24 network, which is probably not your original intent. You can then modify the named ACL to correct this error. On the other hand, if you do not perform the ACL validation, the appliance is not notified about the new network entry in “foo.” As a result, you are not notified about the denial of zone transfers to 10.0.0.0/24.

Example 2

You add a nested named ACL “bar” as the first entry to the named ACL “foo.” You then add a “Deny All” entry right after “bar” (as the second entry). You later import a new “Allow All” entry to “bar” through CSV import. The “Allow All” entry will be appended to the end of “bar.” When you perform an ACL validation on “foo” after the CSV import, the appliance detects a conflict between the “Allow All” (in “bar”) and “Deny All” (right after “bar”) permissions and displays a warning. Imagine if you do not perform the ACL validation on “foo,” the appliance is not notified about the new “Allow All” entry in “bar” and therefore cannot detect the conflict between the “Allow All” and “Deny All” entries. As a result, almost all hosts will get zone transfers, which may not be the outcome you have intended.

Note: It is important that you manually validate each named ACL after a CSV import to ensure data and performance integrity. The appliance does not automatically perform ACL validation.

To validate a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named_acl*/check box, and then click the Validate icon.
or
In the *Add Named ACL* wizard or *Named ACL* editor, click the Validate icon.
2. Grid Manager validates all the ACEs in the named ACL and displays a system message at the top of the screen indicating whether all ACEs in the named ACL are valid or not, depending on the validation results. When the appliance detects conflicts or issues related to specific ACEs, it displays the results in a CSV file. You can save the file or open it. Grid Manager displays the following information in the file:
 - **Defined ACL:** The name of the named ACL.
 - **Type of Issue:** The type of issue found. This can be one of the following:
 - **Optimize:** An ACE is a duplicate of a previous entry or an ACE configuration can be a subset of another entry. See optimized suggestions in the **Issue** field.
 - **Conflict:** The same IP address or network has a conflicting permission. Re-configure the ACE based on your requirements.
 - **Warning:** An ACE is a subset of a previously configured entry, but it has a conflicting permission.
 - **ACE A:** The ACE that has a conflict or an optimized issue with ACE B.
 - **ACE B:** The ACE that has a conflict or an optimized issue with ACE A.
 - **Issue:** Detailed information and optimized suggestions about the conflict or issue.

Note: It may take a long time to validate a named ACL that contains a large number of ACEs.

Viewing Named ACLs

To view a list of named ACLs:

- From the **Administration** tab, select the **Named ACLs** tab. Grid Manager displays the following information for each named ACL:
 - **Name:** The name of the named ACL.
 - **Comment:** Information about the named ACL.
 - **Site:** The site to which the named ACL belongs. This is one of the predefined extensible attributes.

You can also do the following in the **Named ACLs** tab:

- Modify data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes or **Cancel** to exit.
- Sort the named ACLs in ascending or descending order by column.
- Select a named ACL and click the Edit icon to modify data, or click the Delete icon to delete it.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the Go to field and select the object from the possible matches.

- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print and export data in this tab.

Modifying Named ACLs

You can modify ACEs in an existing named ACL. When you update a named ACL, the appliance validates the updates to ensure that ACEs in the named ACL are valid for the operations to which the name ACL has been applied. For example, if a named ACL is used for file distribution access, you are not allowed to add IPv6 address access control to it because the file transfer operation does not support IPv6 addresses.

To modify a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named_acl* check box, and then click the Edit icon.
2. The *Named ACL* editor provides the following tabs from which you can modify data:
 - **General Basic:** You can modify data in this tab as described in [Defining Named ACLs](#) on page 307.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific named ACL. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.

Deleting Named ACLs

When you delete a named ACL, the appliance puts it in the Recycle Bin, if enabled. You can restore the named ACL later if needed.

Note: You cannot delete a named ACL that has been applied to an operation or is currently in use by another operation.

To delete a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named_acl* check box, and then click the Delete icon. You can select multiple named ACLs for deletion.
2. In the *Delete Confirmation* dialog box, click **Yes**.

Applying Access Control to Operations

When you apply access control to NIOS operations, you can use anonymous ACEs or a named ACL. You cannot combine ACEs and named ACLs for access control. Depending on the access control types each operation supports, you may or may not be able to apply a named ACL to a specific operation. For information about which access control types each operation supports, see [Table 7.1](#) on page 307.

Note: If you disable access control or select **None** or **Any** for an operation, the appliance removes the previously applied named ACL or the configured anonymous ACEs. To avoid losing your ACE configuration, Infoblox recommends that you convert the ACEs to a named ACL.

For information about how to apply access control to each supported operation, see the following:

- DNS zone transfers, as described in [Enabling Zone Transfers](#) on page 583
- DNS queries, as described in [Controlling DNS Queries](#) on page 570
- Recursive queries, as described in [Enabling Recursive Queries](#) on page 571
- Dynamic DNS updates, as described in [Enabling DNS Servers to Accept DDNS Updates](#) on page 706
- AAAA filtering, as described in [Controlling AAAA Records for IPv4 Clients](#) on page 573
- Blackhole list, as described in [Configuring a DNS Blackhole List](#) on page 590
- Match clients list for DNS views, as described in [Defining Match Clients Lists](#) on page 605

- Match destinations for DNS views, as described in [Defining a Match Destinations List](#) on page 607
- DNS64 clients, DNS64 mapped IPv4 addresses, and DNS64 excluded IPv6 addresses, as described in [Setting DNS64 Group Properties](#) on page 597
- File distribution services, as described in [Configuring Access Control for File Distribution](#) on page 394
- Grid Manager and API access, as described in [Configuring Security Features](#) on page 344
- NTP access control, as described in [Defining NTP Access Control](#) on page 320
- Syslog proxy access, as described in [Configuring Syslog for Grid Members](#) on page 1014

MANAGING TIME SETTINGS

You can define the date and time settings for your NIOS appliance using the Infoblox Appliance Startup Wizard. Alternatively, you can set the date and time of the appliance anytime after you first configure it if you did not do so using the startup wizard or if you need to change it if, for example, you move an appliance from a location in one time zone to a location in a different time zone. To set the date and time of the appliance, you can either manually enter the values or configure the appliance to synchronize its time with a public NTP server.

Changing Time and Date Settings

If you do not use the NTP service, you can set the date and time for a Grid.

Note: You cannot manually set the date and time if the NTP service is enabled.

To set the time and date for a Grid using the *Grid Properties* editor:

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
2. In the **General** tab of the *Grid Properties* editor, complete the following:
 - **Date:** Click the calendar icon to select a date or enter the date in YYYY-MM-DD format.
 - **Time:** Click the clock icon to select a time or enter the time in HH:MM:SS format. For afternoon and evening hours, use the integers 13-24.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: Changing the date and time resets the application and terminates the management session.

Changing Time Zone Settings

Whether you enable NTP (Network Time Protocol) or manually configure the date and time, you must always set the time zone manually. You can set the time zone for a Grid, which then applies to all members. If different members are in different time zones, you can choose the time zone that applies to most members at the Grid level, and then override the setting for the remaining members.

Note: Changing the time zone does not reset the application nor does it terminate the management session.

To set the time zone for a Grid or member:




1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **General** tab of the editor, select the appropriate time zone.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Monitoring Time Services

In a Grid, the Grid Master and its members use an internal NTP daemon to synchronize their time. It is not user-configurable and functions regardless of how you set the time on the Grid Master. The *Detailed Status* panel contains an NTP Synchronization icon so you can monitor the internal NTP daemon that runs within a Grid to ensure the time among its members is synchronized.

To display the *Detailed Status* panel, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Detailed Status icon in the table toolbar of the *Members* panel.

The following are descriptions of the NTP status icons in the *Detailed Status* panel:

Icon	Color	Meaning
	Green	The NTP service is running properly.
	Yellow	The appliance is synchronizing its time.
	Red	The NTP service is not running properly. View the corresponding description for additional information.

USING NTP FOR TIME SETTINGS

Note: vNIOs Grid members on Riverbed can be NTP clients only.

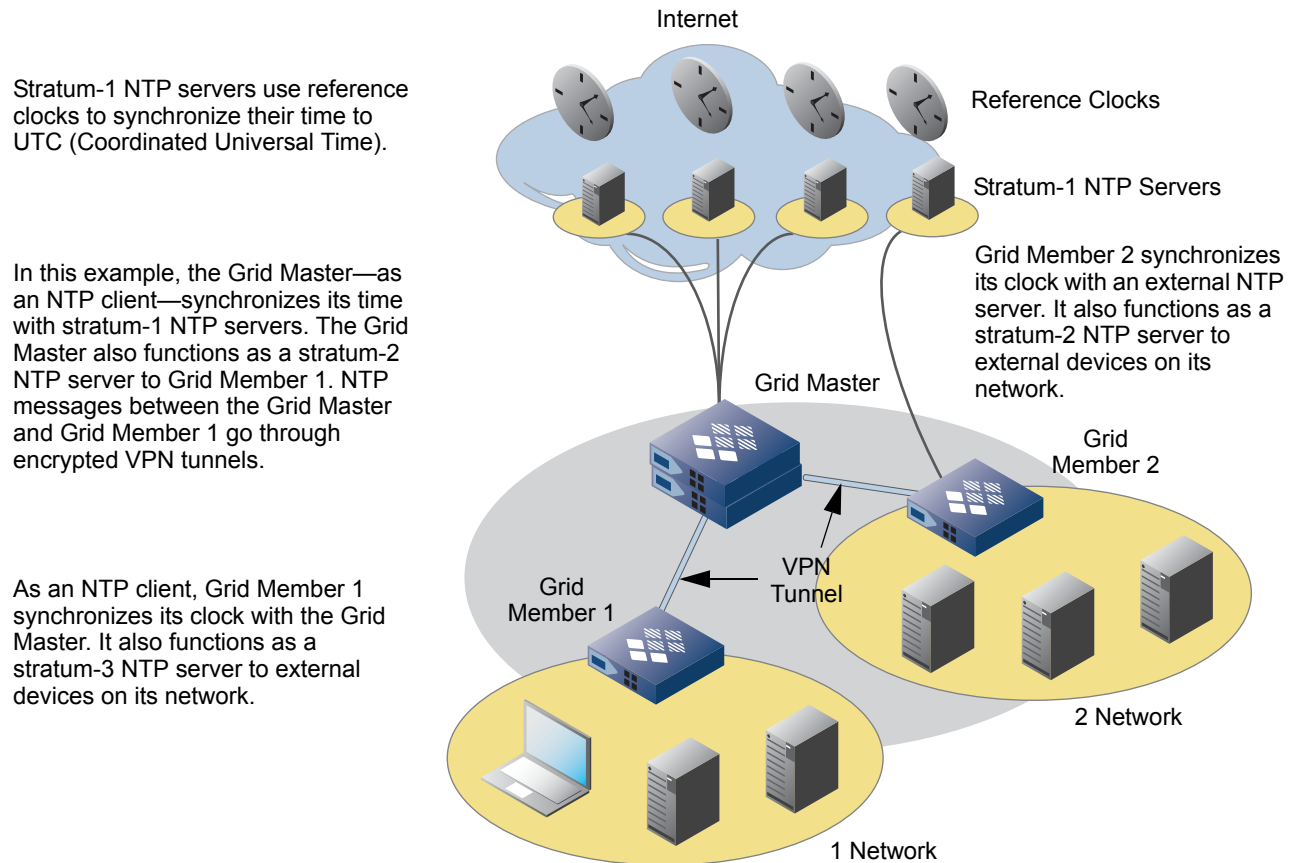
NTP (Network Time Protocol) is a standard protocol that system clocks use to ensure their time is always accurate. Appliances that use NTP try to get their time as close as possible to UTC (Coordinated Universal Time), the standard timescale used worldwide. NTP uses UDP (User Datagram Protocol) on port 123 for communications between clients and servers.

NTP is based on a hierarchy where reference clocks are at the top. Reference clocks use different methods such as special receivers or satellite systems to synchronize their time to UTC. NTP servers on the first level of the hierarchy synchronize their time with the reference clocks, and serve time to clients as well. Each level in the hierarchy is a stratum; stratum-0 is a reference clock. Stratum-1 servers synchronize their clocks with reference clocks. Stratum-2 servers synchronize their clocks with stratum-1 servers, and so forth. The stratum number indicates the number of levels between the NTP server and the reference clock. A higher stratum number could indicate more variance between the NTP server and the reference clock.

You can configure a NIOS appliance to function as an NTP client that synchronizes its clock with an NTP server. NTP clients typically use time information from at least three different sources to ensure reliability and a high degree of accuracy. There are a number of public NTP servers on the Internet with which the NIOS appliance can synchronize its clock. For a list of these servers, you can access <http://www.ntp.org>. When NTP is configured, it listens on all interfaces, including the loopback interface on the NIOS appliance.

In a Grid, the Grid Master and Grid members can function as NTP clients that synchronize their clocks with external NTP servers. They can in turn function as NTP servers to other appliances in the network. Note that when the Grid Master functions as an NTP server, it synchronizes its local clock with its NTP clients and does not synchronize time with any other external NTP server. This allows you to deploy multiple NTP servers to ensure accurate and reliable time across the network. To configure the Grid Master and Grid members as NTP clients, you must first enable the NTP service and configure external NTP servers at the Grid level. You can then configure the Grid Master and Grid members to override the Grid-level NTP servers and use their own external NTP servers. A Grid member synchronizes its clock with the Grid Master if you do not configure it to use external NTP servers.

Figure 7.1 Infoblox Appliances as NTP Servers



Authenticating NTP

To prevent intruders from interfering with the time services on your network, you can authenticate communications between a NIOS appliance and a public NTP server, and between a NIOS appliance and external NTP clients. NTP communications within the Grid go through an encrypted VPN tunnel, so you do not have to enable authentication between members in a Grid.

NTP uses symmetric key cryptography, where the server and the client use the same algorithm and key to calculate and verify a MAC (message authentication code). The MAC is a digital thumbprint of the message that the receiver uses to verify the authenticity of a message.

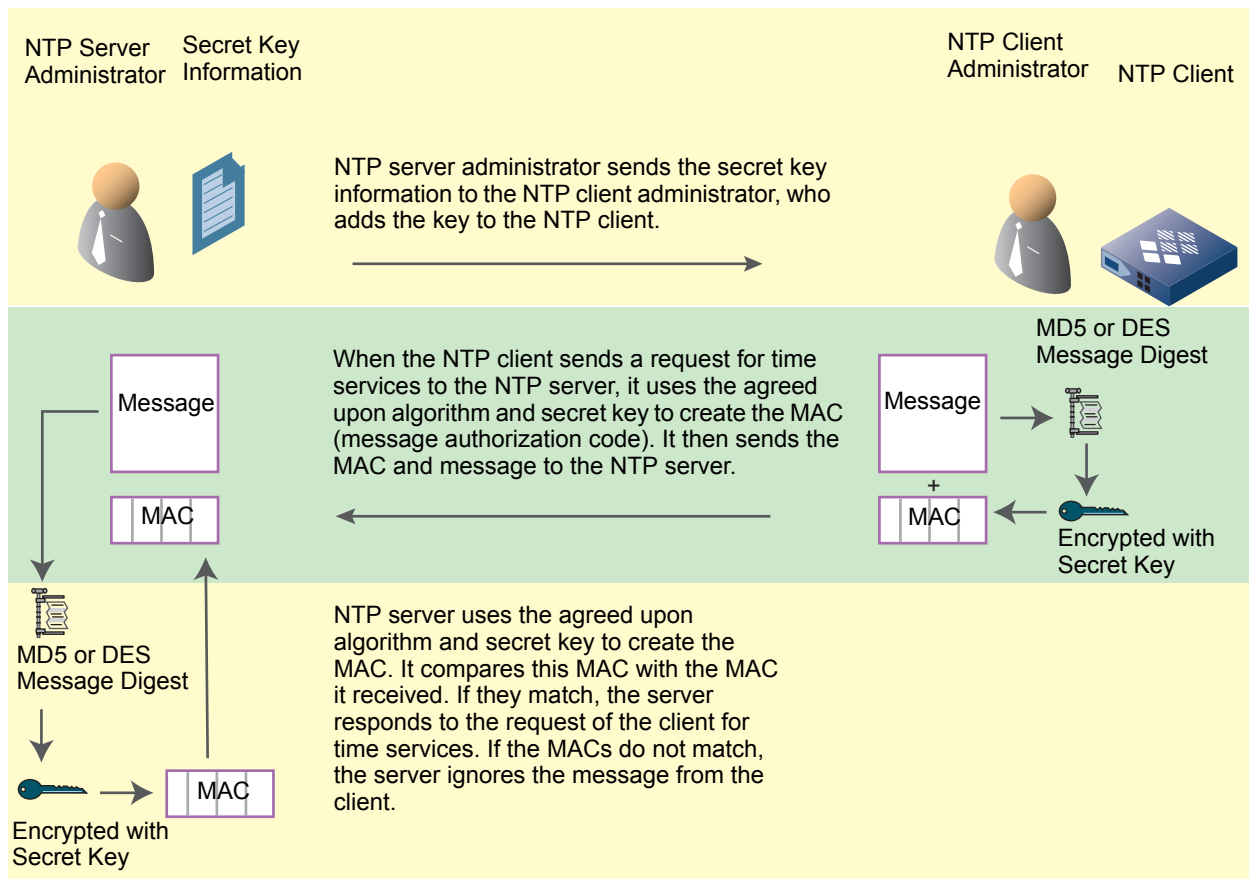
As shown in [Figure 7.2](#), the NTP client administrator must first obtain the secret key information from the administrator of the NTP server. The server and the client must have the same key ID and data. Therefore, when you configure the NIOS appliance as an NTP client and want to use authentication, you must obtain the key information from the administrator of the external NTP server and enter the information on the NIOS appliance. When you configure a NIOS appliance as an NTP server, you must create a key and send the key information to clients in a secure manner. A key consists of the following:

- **Key Number:** A positive integer that identifies the key.
- **Key Type:** Specifies the key format and the algorithm used to calculate the MAC (message authentication code) of a message.
 - **M:** The key is a 1-31 character ASCII string using MD5 (Message Digest).
 - **S:** The key is a 64-bit hexadecimal number in DES (Data Encryption Standard) format. The high order 7 bits of each octet form the 56-bit key, and the low order bit of each octet is given a value so that the octet maintains odd parity. You must specify leading zeros so the key is exactly 16 hexadecimal digits long and maintains odd parity.

- A: The key is a DES key written as a 1-8 character ASCII string.
- N: The key is a 64-bit hexadecimal number in NTP format. It is the same as the S format, but the bits in each octet have been rotated one bit right so the parity bit is in the high order bit of the octet. You must specify leading zeros and odd parity must be maintained.
- **Key String:** The key data used to calculate the MAC. The format depends on the Key Type you select.

When the NTP client initiates a request for time services to the NTP server, it creates the MAC by using the agreed upon algorithm to compress the message and then encrypts the compressed message (which is also called a message digest) with the secret key. The client appends the MAC to the message it sends to the NTP server. When the NTP server receives the message from the client, it performs the same procedure on the message — it compresses the message it received, encrypts it with the secret key and generates the MAC. It then compares the MAC it created with the MAC it received. If they match, the server continues to process and respond to the message. If the MACs do not match, the receiver drops the message.

Figure 7.2 NTP Client Administrator Obtaining Secret Key from NTP Server Administrator



NIOS Appliance as NTP Client

You can configure an independent NIOS appliance, a Grid Master, or any Grid member in a Grid as an NTP client that synchronizes its system clock with an external NTP server.

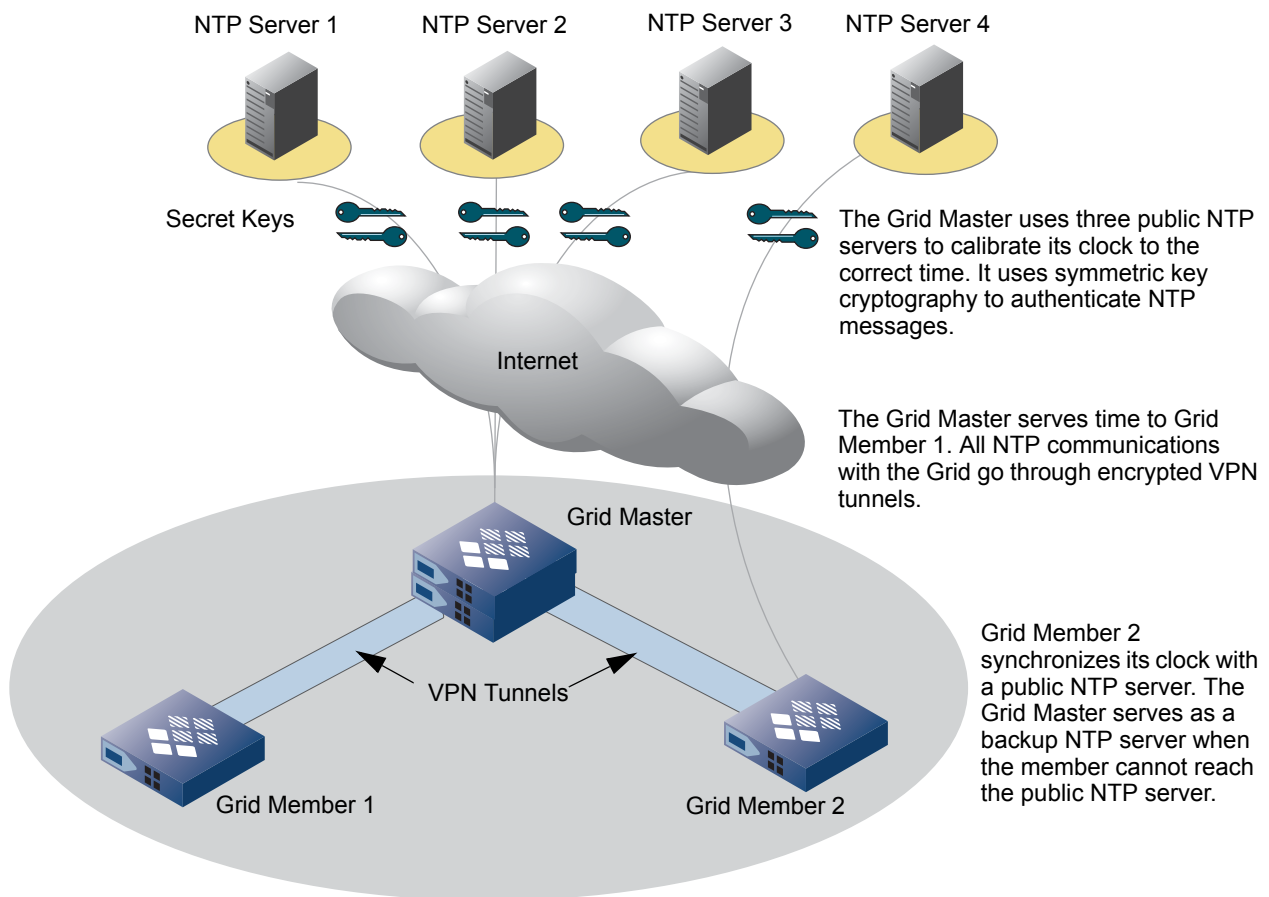
When you enable a NIOS appliance to function as an NTP client, you must specify at least one NTP server with which the appliance can synchronize its clock. Infoblox recommends that you specify multiple NTP servers that synchronize their time with different reference clocks and that have different network paths. This increases stability and reduces risk in case a server fails. For a list of public NTP servers, you can access www.ntp.org.

When you specify multiple NTP servers, the NTP daemon on the appliance determines the best source of time by calculating round-trip time, network delay, and other factors that affect the accuracy of the time. NTP periodically polls the servers and adjusts the time on the appliance until it matches the best source of time. If the difference between the appliance and the server is less than five minutes, the appliance adjusts the time gradually until the clock time matches the NTP server. If the difference in time is more than five minutes, the appliance immediately synchronizes its time to match that of the NTP server.

To secure communications between a NIOS appliance and an NTP server, you can authenticate communications between the appliance and the NTP server. When you configure authentication, you must obtain the key information from the administrator of the NTP server and enter the key on the appliance. For information, see [Authenticating NTP](#) on page 314.

In a Grid, you can configure the Grid Master and Grid members to synchronize their clocks with external NTP servers. When you enable the NTP service on the Grid, the Grid Master automatically functions as an NTP server to the Grid members. A Grid member can synchronize its time with the Grid Master, an external NTP server, or another Grid member. When Grid members synchronize their times with the Grid Master, the Grid Master and its members send NTP messages through an encrypted VPN tunnel, as shown in [Figure 7.3](#). When a Grid member synchronizes its time with another Grid member, the NTP messages are not sent through a VPN tunnel.

Figure 7.3 Grid Master as NTP Client



Configuring a Grid to Use NTP

In a Grid, the Grid Master and Grid members can synchronize their clocks with external NTP servers. They then forward the clock time to other appliances in the network. Likewise, in an independent HA pair, the active node communicates directly with an external NTP server. The passive node then synchronizes its clock with the active node.

In a Grid, you must first enable the NTP service and configure external NTP servers at the Grid level before you configure the Grid Master and Grid members as NTP clients.

To configure a Grid Master as an NTP client, perform the following tasks:

- If you want to enable authentication between the Grid members and NTP servers, you must specify the authentication keys before enabling the NTP service. You can specify authentication keys at the Grid and member levels. For information, see [Adding NTP Authentication Keys](#).
- Enable the NTP service on the Grid and specify one or more external NTP servers. For information, see [Enabling the NTP Service](#) on page 317.

Adding NTP Authentication Keys

To enable authentication between the appliance and the NTP servers, add the authentication keys before enabling the NTP service on the Grid. You can specify authentication keys at the Grid and member levels.

To add NTP authentication keys:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP -> NTP Grid Config**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box. Expand the Toolbar and click **NTP -> NTP Member Config**.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Click the Add icon in the NTP Keys section and enter the following information.
 - **Key Number:** A positive integer that identifies a key.
 - **Type:** Specifies the key format and the algorithm used to calculate the MAC (message authentication code) of a message.
 - **MD5 in ASCII format (M):** The key is a 1-31 character ASCII string using MD5 (Message Digest).
 - **DES in hex format (S):** The key is a 64-bit hexadecimal number in DES (Data Encryption Standard) format. The high order 7 bits of each octet form the 56-bit key, and the low order bit of each octet is given a value so that the octet maintains odd parity. You must specify leading zeros so the key is exactly 16 hexadecimal digits long and maintains odd parity.
 - **DES in ASCII format (A):** The key is a DES key written as a 1-8 character ASCII string.
 - **DES in NTP format (N):** The key is a 64-bit hexadecimal number in NTP format. It is the same as the S format, but the bits in each octet have been rotated one bit right so the parity bit is in the high order bit of the octet. You must specify leading zeros and odd parity must be maintained.
 - **String:** The key data used to calculate the MAC. The format depends on the Key Type you select.
3. Click **Save** to save the entry and keep the editor open so you can enable the Grid to synchronize its time with external NTP servers, as described in [Enabling the NTP Service](#).

Note that if you enter a new key, the appliance checks if the key already exists in the key list. If the key exists, but either the key type or key string does not match, the NIOS appliance sends an error message.

After you enter an authentication key, you can modify or delete it. Note that you cannot delete a key that an NTP server references. You must first delete all NTP servers that reference that key and then delete the key.

Enabling the NTP Service

To enable the Grid to synchronize its time with external NTP servers:

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP -> NTP Grid Config**.
2. In the *Grid NTP Properties* editor, select **Synchronize the Grid with NTP Servers**.
3. Click the Add icon in the External NTP Servers table.
4. In the *Add NTP Server* dialog box, enter the following information, and then click **Add**.
 - **NTP Server:** Enter either the IP address or the resolvable host name of an NTP server. You can view a list of public NTP servers at ntp.isc.org. To check whether the DNS server can resolve the NTP server host name, click **Resolve Name**. You must have a DNS name resolver configured. For information, see [Enabling DNS Resolution](#) on page 376.

- **Enable Authentication:** Select this option to enable authentication of NTP communications between the external NTP server and the NIOS appliance (the Grid Master or Grid member in a Grid, an independent NIOS appliance, or the active node in an independent HA pair).

Note: To prevent intruders from interfering with the time services on your network, you can authenticate communications between a Grid member and an external NTP server, as well as between a Grid member and external NTP clients. NTP communications within the Grid go through an encrypted VPN tunnel, so you do not have to enable authentication between the Grid Master and Grid members.

- **Authentication Key:** Select a key that you previously entered, and then click **OK**. For information, see [Adding NTP Authentication Keys](#) on page 317.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring Grid Members to Use NTP

To configure Grid members to synchronize their time with external NTP servers:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box.
2. Expand the Toolbar and click **NTP** -> **NTP Member Config**.
3. In the *Member NTP Configuration* editor, do the following:
 - **Synchronize this Member with other NTP Servers:** Select this option to enable this Grid member to use external NTP servers. When you select this check box, you must enter at least one external NTP server for the member.
 - **Exclude Grid Master as NTP Server:** Select this option if you want to exclude the Grid Master from being one of the time sources. By default, the appliance automatically configures the Grid Master as the backup NTP server for a Grid member. When the member cannot reach any of its configured NTP servers, it uses the Grid Master as the NTP server.
4. Click **Override**, and then click the Add icon in the External NTP Servers table.
5. In the *Add NTP Server* dialog box, enter the following information, and then click **Add**.
 - **NTP Server:** Enter either the IP address or the resolvable host name of an NTP server. You can view a list of public NTP servers at ntp.isc.org. To check whether the DNS server can resolve the NTP server host name, click **Resolve Name**. You must have a DNS name resolver configured. For information, see [Enabling DNS Resolution](#) on page 376.
 - **Enable Authentication:** Select this check box to enable authentication of NTP communications between the external NTP server and the NIOS appliance (the Grid Master or Grid member in a Grid, an independent NIOS appliance, or the active node in an independent HA pair).

Note: To prevent intruders from interfering with the time services on your network, you can authenticate communications between a Grid member and an external NTP server, as well as between a Grid member and external NTP clients. NTP communications within the Grid go through an encrypted VPN tunnel, so you do not have to enable authentication between the Grid Master and Grid members.

- **Authentication Key:** Select a key that you previously entered, and then click **OK**. For information, see [Adding NTP Authentication Keys](#) on page 317.

6. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing External NTP Servers

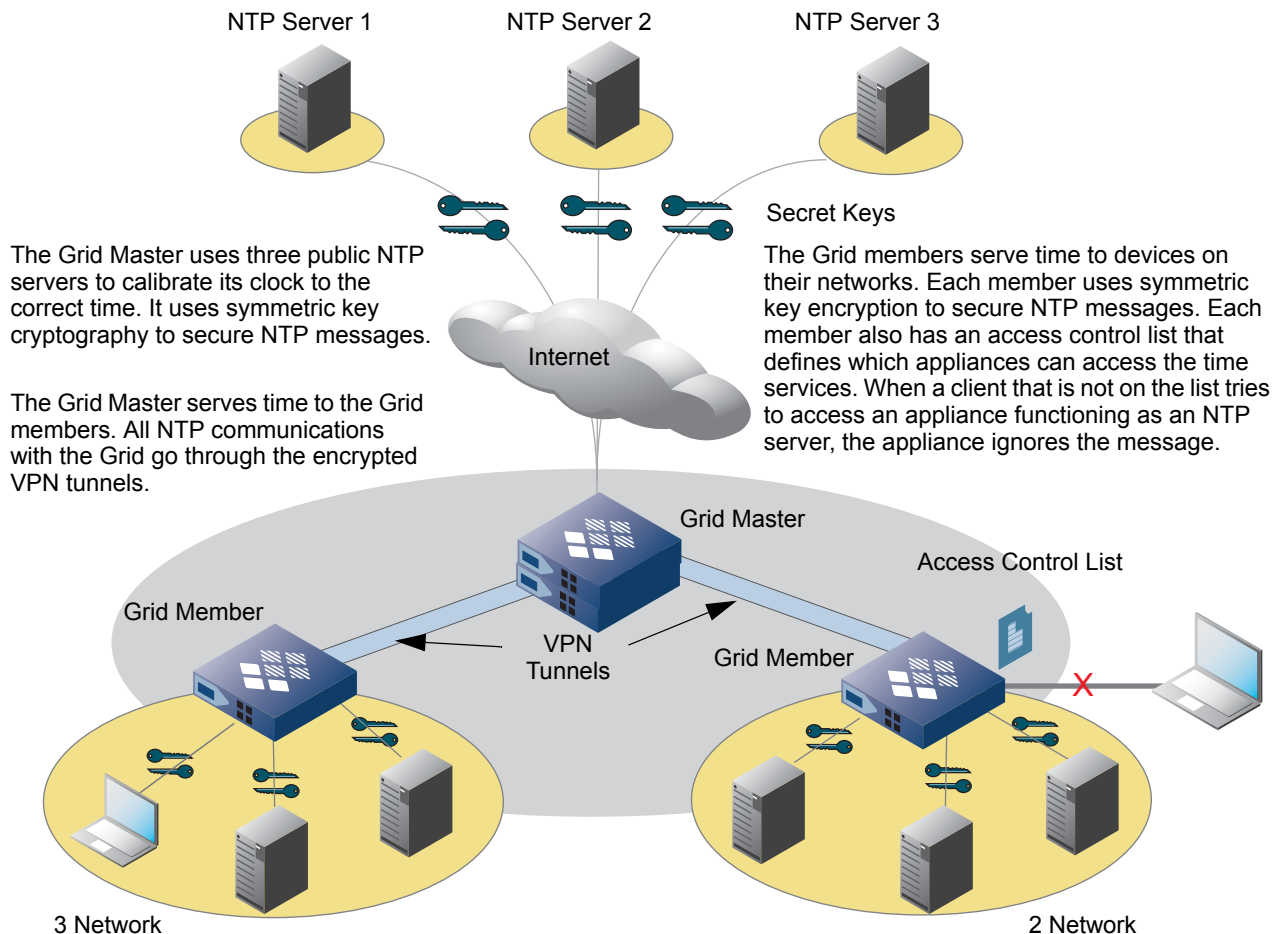
You can specify multiple NTP servers for failover purposes. The NIOS appliance attempts to connect to the NTP servers in the order they are listed. A Grid member uses the Grid Master as the NTP server when it cannot reach any of its external NTP servers.

You can change the order of the list by selecting an NTP server and dragging it to its new location or by clicking the up and down arrows. You can add and delete servers and modify their information as well.

NIOS Appliances as NTP Servers

After you enable NTP on a Grid, the Grid members—including the Grid Master—can function as NTP servers to clients in different segments of the network. Similarly, after you enable NTP on an independent appliance or an HA pair, and it synchronizes its time with an NTP server, you can configure it to function as an NTP server as well. When you configure DNS anycast addressing on a Grid member and use it as an NTP server, the member can answer NTP requests from other NTP clients through the anycast IP address.

Figure 7.4 Grid Members as NTP Servers



To configure a NIOS appliance as an NTP server, perform the following tasks:

- Enable the appliance as an NTP server.
- Enable authentication between the appliance and its NTP clients.
- Optionally, specify which clients can access the NTP service of the appliance.
- Optionally, specify which clients can use `ntpq` to query the appliance.

Configuring a NIOS Appliance as an NTP Server

You can configure a Grid member—including the Grid Master—or an independent appliance or HA pair to function as an NTP server. When you enable a NIOS appliance to function as an NTP server, you can enable authentication between a NIOS appliance functioning as an NTP server and its NTP clients. When you enable authentication, you must specify the keys that the appliance and its clients must use for authentication. In a Grid, you can enter NTP authentication keys at the Grid level so that all the members can use them to authenticate their clients. You can also enter keys at the member level, if you want that member to use different keys from those set at the Grid level. After you enter the keys, you can download the key file and distribute the file to the NTP clients.

To enable an appliance as an NTP server and authenticate NTP traffic between a NIOS appliance and an NTP client, perform the following tasks:

- Enable an appliance as an NTP server and define authentication keys. For information, see [Enabling an Appliance as an NTP Server](#) on page 320.
- Optionally, define NTP access control. For information, see [Defining NTP Access Control](#) on page 320.
- Optionally, configure anycast addressing for DNS and use the anycast IP address for NTP requests. For information about how to configure DNS anycast, see [Configuring Anycast Addresses](#) on page 762.

Enabling an Appliance as an NTP Server

To enable an appliance as an NTP server and add authentication keys:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box.
2. Expand the Toolbar and click **NTP** -> **NTP Member Config**.
3. In the *Member NTP Properties* editor, do the following:
 - **Enable this Member as an NTP Server:** Select this option to configure a Grid Master or a Grid member as an NTP server. If you have configured DNS anycast on the appliance, it can answer NTP requests through the anycast IP address.
 - Click **Override** in the NTP Keys section to enter NTP authentication keys at the member level. The member uses these keys when acting as an NTP server and authenticates requests from NTP clients. Clear the check box to use the Grid-level authentication keys.
4. Click **Add** in the NTP Keys section. For information, see [Adding NTP Authentication Keys](#) on page 317.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

After you enter the authentication keys, you can download the key file (usually called *ntp.keys*) and distribute it to NTP clients as follows:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box.
2. Expand the Toolbar and click **NTP** -> **Download NTP Keys**.
3. In the *Opening ntp.keys* dialog box, save the file, and then click **OK**.
4. Distribute this to the NTP clients using a secure transport.

Defining NTP Access Control

The NTP access control list specifies which clients can use a NIOS appliance as an NTP server. If you do not configure access control, then the NIOS appliance allows access to all clients. You can configure access control at the NTP Grid level and override that at the member level.

In addition, the NIOS appliance can accept queries from clients using *ntpq*, the standard utility program used to query NTP servers about their status and operational parameters. You can specify from which clients the NIOS appliance is allowed to accept *ntpq* queries. The appliance does not accept *ntpq* queries from any client, by default.

You can use an existing named ACL (access control list) or multiple ACEs (access control entries) to control which clients can use the NIOS appliance as an NTP server, as well as those clients from which the appliance can accept queries using *ntpq*. For information about access control, see [Configuring Access Control](#) on page 306.

To specify which clients can access the NTP service of a NIOS appliance and from which clients a NIOS appliance can accept ntpq queries:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP -> NTP Grid Config**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box. Expand the Toolbar and click **NTP -> NTP Member Config**.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **Access Control** tab of the *Grid* or *Member NTP Properties* editor, select one of the following to configure NTP access control:
 - **None:** Select this if you do not want to configure access control for NTP service. When you select **None**, the appliance allows all clients to access the NTP service. This is selected by default.
 - **Use Named ACL for Time only:** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. NTP queries do not support IPv6 addresses/networks and TSIG key based ACEs. When you select this, the appliance allows clients that have the **Allow** permission in the named ACL to use its NTP service. You can click **Clear** to remove the selected named ACL. The appliance accepts ntpq queries from specified NTP clients.
 - **Use Named ACL for Time + NTP Control (NTPQ):** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. NTP queries do not support IPv6 addresses/networks and TSIG key based ACEs. When you select this, the appliance allows clients that have the **Allow** permission in the named ACL to use its NTP service, and for the appliance to accept ntpq queries from those clients as well. You can click **Clear** to remove the selected named ACL.
 - **Use this set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:
 - **IPv4 Address:** Select this to add an IPv4 address. Click the **Value** field and enter the IPv4 address. The default permission is **Allow**, which means that the appliance allows access to and from this IPv4 client. You cannot change the default permission. In the **Service** field, select **Time only** to allow this client for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from this client.
 - **IPv4 Network:** Select this to add an IPv4 network. Click the **Value** field and enter the IPv4 network. The default permission is **Allow**, which means that the appliance allows access to and from this IPv4 network. You cannot change the default permission. In the **Service** field, select **Time only** to allow this network for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from this network.
 - **Any Address/Network:** Select this to allow access to all IPv4 addresses and networks. The default permission is **Allow**, which means that the appliance allows access to and from all IPv4 clients. You cannot change the default permission. In the **Service** field, select **Time only** to allow clients for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from all clients.





After you have added access control entries, you can do the following:

 - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Monitoring NTP

When you enable the Grid to synchronize its time with external NTP servers, you can monitor the status of the NTP service by checking the NTP status icons in the *Member Services* panel. To access the panel, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then select the Manage Member Services icon in the table toolbar of the Members tab.

The following are descriptions of the NTP status icons in the *Members Services* panel. The type of information that can appear in the **Description** column corresponds to the SNMP trap messages. For information about the Infoblox SNMP traps, see [Chapter 37, Monitoring with SNMP](#), on page 1037.

Icon	Color	Meaning
	Green	The NTP service is enabled and running properly.
	Yellow	The NTP service is enabled, and the appliance is synchronizing its time.
	Red	The NTP service is enabled, but it is not running properly or is out of synchronization.
	Gray	The NTP service is disabled.

After you upgrade the Grid to 6.6.x or later, the color of the Grid status icon changes based on the following:

- If you activate an external synchronization, or start the NTP service using the Grid Manager, or do not configure any external NTP servers, except local, then the NTP behavior remains the same and the NIOS appliance displays the Grid status icon in green.
- If you activate an external synchronization and configure one or more external NTP servers, or if the servers are in synchronization with the Grid Master, then the Grid status icon is as follows:
 - Green: NTP is synchronizing with an external server.
 - Red: NTP is synchronizing with the local server and none of the external NTP servers are reachable. This status icon also indicates if there are problems with the NTP service.
 - Yellow: NTP is synchronizing with the local server and at least one external NTP server is reachable; but there could be problems on the external server, such as an exceeded root distance error.

ABOUT EXTENSIBLE ATTRIBUTES

Extensible attributes are identifiers that you use to further define and track a NIOS object. For example, to specify the location of a network, you can add the predefined attribute **Site** and enter a specific location for the network. You can also specify whether an extensible attribute is required for an object or restrict the values that can be entered when you create a new object.

You can also specify if an extensible attribute is inheritable by other objects in an inheritance chain. When you enable the inheritance of an extensible attribute, all descendants in the inheritance chain can inherit the extensible attribute so you do not have to configure it at all object levels. For example, if you define an extensible attribute for a network, the attribute and its value can be automatically added for DHCP ranges and fixed addresses in the network.

An extensible attribute is inheritable by descendants in an inheritance chain if its definition does not restrict it to objects that are not part of an inheritance chain. The appliance supports this inheritance chain: Network View -> Network Container -> Network -> Range -> Host/Fixed Address/Reservation. A parent object can have descendants at one or more levels. For example, a network view, network container, network, or DHCP range can be a parent object and have descendants at one or more levels, while a host, fixed address, and reservation can only be a descendant, not a parent. You can set an extensible attribute to be inheritable by selecting the **Enable Inheritance** option when you define an attribute. For more information, see [Configuring Inheritable Extensible Attributes](#) on page 326.

Note: Only superusers can configure extensible attributes.

You can use predefined extensible attributes or specify new ones for different objects. The appliance provides the following predefined extensible attributes that you can customize:

- Region
- Country
- State
- Site
- Building
- VLAN

When you use a predefined attribute, you can modify it and change its name, but you cannot change the type of data it accepts. You can also delete predefined attributes that you do not use. All predefined attributes accept text strings. You can define other settings though, as described in [Modifying Extensible Attributes](#) on page 330. You can also create your own extensible attributes, as described in [Adding Extensible Attributes](#) on page 324.

For example, you can configure the predefined attribute **Site** for fixed addresses and hosts, and define a new attribute **Department** for admin groups.

When you configure an extensible attribute, you can specify the following:

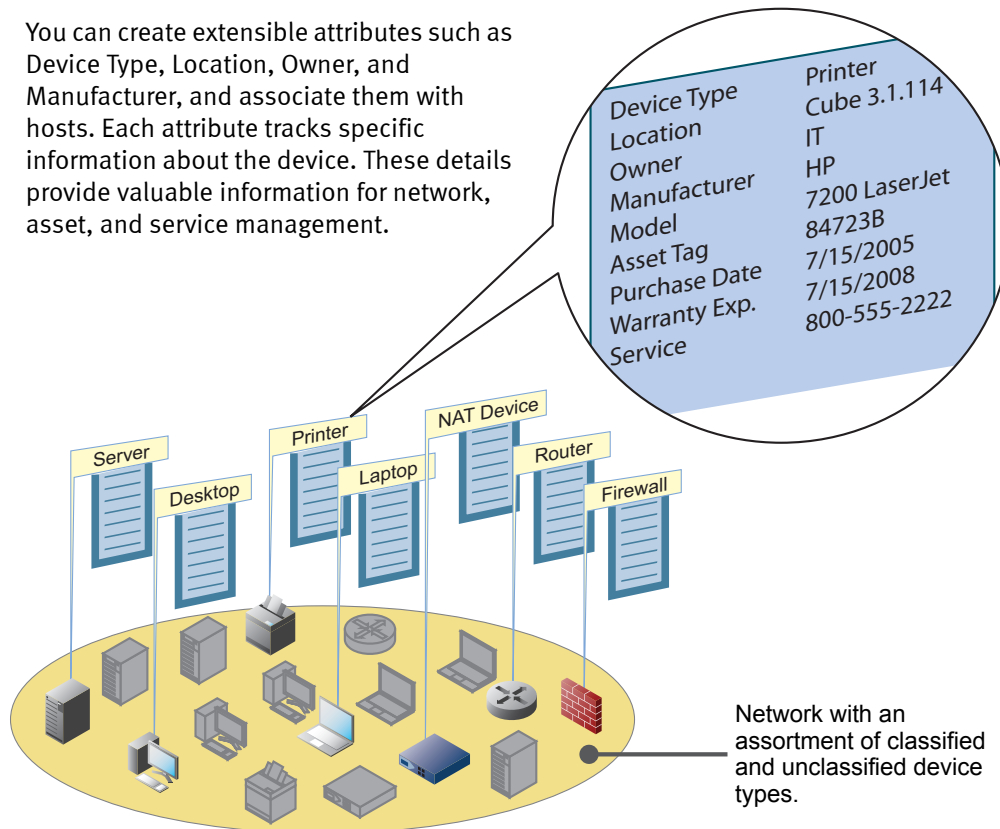
- The type of data that admins enter, such as text strings, integers, or email addresses. You can also restrict admins to a list of values.
- Whether admins can enter multiple values
- A default value
- Whether the attribute is required
- Whether the attribute is inheritable
- The objects associated with the attribute, such as admin groups, DNS views, or DHCP networks.
- Whether the appliance makes an entry in the audit log each time an object with the attribute is added or modified.

Activities such as additions, modifications, and deletions of inheritable extensible attributes, are recorded in the audit log. For information about how to use the audit log, see [Using the Audit Log](#) on page 1018.

[Figure 7.5](#) illustrates a network with different device types. Each device is represented as a host in the NIOS appliance database. You can configure **Device Type**, **Location** and **Owner** as required attributes for hosts. Then when admins add hosts, they will be required to enter values for these attributes in the **Extensible Attributes** tab of the *Add Host* wizard.

Figure 7.5 Using Extensible Attributes to Define Network Devices

You can create extensible attributes such as Device Type, Location, Owner, and Manufacturer, and associate them with hosts. Each attribute tracks specific information about the device. These details provide valuable information for network, asset, and service management.



After you configure extensible attributes for an object, the attributes become available in the **Extensible Attributes** tab of the wizard and editor of the corresponding object. Users then add or edit the attribute values, based on your configuration. Users can also specify attributes when searching for data and add attributes as columns in the tables of Grid Manager. For example, you can add the predefined **Site** attribute as a column in the Records panel of the **Zones** tab. For information about adding columns to tables, see [Customizing Tables](#) on page 60.

Users can also group objects in smart folders according to their attributes. For example, a user can create a smart folder that contains all networks in a certain site.

Adding Extensible Attributes

To add a new extensible attribute:

1. From the **Administration** tab, select the **Extensible Attributes** tab.
2. Click the Add icon on any of the toolbars.
3. In the *Add Extensible Attribute* wizard, complete the following:
 - **Name:** Enter the name of the attribute. This is a required field and is case-sensitive. You can enter up to 128 UTF-8 characters.
 - **Type:** Specify the type of data that you want to capture for an object. Select one of the following:
 - **String:** Select this when the attribute is used to define string values, such as names. When you select this type, the wizard displays the **Number of Characters** field where you can enter the minimum and maximum number of characters that users can enter.
 - **List:** Select this when you want to define a list of values for the attribute. Users can then select a value from this list. For example, if you want to restrict an attribute to five specific values, you can define the attribute as a **List** and then list the five values in the *List Values* section. When a user uses the attribute, they are limited to selecting from one of the five values.

When you select **List**, the wizard displays the List of values table, where you add the allowed values. These values appear in the drop-down list when a user defines the attribute. Click the Add icon to enter values in the table. You can enter up to 64 UTF-8 characters for each value.

You can also modify list values at a later time. When you modify list values, all object attributes using the modified values are updated to the new values.

You can also delete values from the list. Note that when you delete a list value, all attributes using the deleted values are removed from the objects. For objects with multiple attribute values, only the deleted values are removed.

You can also move a value up or down in the list.

- **Integer:** Select this when the attribute is used to track whole numbers, such as serial numbers. When you select this type, the wizard displays the **Value Limits** fields where you can enter the range of allowed values. Note that you cannot change your entries in the **Value Limits** fields if you modify the attribute at a later date.
- **Email:** Select this when the attribute is used for email addresses. Email addresses are entered in the format *user@domain.com*.
- **URL:** Select this when the attribute is used for tracking URLs (Uniform Resource Locators). URLs must be entered in a valid format.
- **Date:** Select this when the attribute is used for dates. The date value is in YYYY-MM-DD format.
- **Comment:** Enter additional information about the attribute. You can enter up to 256 UTF-8 characters.

4. Click **Next** and complete the following:

- **Enable Inheritance:** Select this check box if you want to allow the extensible attribute and its values to be inherited by descendants in an inheritance chain. When you select this check box, inheritance is enabled for network related objects only. When you select this check box and restrict an attribute to certain objects, then the extensible attribute and its value will be inherited by those objects only.

Note the following:

- If you create an extensible attribute with inheritance disabled and later enable it, the *Descendant Actions* dialog box may be displayed with the available options for adding an extensible attribute. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329.
- If you create an extensible attribute with inheritance enabled and later disable it, the *Descendant Actions* dialog box may be displayed with the available options for deleting an extensible attribute. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.
- **Allow multiple values:** Select this check box if you want to allow multiple values for this attribute to be set on an object. You cannot change this value for predefined attributes.
- **Default Value:** Enter the default value that the appliance displays for the attribute. Leave this blank if there is no default value for this attribute. If the attribute type is **String**, you can enter up to 256 UTF-8 characters. If the attribute type is **List**, the value must be one of the list values and can be up to 64 UTF-8 characters.
- **Required:** If you select this option, it is required to enter a value for this attribute when adding or modifying the corresponding object in the GUI.
- **Recommended:** If you select this option, it is recommended to enter a value for this attribute when adding or modifying the corresponding object in the GUI.
- **Optional:** This is selected by default. By selecting this option, you may or may not enter a value for this attribute when adding or modifying the corresponding object in the GUI.
- **Restrict to Specific Object Types:** Click the Add icon to select the object type with which you want to associate the attribute. The appliance adds a row to the table. To delete an object type, select an object type and click the Delete icon. By default, the appliance associates an extensible attribute with all the supported object types.

- **Log Attribute Values When Objects are Updated:** Select this check box if you want the appliance to make an entry in the audit log each time an object with this attribute is added or modified. When you select attribute values for audit, they are included in all the audit log entries. For information about the audit log, see [Using the Audit Log](#) on page 1018.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

Grid Manager adds the attribute to the **Extensible Attributes** tab of the wizard and editor of the specified object types.

Configuring Inheritable Extensible Attributes

An extensible attribute can be inherited by descendants when it is at the top or in the middle of the inheritance chain.

When you add a new extensible attribute to a parent object, the same extensible attribute may or may not already exist at the descendant levels. If the extensible attribute exists on a descendant, you can choose to have the descendant inherit the value from the parent, or retain the original value from the descendant. When the extensible attribute does not exist on the descendant, you can choose to have the descendant either inherit the extensible attribute and its value from the parent or not inherit anything from the parent.

When you add a range, host, fixed address or IPv4 reservation to a parent object which has inheritable extensible attributes, the newly added object can inherit extensible attributes from the parent object. For example, if you create an IPv4 network with inheritable extensible attributes, and then add a host, the values you specified for the extensible attributes while creating the network can be inherited by the host.

To assist you in identifying whether an extensible attribute value is inherited or overridden, the appliance displays the inheritance state of an attribute in the **Inheritance State** column of an extensible attribute. This column is displayed only for objects that support inheritance. For information about how to view inheritance states, see [Modifying Inheritable Extensible Attributes](#) on page 330.

Following are the supported inheritance states:

Table 7.2 Inheritance States

Inheritance State	Description
Inherited	The extensible attribute inherits its value from the parent. You cannot edit the value of an attribute when the inheritance state is set to Inherited . You can change the state to Overridden and then change the value of the attribute or change the state to Not Inherited to remove the inherited value.
Overridden	The extensible attribute overrides the value inherited from the parent. You can change the state to Inherited and restore the original inherited value or change the state to Not Inherited and remove the inherited value.
Not Inherited	The extensible attribute can inherit its value from the parent, but the attribute does not exist on the descendant. You can change the state to Inherited and restore the original inherited value or change the state to Overridden and change the value of the attribute. Note that when the state of an inheritable extensible attribute is Not Inherited , the corresponding attribute will not be added as a new extensible attribute for objects that are currently not inheriting this extensible attribute.
No Parent	The inheritance state is set to No Parent when an object has a parent, but the parent does not have an extensible attribute or the parent's extensible attribute is set to Not Inherited .
Disabled	Extensible attribute inheritance is not enabled for the attribute.
No Change	The extensible attributes of the selected objects do not have the same inheritance state for all objects. This state allows you to retain the current state on the selected objects. Note that this state is only seen in the <i>Multi-object Extensible Attributes</i> editor.

When you add an inheritable extensible attribute to an object, if there are descendants of this object the *Descendant Actions* dialog box is displayed which will provide options for the descendants. Following is a summary of these options:

- Retain the original value of the attribute for all descendants.
- Inherit the extensible attribute and its value from the parent object.
- Inherit the extensible attribute and its value when it does not exist on descendants.
- If the extensible attribute does not exist on the descendant, do not add it.
- If you are deleting an inherited extensible attribute from a parent object you can retain or remove the extensible attribute from the object's descendants.

To configure default descendant actions for inheritable extensible attributes:

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
2. In the **Extensible Attribute Inheritance** tab, complete the following:

When adding an extensible attribute that already exists on a descendant:

- **Keep the descendant's existing value and change the inheritance state to Override:** Select this if you want to retain the existing extensible attribute values for all direct descendants, irrespective of the values you define at the parent level. The inheritance state for all direct descendants will be set to **Overridden**. Note that this is applicable only when you add a new extensible attribute to the parent object that already exists on the descendant. If you modify the value of an existing extensible attribute that is already inherited by the descendant, and select the above option in the Descendant Actions dialog box, then the new value will be inherited by the descendant, but the inheritance state will remain **Inherited**. For example, consider a network 10.0.0.0/16 that has an extensible attribute **Site** with the value **SA** (native). When you add another network 10.0.0.0/24, extensible attribute **Site** inherits its value, **SA**, from the parent object. Now, if you add network 10.0.0.0/8, assign extensible attribute **Site** and set its value to **NY**, then when you choose this option, the value of **Site** will remain as **SA**, but the inheritance state will be changed to **Overridden** for network 10.0.0.0/16; however, network 10.0.0.0/24 will still have its value as **SA** for **Site** with the inheritance state set to **Inherited**.
- **Inherit the parent's value and change the inheritance state to Inherit:** Select this to inherit the extensible attribute values from the parent for all descendants. The inheritance state for all descendants will be set to **Inherited**.
- **Change the inheritance state to Inherit only if the descendant's value is the same as the parent's value. Otherwise, change the state to Override:** Select this to set the inheritance state to **Inherit** if the descendants have the same extensible attribute value as the parent. Otherwise, retain the original extensible attribute value on the descendants and change the inheritance state to **Overridden**.

When adding an extensible attribute that does not exist on a descendant:

- **Do not inherit the value from the parent and change the inheritance state to Not Inherited:** Select this if the extensible attributes do not exist on the descendants and you do not want them to inherit the attributes from the parent. The inheritance state is set to **Not inherited**.
- **Inherit the value from the parent and change the inheritance state to Inherited:** Select this if you want all descendants to inherit extensible attributes from the parent, and the inheritance state for all descendants will be set to **Inherited**.

When deleting an extensible attribute:

- **Keep the descendant's value and change the inheritance state to No Parent:** Select this if you want to preserve extensible attributes on all descendants when you delete an inheritable extensible attribute. The inheritance state for direct descendants will be set to **No Parent**.
- **If the current inheritance state is Inherited, remove the extensible attribute. If the current inheritance state is Overridden, keep the value and change the inheritance state to No Parent:** Select this if you want to remove the extensible attributes that are inherited by the descendants. If the inheritance state of the extensible attributes is set to **Inherited** on the descendant, the attributes will be removed; however, if the inheritance state is set to **Overridden**, then the state will be changed to **No Parent**.

3. Save the configuration.

For more information about how to configure inheritable extensible attributes, see [Configuration Examples for Inheritable Extensible Attributes](#) on page 334.

Admin Permissions and Inheritable Extensible Attributes

Permissions for descendant objects can affect the results of the actions that are chosen in the *Descendant Actions* dialog box:

- When you add an extensible attribute to the parent object: The descendants to which you have read-write permission will behave as expected with any of the chosen options in the *Descendant Actions* dialog box.
- When you change the extensible attribute value on the parent object: The descendants that have the same extensible attribute set to **Inherited** will be automatically changed to the new value, even though you may not have write permission for those descendants.
- When you select to preserve descendant values while removing an extensible attribute associated with the parent object, values will be preserved even if you do not have write read-write permission for those descendants.
- When you select to remove an extensible attribute on descendants when removing a parent's extensible attribute, an error message will be displayed if you do not have read-write permission to some of the descendants.

Guidelines for Configuring Inheritable Extensible Attributes

- When you add an inheritable extensible attribute to a parent object, you can choose to have descendants inherit or override the parent's extensible attribute value. You can also choose that the extensible attribute not be added to a descendant.
- When you add a new parent with an inheritable extensible attribute, the options for changes to descendants remain the same as when you add a new inheritable extensible attribute to a parent. For more information, see [Configuring Inheritable Extensible Attributes](#) on page 326.
- When you add a new descendant to the existing parent with inheritable attributes, the descendant inherits all the extensible attributes. However, you can select if you want to inherit or override the values. If you set the inheritance state to **Not Inherited**, then the extensible attribute will not exist on the descendant, but you can later change the state to **Inherited** or **Overridden**. For more information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329.
- When you delete an inheritable extensible attribute associated with the parent, you can either preserve the extensible attribute values on the descendants or delete the inherited extensible attributes. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.
- When you delete a parent object and if there is grandparent, then the extensible attribute will be re-parented when you choose preserve. The current inheritance state of the attribute will be retained. If you delete a parent object and if there is no grandparent, then the inheritance state of the extensible attribute is changed to **No Parent** when you choose preserve.
- When you split a network, the extensible attribute will be copied to the newly created networks. For inheritable extensible attributes, the newly created network inherits the extensible attributes and the state is set to **Inherited**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329.
- When you join two networks to form a larger network, the *Descendant Actions* dialog box is displayed with the following options:

When joining networks, select the action(s) you want to apply to the merged networks:

- **Preserve extensible attributes for all descendants of the merged networks and change the inheritance state to No Parent:** Select this if you want to preserve the extensible attributes for all descendants of the merged networks. The inheritance state of the attributes will be changed to **No Parent**.
- **Remove extensible attributes from descendants of the merged networks:** Select this if you want to remove extensible attributes that are inherited by descendants.

Note: The options above apply only to extensible attributes which no longer have a parent, due to the merge. If the extensible attributes on descendants are also on the resulting merged network, then they will retain their current state.

- When you add multiple inheritable networks, new networks will automatically inherit all extensible attributes from the parent.

Managing Inheritable Extensible Attributes at the Parent and Descendant Level

You can define if descendants will inherit values from the parent when a new extensible attribute is added to the parent. You can also choose to override the values of the extensible attributes on the descendants.

When you delete an existing attribute, you can choose to either preserve the values at the descendant level or delete the values inherited by the descendants.

Note: The *Descendant Actions* dialog box is displayed only when an object has descendants and you are modifying extensible attributes that affects those descendants. However, the dialog box is always displayed when a join is performed for a network that has inheritable extensible attributes.

The following section describes configuration changes for inheritable extensible attributes:

1. **Network Container:** From the **Dashboards** tab, select the **Tasks** tab -> click **Add Networks**. Select a network, enter the required details. You can edit the inheritable extensible attributes that are displayed automatically. If this is a parent object, then you can add extensible attributes.
IPv4 Network: From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab. In the **Networks** section, select **IPv4 Network** from the Add drop-down menu. In the *Add IPv4 Network* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
IPv6 Network: From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab. In the **Networks** section, select **IPv6 Network** from the Add drop-down menu. In the *Add IPv6 Network* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
IPv4 Range: From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> *network* -> click *addr_range*, select **Range** from the Add drop-down menu. In the *Add IPv4 Range* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
IPv6 Range: From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> *network* -> click *addr_range*, select **Range** from the Add drop-down menu. In the *Add IPv6 Range* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
2. You can either add new extensible attributes to the parent object or modify original extensible attribute values. Click on the extensible attribute value displayed in the **Value** column of the respective attribute to modify the original value or click the Add icon to add a new attribute.
3. Select a state from the drop-down list displayed in the **Inheritance State** column. Note that you can only change the inheritance state of a descendant. You must select **Overridden** from the drop-down list to enter a new value. For more information about inheritance states, see [Table 7.2](#) on page 326. When an object has a parent and the parent does not have the object's inheritable extensible attribute, then the inheritance state of the extensible attribute is set to **No Parent** and the state cannot be changed.
4. **Select the inheritable extensible attributes for which you want to modify descendant actions:** Select this check box if you would like to apply the actions of the *Descendant Actions* dialog box for existing extensible attributes. Before you select this check box, you must select the extensible attributes which will be affected by the actions of the *Descendant Actions* dialog box.

Note: This check box is not displayed for hosts, fixed addresses, and reservations since they do not have descendants.

5. In the *Descendant Actions* dialog box, select options that will be applied for descendant objects as described in [Configuring Inheritable Extensible Attributes](#) on page 326.

The *Descendant Actions* dialog box displays all the mentioned options when you perform add and delete operations simultaneously. Consider an example where you add a new inheritable extensible attribute **Site**, and delete an existing inheritable attribute **Region** from the parent object, and then click **Save** to save both changes. In this case, the *Descendant Actions* dialog box displays all the options.

6. Save the configuration.

Viewing Extensible Attributes

To view the configured extensible attributes, from the **Administration** tab, select the **Extensible Attributes** tab. The panel displays the following information:

- **Name:** The name of the extensible attribute.
- **Type:** The type of data defined by the attribute.
- **Comment:** Comments entered for the extensible attribute.
- **Required:** Indicates whether users are required to complete this field.
- **Restricted to Objects:** The object types that are associated with the attribute.
- **Inheritance Enabled:** Indicates whether inheritance is enabled or not.

You can do the following in this panel:

- Sort the displayed data in ascending or descending order by column.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Add or delete extensible attributes.
- Print or export the data.

Modifying Extensible Attributes

When you modify an extensible attribute, all objects using the modified attributes are updated. You can perform inline editing by double-clicking the row of data that you want to modify. The appliance displays the inline editing editor in the selected row. Click **Save** after modifying the data. Note that you cannot edit extensible attributes that have multiple values.

To modify an extensible attribute:

1. In the **Administration** tab, select the **Extensible Attributes** tab.
2. Select the attribute and click the Edit icon.
3. In the **General** tab of the *Extensible Attributes* editor, you can only change the name of the attribute. You cannot change the data type. The data type for predefined attributes is string.
4. In the **Additional Properties** tab, you can modify any of the fields described in the step 4 of [Adding Extensible Attributes](#) on page 324.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying Inheritable Extensible Attributes

When values are inherited by a descendant, the inheritance state of the inherited extensible attribute is displayed as **Inherited**. You can select **Overridden** and specify a new value or select **Inherited** to retain the same value as the parent. If you select **Not Inherited**, the extensible attribute and its value will not be inherited. The inherited value will have a strike-through and you cannot edit the value when the state is set to **Not Inherited**.

In addition to the attribute values, the **Value** column of an inheritable extensible attribute also displays the name of the source and the object type of the extensible attribute. For example, a *Network Container* has a descendant, *Network*, which inherits an extensible attribute value from *Network Container* and *Network* has a descendant, *Fixed Address* that inherits the same extensible attribute value. In this case, *Fixed Address* shows *Network Container* as the source.

The following table displays various inheritance states and corresponding changes to source and object types that are displayed in the **Value** column of an extensible attribute.

Inheritance State	Source and Object Type in the Value Column
If an extensible attribute is a native attribute (an object which is at the top of the inheritance chain, or does not have ancestors),	Source is not displayed in the Value column. This column will not display the source details, if none of the selected objects support inheritance.
If the state of an extensible attribute is set to Inherited ,	then the Source and object is displayed. You cannot change the value of the extensible attribute.
If the state of an extensible attribute is set to Overridden or Not Inherited ,	then the Source will have a strike-through. You can change the state of such extensible attributes. You cannot change the value of the extensible attribute when the inheritance state is set to Not Inherited .

To modify the value and inheritance state of an inheritable extensible attribute:

1. **For IPv4 and IPv6 Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.
For IPv4 Range, IPv6 Range, Fixed Address, Reservation, and Host: From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> *network* -> click *addr_range*, click the Edit icon.
2. In the editor, click the **Extensible Attributes** tab, select the check box of the respective attribute.
3. At the parent level, click on the value you want to change and enter the new value.
At the descendant level, click on the value you want to change and enter the new value. Note that you can change the value only when the inheritance state is set to **Overridden**.
4. Select a state from the drop-down list displayed in the **Inheritance State** column. Note that you can only change the inheritance state of a descendant. You must select **Overridden** from the drop-down list to enter a new value. For more information, see [Table 7.2](#) on page 326.
5. **Select the inheritable extensible attributes for which you want to modify descendant actions:** Select this check box if you would like to apply the actions of the *Descendant Actions* dialog box for existing extensible attributes. Before you select this check box, select the extensible attributes which will be affected by the actions of the *Descendant Actions* dialog box. For more information about the *Descendant Actions* dialog box, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329.
6. Save the configuration.

Deleting Extensible Attributes

When you delete an extensible attribute, the appliance removes the attribute. All the attribute values set on the selected object types are removed from those objects. Once deleted, the attribute no longer exists in the system. Deleted attributes are not moved to the Recycle Bin. This operation might take a long time depending on the amount of data that needs to be deleted.

To delete extensible attributes:

1. In the **Administration** tab, select the **Extensible Attributes** tab.
2. Select the attribute and click the Delete icon.

- When the confirmation dialog box appears, click **Yes**.

Deleting Inheritable Extensible Attributes Associated with Parent Objects

When you remove an inheritable extensible attribute, which is associated with a parent object, you can choose to retain the descendant extensible attribute or remove it from all the descendants.

Note that you cannot delete extensible attributes that have the inheritance state set to **Overridden**, **Inherited**, and **Not Inherited**. You can delete an extensible attribute that is directly assigned to the object and has its inheritance state set to **No Parent** or if the inheritance state is **Disabled**.

To remove an inheritable extensible attribute associated with a parent object:

- IPv4 and IPv6 Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.
For IPv4 Range, IPv6 Range, Fixed Address, Reservation, and Host: From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> *network* -> click *addr_range*, select the object and click Edit icon.
- In the editor, click the **Extensible Attributes** tab, select the attributes and then click the Delete icon.
- When you click **Save and Close**, the *Descendant Actions* dialog box is displayed automatically with the following options:
Select the action(s) you want to apply to descendant objects that have the following extensible attribute(s):
When deleting an extensible attribute:
 - **Keep the descendant's value and change the inheritance state to No Parent:** Select this if you want to preserve extensible attributes for all descendants. The inheritance state of the extensible attribute changes to **No Parent**.
 - **If the current inheritance state is Inherited, remove the extensible attribute. If the current inheritance state is Overridden, keep the value and change the inheritance state to No Parent:** Select this if you want to remove the extensible attributes that are inherited by the descendants. If the inheritance state of the extensible attributes is set to **Inherited** on the descendant, the attributes will be removed; however, if the inheritance state is set to **Overridden**, then the state will be changed to **No Parent**.
- Click **Yes** to save the configuration or **No** to exit.

Note: The deleted extensible attributes will not be moved to the Recycle Bin and you cannot restore extensible attributes that are deleted.

Using Extensible Attributes

After a superuser admin configures the attributes of an object, they become available in the wizard and editor of the object. This section describes how users can then add and manage the attributes that were configured.

Grid Manager displays the required extensible attributes in the **Extensible Attribute** tab. You must enter values for all required attributes. If an object does not have required attributes, you can add the available optional attributes.

In the **Extensible Attribute** tab of an object, such as a network or host record, you can do the following:

- Enter values for extensible attributes
- Add attributes
- Change the inheritance state of an extensible attribute
- Select if descendants must inherit extensible attribute values from its parent
- Delete optional attributes

To enter values for the extensible attributes of an object:

- Open the editor of the object. For example, to enter values for the attributes of a network, select it and click its **Extensible Attributes** tab.
- Click the Value column of the attribute. You must enter values for all required attributes.
- Depending on the required attribute type, either enter or select a value for the attribute from the Value column.

- Based on whether the attribute is inheritable, the values are displayed in the **Inheritance State** column. This value can be set to **Inherited**, **Overridden** or **Not inherited**. If the object is at the top of the inheritance chain (Network View), then the inheritance state is not displayed. The inheritance state is set to **No Parent** only if an object has a parent, but the parent does not have the inherited extensible attribute. This column is not displayed if all selected objects do not belong to the supported inheritance chain. Example: Zones, DNS View, DNS records, etc.

To add attributes:

- Click the Add icon. Grid Manager adds a row to the table with the default attribute displayed.
- Click the default attribute and expand the list of available attributes.
- Select an attribute from the drop-down list.
- Enter or select a value for the attribute from the Value column.

To delete an attribute:

- Click the check box beside the attribute you want to delete.
- Click the Delete icon.

Note: You cannot delete an extensible attribute which has its inheritance state set to **Inherited**, **Overridden**, and **Not Inherited**. You can delete an extensible attribute that is directly assigned to the object and has its inheritance state set to **No Parent** or if the inheritance state is **Disabled**.

To delete all attributes:

- Click the **Attribute Name** check box.
- Click the Delete icon.

Note: You can delete only attributes that are not required. If you have one or more required attributes, you cannot use the delete all function.

- Save the configuration and click **Restart** if it appears at the top of the screen.

Editing Multiple Extensible Attribute Values

You can also manage the extensible attributes of multiple objects at the same time. For example, you can select several zones, and view and modify their extensible attributes all at once in the *Multi-Select Edit Extensible Attributes* editor.

Note that Grid Manager may not apply the changes you made to all the selected objects. It applies the change to objects that meet the following criteria:

- You have read/write permission to the object.
- The selected object is not locked by another user or does not have a scheduled pending task.
- If the attribute was restricted to certain object types, the object must be one of those types.

To edit multiple extensible attribute values:

- Select the objects whose extensible attributes you want to modify. You can select specific objects or select all objects in a dataset, as described in [Selecting Objects in Tables](#) on page 60.
- Expand the Toolbar and click **Extensible Attributes**.

Grid Manager displays the *Multi-Select Edit Extensible Attributes* dialog box which lists the extensible attributes of the selected objects. It displays the following information for each attribute:

- **Attribute Name:** This field displays the name of the extensible attribute associated with the selected object.
- **Value:** If the selected objects have the same value for the attribute, Grid Manager displays that value in this field. If the selected objects have different values for the attribute or if some have values and others do not, this field displays **Multiple Values** and the cell is highlighted in gray.

An attribute can have multiple rows if it allows multiple values. Grid Manager displays the values that all objects have in common, if any. Otherwise, it displays **Multiple Values**. This column displays the source for inherited extensible attributes only. Note that when you add new extensible attributes, edit values of existing extensible attributes or delete an extensible attribute, then the *Descendant Actions* dialog box is displayed, even if the objects do not have any descendants. For more information about Source values, see [Modifying Inheritable Extensible Attributes](#) on page 330.

- If you select objects that have the same inherited extensible attributes, but objects have different parents, then the **Source** column will display **Multiple Ancestors**.
- If the inheritance state of an extensible attribute is **Not Inherited**, then the extensible attribute will not be added as a new extensible attribute to objects that are currently not inheriting this extensible attribute.
- **Inheritance State:** This field displays the inheritance state of an extensible attribute. The column value can be **Inherited**, **Not Inherited**, **No Parent**, **No Change** or **Overridden**. This column is not displayed if all selected objects do not belong to the supported inheritance chain. Example: Zones, DNS View, DNS records, etc.
If extensible attributes for the selected objects have the same inheritance state, then the respective inheritance state is displayed in this column. When objects have different inheritance states, this column displays **No Change**, so that the current inheritance state is retained on the selected objects. If you then change the inheritance state of an extensible attribute to a specific state, the corresponding attribute will be changed to the selected inheritance state on all selected objects where the extensible attribute is currently inherited. If the object is at the top of the inheritance chain (Network View), then the inheritance state is not displayed. The inheritance state is set to **No Parent** only if an object has a parent, but the parent does not have the inherited extensible attribute. For more information about inheritance states, see [Table 7.2](#) on page 326.
- **Required:** This field displays **Yes** if the attribute is required in at least one object associated with the attribute. It displays **No** if the attribute is not required in any of the objects.

3. You can do the following:

- Change the value of an attribute. Depending on the attribute type, select the value and either enter a new value or select one from the drop-down list.
- Add an attribute to the selected objects. Click the Add icon. In the **Attribute Name** field of the new row, select an attribute from the list of available attributes and specify its value. If the attribute that you added was configured as a required attribute, the **Required** field displays **Yes**. Otherwise, it displays **No**.
- Delete an attribute. You can delete an attribute if it is not required. Select the attribute and click the Delete icon.

4. Click **OK** when you are finished modifying the extensible attributes.

Grid Manager applies your changes to the applicable objects. This operation might take a long time, depending on the amount of data being modified. You can choose to run this operation in the background, as described in [About Tasks](#) on page 72.

Configuration Examples for Inheritable Extensible Attributes

All examples in this section are based on the inheritance chain Network View -> Network Container -> Network -> Range -> Host/Fixed Address/Reservation, in which network view is at the top level and host, fixed address and reservation at the bottom of the inheritance chain.

Example 1

When you add an extensible attribute to the top object, the inheritance state is set to **No Parent**. For example, if you add a new inheritable extensible attribute, **Building**, to a network view, the inheritance state of this extensible attribute is set to **No Parent** for the network view.

Example 2

When you add an extensible attribute **Site** to the parent object **Network** that has a descendant **Range**, you can define **Site** as inheritable and add it to the **Network**. The descendant, **Range**, may or may not have the same extensible attribute. Infoblox displays a list of options that lets you either inherit the value or retain or override the existing value of the extensible attribute at the descendant level. Another option is to inherit the value of **Site**, only if the value for this attribute in **Range** is same as that in **Network**. You can also decide if **Range** should acquire the same value for **Site**, if it is not defined for **Range**. This change can be inherited by the descendants of **Range**.

Depending on your configuration, the inheritance state of the extensible attribute can display **Inherited**, **Overridden** or **Not Inherited**. If the object is at the top of the inheritance chain (Network View), then the inheritance state is not displayed. The inheritance state is set to **No Parent** only if an object has a parent, but the parent does not have the inherited extensible attribute.

Example 3

Examples in this section show different results when you add a new inheritable extensible attribute to an object located at the top or in the middle of the inheritance chain based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network	Region	ABC	Native
10.1.0.0/16	Network			

Example 3.1: Add an extensible attribute **Region** with value **DEF** to **10.0.0.0/8**

You select the following options for the existing extensible attribute:

- For descendants that already have this extensible attribute, the existing extensible attribute will always be set to **Inherit**.
- For descendants that do not have this extensible attribute, the descendants will inherit this extensible attribute.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	DEF	Inherited from 10.0.0.0/8
10.1.0.0/16	Network	Region	DEF	Inherited from 10.0.0.0/8

Example 3.2: Add an extensible attribute **Region** with value **DEF** to **10.0.0.0/8**

You select the following options for the existing extensible attribute:

- For descendants that already have this extensible attribute, the existing extensible attribute will always be set to **Override**.
- For descendants that do not have this extensible attribute, the descendants will not inherit this extensible attribute (extensible attribute is set to **Do not Inherit**).

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	ABC	Overridden
10.1.0.0/16	Network	Region		

Example 3.3: Add an extensible attribute Region with value DEF to 10.0.0.0/8

You select the following options for the existing extensible attributes:

- For descendants that already have this extensible attribute, the existing extensible attribute will always be set to **Inherit**.
- For descendants that do not have this extensible attribute, the descendants will not inherit this extensible attribute (extensible attribute is set to **Do not Inherit**).

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	DEF	Inherited, Source 10.0.0.0/8
10.1.0.0/16	Network	Region		

Example 4

Examples in this section show different results when you remove an existing inheritable extensible attribute from an object located at the top or in the middle of the inheritance chain based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	DEF	Inherited from 10.0.0.0/8
10.1.0.0/16	Network	Region	ABC	Overridden

Example 4.1: Remove extensible attribute Region with value DEF from 10.0.0.0/8

You select the following option for the existing extensible attribute:

- Remove extensible attributes with inheritance state set to **Inherited**. Extensible attributes with the state set to **Overridden** are not removed.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network			
10.1.0.0/16	Network	Region	ABC	Overridden

Example 4.2: Remove extensible attribute Region with value DEF from 10.0.0.0/8

You select the following option for the existing extensible attribute:

- Preserve all descendant extensible attributes, whether the state is set to **Inherited** or **Overridden**.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network	Region	DEF	Native
10.1.0.0/16	Network	Region	ABC	Native

Example 5

Examples in this section show different results when you remove parent object based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network Container	Region	DEF	Inherited from 10.0.0.0/8
10.0.0.0/24	Network	Region	DEF	Inherited from 10.0.0.0/8
10.0.1.0/24	Network	Region	ABC	Overridden
10.10.0.0/16	Network Container	Region	GHI	Overridden
10.10.0.0/24	Network	Region	GHI	Inherited from 10.10.0.0/16
10.10.1.0/24	Network	Region	JKL	Overridden

Example 5.1: Removing object 10.0.0.0/8 from the parent level

You select the following option for the existing extensible attribute:

- Remove extensible attributes with the inheritance state set to **Inherited**. Extensible attributes with the state set to **Overridden** are not removed.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/16	Network Container			
10.0.0.0/24	Network			
10.0.1.0/24	Network	Region	ABC	Overridden
10.10.0.0/16	Network Container	Region	GHI	Overridden
10.10.0.0/24	Network	Region	GHI	Inherited from 10.10.0.0/16
10.10.1.0/24	Network	Region	JKL	Overridden

Example 5.2: Removing object 10.0.0.0/8 from the parent level

You select the following option for the existing extensible attribute on descendants:

- Preserve all extensible attributes on the descendant.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/16	Network Container	Region	DEF	Native
10.0.0.0/24	Network	Region	DEF	Inherited from 10.0.0.0/16
10.0.1.0/24	Network	Region	ABC	Overridden
10.10.0.0/16	Network Container	Region	GHI	Native
10.10.0.0/24	Network	Region	GHI	Inherited from 10.10.0.0/16
10.10.1.0/24	Network	Region	JKL	Overridden

Example 5.3: Remove object 10.10.0.0/16 from the parent level

You select the following option for the existing extensible attribute on descendants:

- Remove extensible attributes with the inheritance state set to **Inherited**. Extensible attributes with the state set to **Overridden** are retained.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network Container	Region	DEF	Inherited from 10.0.0.0/8

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/24	Network	Region	DEF	Inherited from 10.0.0.0/8
10.0.1.0/24	Network	Region	ABC	Overridden
10.10.0.0/24	Network	Region		
10.10.1.0/24	Network	Region	JKL	Overridden

Example 6

Examples in this section show different results after you add an object in the middle of the inheritance chain based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	Admin	Native
10.0.0.0/16	Network Container	Owner	Admin	Inherited from 10.0.0.0/8
10.0.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.0.1.0/24	Network	Owner	Joe	Overridden
10.10.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.10.1.0/24	Network	Owner	Annie	Overridden

Example 6.1: Adding object 10.10.0.0/16 without extensible attributes

You select the following option for the existing extensible attribute:

- Retain values if the extensible attribute already exists, and inherit the attribute from the parent object if it does not exist.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	Admin	Native
10.0.0.0/16	Network Container	Owner	Admin	Inherited from 10.0.0.0/8
10.0.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.0.1.0/24	Network	Owner	Joe	Overridden
10.10.0.0/16	Network Container	Owner	Admin	Inherited from 10.0.0.0/8

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.10.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.10.1.0/24	Network	Owner	Annie	Overridden

Example 7

Examples in this section show different results after you modify inheritable extensible attributes with multiple values based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network Container	Region	MNO	Native
		Region	PQR	Native

Example 7.1: Adding extensible attribute Region with value GHI to 10.0.0.0/8

You select the following option for the existing extensible attributes:

- The descendants that already have this extensible attribute will inherit the value from the parent object.

Result: Multiple values will be replaced with the single inherited value.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	GHI	Native
10.0.0.0/16	Network Container	Region	GHI	Inherited from 10.0.0.0/8

Example 7.2: Adding extensible attribute Region with value GHI to 10.0.0.0/8

You select the following option for the existing extensible attributes:

- The descendants that already have this extensible attribute will override the value.

Result: Extensible attribute values on descendants are retained, but single attribute value will be overridden.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	GHI	Native
10.0.0.0/16	Network Container	Region	MNO	Overridden
		Region	PQR	Overridden

Example 8

Examples in this section show different results after you modify existing inheritable extensible attribute of an object, but you do not have required permission to modify some descendants. For information about admin permissions, see [About Administrative Permissions](#) on page 160.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.0.0.0/8	Network Container	Owner	Sam	Native	Write
10.0.0.0/16	Network Container	Owner	Sam	Inherited from 10.0.0.0/8	Read
10.0.0.0/24	Network	Owner	Sam	Inherited from 10.0.0.0/8	Read
10.0.1.0/24	Network	Owner	Bob	Overridden	Write
10.10.0.0/16	Network Container	Owner	John	Overridden	Read
10.10.0.0/24	Network	Owner	John	Inherited from 10.10.0.0/16	Read
10.10.1.0/24	Network	Owner	Max	Overridden	Read
10.20.0.0/16	Network Container	Owner	Sam	Inherited from 10.0.0.0/8	Write
10.20.0.0/24	Network	Owner			Read
10.20.1.0/24	Network	Owner			Read

Example 8.1: Removing object 10.0.0.0/8

You select the following option for the existing inheritable extensible attribute:

- Retain extensible attribute values on descendants that are inherited from this parent object.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.0.0.0/16	Network Container	Owner	Sam	Native	Read
10.0.0.0/24	Network	Owner	Sam	Native	Read
10.0.1.0/24	Network	Owner	Bob	Overridden	Write
10.10.0.0/16	Network Container	Owner	John	Overridden	Read
10.10.0.0/24	Network	Owner	John	Inherited from 10.10.0.0/16	Read
10.10.1.0/24	Network	Owner	Max	Overridden	Read

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.20.0.0/16	Network Container	Owner	Sam	Native	Write
10.20.0.0/24	Network				Read
10.20.1.0/24	Network				Read

Example 8.2: Removing object 10.0.0.0/8

You select the following option for the existing inheritable extensible attribute:

- Remove extensible attribute values from descendants that are inherited from this parent object.

The appliance displays an error message when you remove an extensible attribute that is associated with a descendant for which you do not have required permission.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.0.0.0/16	Network Container				
10.0.0.0/24	Network				
10.0.1.0/24	Network	Owner	Bob	Overridden	Write
10.10.0.0/16	Network Container	Owner	John	Overridden	Read
10.10.0.0/24	Network	Owner	John	Inherited from 10.10.0.0/16	Read
10.10.1.0/24	Network	Owner	Max	Overridden	Read
10.20.0.0/16	Network Container				Write
10.20.0.0/24	Network				Read
10.20.1.0/24	Network				Read

Example 9

Examples in this section show different results after you join multiple networks, based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.1.0.0/16	Network	Owner	John	Native
10.1.0.1	Fixed Address	Owner	John	Inherited from 10.1.0.0/16
10.2.0.0/16	Network	Owner	Sam	Native
10.2.0.1	Fixed Address	Owner	Jane	Overridden

Example 9.1: Joining networks 10.0.0.0/8 with 10.1.0.0/16

You select the following option for the existing extensible attribute:

- Join networks 10.0.0.0/8 with 10.1.0.0/16.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	John	Native
10.1.0.1	Fixed Address	Owner	John	Inherited from 10.0.0.0/8
10.2.0.1	Fixed Address	Owner	Jane	Overridden

Example 9.2: Joining networks 10.0.0.0/8 with 10.2.0.0/16

You select the following option for the existing extensible attribute:

- Join networks 10.0.0.0/8 with 10.2.0.0/16.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	Sam	Native
10.1.0.1	Fixed Address	Owner	Sam	Inherited from 10.0.0.0/8
10.2.0.1	Fixed Address	Owner	Jane	Overridden

MANAGING SECURITY OPERATIONS

The Grid provides certain security-related features. The following sections describe the different security-related features that you can set. For information about how to configure these features,

Enabling Support Access

Infoblox Technical Support might need access to your NIOS appliance to troubleshoot problems. This function enables an SSH (Secure Shell) daemon that only Infoblox Technical Support can access. If you have any questions, contact Infoblox Technical Support at support@infoblox.com. By default, this option is disabled.

Enabling Remote Console Access

This function makes it possible for a superuser admin to access the Infoblox CLI from a remote location using an SSH (Secure Shell) v2 client. The management system must have an SSH v2 client to use this function. After opening a remote console connection using an SSH client, log in using a superuser name and password. By default, this option is disabled. Note that only superusers can log in to the appliance through a console connection.

Permanently Disabling Remote Console and Support Access

You can permanently disable remote console (Secure Shell v2) access for appliance administration and for Infoblox Technical Support to perform remote troubleshooting. Disabling this type of access might be required in a high-security environment.

WARNING: AFTER PERMANENTLY DISABLING REMOTE CONSOLE AND SUPPORT ACCESS, YOU CANNOT RE-ENABLE THEM! NOT EVEN RESETTING AN APPLIANCE TO ITS FACTORY DEFAULT SETTINGS CAN RE-ENABLE THEM.

Restricting GUI/API Access

You can specify the IP addresses from which administrators are allowed to access the NIOS appliance. When the NIOS appliance receives a connection request, it tries to match the source IP address in the request with IP addresses in the list. If there is at least one item in the HTTP Access Control list and the source IP address in the request does not match it, the NIOS appliance ignores the request.

Caution: If you specify an address or network other than the one from which you are currently accessing the appliance, when you save your configuration, you will lose your administrative session and be unable to reconnect.

Enabling HTTP Redirection

You can enable the NIOS appliance to redirect administrative connection requests using HTTP to the secure HTTPS protocol. When you disable redirection, the NIOS appliance ignores any administrative connection requests not using HTTPS. By default, the NIOS appliance does not redirect HTTP connection requests to HTTPS. When you change this setting, the application restarts and your management session terminates.

Modifying the Session Timeout Setting

You can set the length of idle time before an administrative session to the Infoblox GUI times out. The default timeout value is 600 seconds (10 minutes).

If a user does not interact with the application for the specified time, the appliance displays a message that a timeout has occurred. Click **OK** to restart the GUI session.

Note: If you change the session timeout value, the new setting takes effect only after you log out and log back in.

Disabling the LCD Input Buttons

By default, the LCD input function is enabled, which allows you to use the LCD buttons on the front panel of a NIOS appliance to change the IPv4 address settings of the LAN port. You can disable this function if the appliance is in a location where you cannot restrict access exclusively to NIOS appliance administrators and you do not want anyone to be able to make changes through the LCD.

Configuring Security Features

You can manage only certain features at the member level. To configure security features for the Grid or an individual member:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the **Security** tab, complete the following:

- **Session Timeout(s):** This field is in the *Grid Properties* editor only. Enter a number between 60 and 31536000 seconds (one minute – one year) in the Session Timeout field. The default session timeout is 600 seconds (10 minutes).
- **Minimum Password Length:** This field is in the *Grid Properties* editor only. Specify the minimum number of characters allowed for an admin password.
- **Redirect HTTP to HTTPS:** This field is in the *Grid Properties* editor only. Select this option to have the appliance redirect HTTP connection requests to HTTPS.
- **Restrict GUI/API Access:** To control access to the GUI and API, select one of the following. You can restrict access using a named ACL or define individual ACEs. For information about named ACLs, see [Configuring Access Control](#) on page 306.
 - **Allow Any:** Select this to allow any clients to access the GUI and API. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 and IPv6 addresses and networks. GUI and API access restriction does not support TSIG key based ACEs. When you select this, the appliance blocks GUI and API access for all ACEs in the named ACL. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or an IPv6 address. Click the **Value** field and enter the IP address. The appliance blocks GUI and API access for this client. Note that if you specify the address from which you are currently accessing the appliance, when you save your configuration, you will lose your administrative session and be unable to reconnect.
 - **IPv4 Network and IPv6 Network:** Select this to add an IPv4 network or IPv6 network. Click the **Value** field and enter the network. The appliance blocks GUI and API access for this network. Note that if you specify the network from which you are currently accessing the appliance, when you save your configuration, you will lose your administrative session and be unable to reconnect.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
- **Enable Remote Console Access:** Select this option to enable superuser admins to access the Infoblox CLI from a remote location using SSH (Secure Shell) v2 clients. You can set this at the Grid and member levels.
- **Enable Support Access:** Select this check box to enable an SSH (Secure Shell) daemon that only Infoblox Technical Support can access. You can set this at the Grid and member levels.
- **Restrict Remote Console and Support Access to the MGMT Port:** This field is in the *Grid Member Properties* editor only. Select this check box to restrict SSH (Secure Shell) v2 access to the MGMT port only. This restricts Infoblox Technical Support and remote console connections—both of which use SSH v2—to just the MGMT port. For an HA pair, you can make an SSH v2 connection to the MGMT port on both the active and passive nodes.
Clear the check box to allow SSH v2 access to both the MGMT and LAN ports.
- **Permanently Disable Remote Console and Support Access:** This field is in the *Grid Properties* editor only. Select this option to permanently disable remote console (Secure Shell v2) access for appliance administration and for Infoblox Technical Support.

- **Enable LCD Input:** Select this check box to allow use of the LCD buttons on the front panel of a NIOS appliance to change the IP address settings of the LAN port. Clear this check box to disable this function. You can set this at the Grid and member levels.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING ETHERNET PORTS

Depending on your deployment and configuration choices, the Ethernet ports on the NIOS appliance perform different functions. The Ethernet ports that handle traffic on the NIOS appliance are as follows:

- **LAN1 port** – A 10/100/1000-Mbps gigabit Ethernet port that connects the appliance to the network. This is the default port for single independent appliances, single Grid members, and passive nodes in HA pairs. You must use the LAN1 port to set up the appliance initially. It handles traffic for all management services if you do not enable the MGMT and LAN2 ports. The passive node in an HA pair uses this port to synchronize the database with the active node.
- **LAN2 port** – A 10/100/1000-Mbps gigabit Ethernet port that connects the appliance to the network. The LAN2 port is not enabled by default. You can enable the LAN2 port and define its use through the GUI after the initial setup. By default, the appliance uses the LAN1 port (and HA port when deployed in an HA pair). To enable and configure the LAN2 port, you must have read/write permission to the Grid member on which you want to enable the port. The LAN2 port is available on Infoblox-250-A, 550-A, -1050-A, -1550-A, -1552-A, -1852-A, -2000-A, -4010, Trinzic 810, Trinzic 820, Trinzic 1410 and Trinzic 1420, Trinzic 2210, and Trinzic 2220 appliances. For information about how to use the LAN2 port, see [Using the LAN2 Port](#) on page 355.
- **HA port** – A 10/100/1000-Mbps gigabit Ethernet port through which the active node in an HA (high availability) pair connects to the network using a VIP (virtual IP) address. HA pair nodes also use their HA ports for VRRP (Virtual Router Redundancy Protocol) advertisements.
- **MGMT port** – A 10/100/1000-Mbps gigabit Ethernet port that you can use for appliance management or DNS service. You can enable the MGMT port and define its use through the GUI after the initial setup. If the MGMT port is enabled, the NIOS appliance uses it for management services (see [Table 7.5](#) on page 351 for specific types).

You can do the following on some of the Ethernet ports, depending on your network requirements and configurations:

- Assign VLANs (Virtual LANs) to the LAN1 and LAN2 ports so that NIOS can provide DNS service to different subnetworks on the same interface. For more information about VLANs, see [About Virtual LANs](#).
- Implement DiffServ (Differentiated Services) on the appliance by configuring the DSCP (Differentiated Services Code Point) value. For more information about DiffServ and DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.

About Virtual LANs

You can assign VLANs (Virtual Local Area Networks) to the LAN1, LAN2, and VIP (for HA pairs) interfaces so the appliance can provide DNS service to different subnetworks on the same interface. VLANs are independent logical networks that are mutually isolated on the interface so that IP packets can pass between them through one or more switches or routers. You can assign VLANs to provide segmentation services to address issues such as scalability, security, and network management. For example, you can partition your network into segments such as DHCP address allocation, DNS service, guest network, and DMZ (demilitarized zone) to achieve a higher level of security and to increase performance by limiting broadcast domains. You can also add quality of service schemes to optimize your network traffic on the VLAN trunk links by configuring the DSCP (Differentiated Services Code Point) value for the corresponding physical and virtual interfaces. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.

VLAN Tagging

When your VLANs span across multiple networks, VLAN tagging is required. VLAN tagging involves adding a VLAN tag or ID to the header of an IP packet so the appliance can identify the VLAN to which the packet belongs. In addition, switches use the VLAN tag to determine the port to which it should send a broadcast packet. The appliance uses the IEEE 802.1Q networking standard to support VLANs and VLAN tagging. On the appliance, you can configure VLANs as tagged networks by adding VLAN tags to them. Untagged networks are those without VLAN tags assigned to them. When you set up a VLAN as either a tagged or untagged network, ensure that you properly configure the corresponding switch for the VLAN to function properly.

VLANs and VLAN tagging are supported on both IPv4 and IPv6 transports. This feature is currently supported on the following Infoblox appliances: Trinzic 2210, Trinzic 2220, and Infoblox-4010. For information about these appliances, refer to the respective installation guides on the Infoblox Support web site at <http://www.infoblox.com/support>.

Currently, only the DNS service can listen on specific VLAN interfaces. The DHCP service listens only on the primary VLAN interface (tagged or untagged). However, if the primary VLAN interface is untagged, DHCP will serve all VLANs on that interface because an untagged primary VLAN receives all broadcast packets. You can also specify VLANs as the source port for sending DNS queries and notify messages. For information about how to configure these, see [Specifying Port Settings for DNS](#) on page 562.

Note: When you join an appliance that supports VLANs to a Grid that does not support VLANs or revert the appliance to a NIOS version that does not support VLANs, the appliance will become unreachable after joining the Grid or being reverted. You must remove VLAN tagging from the corresponding switch in order to reach the downgraded appliance.

Consider the following guidelines when tagging VLANs on the LAN1 and LAN2 ports:

- You can assign VLAN addresses to an interface and add VLAN tags to them. However, you must designate one of the tagged VLANs as a primary address.
- If the primary IPv4 address is tagged with a VLAN ID, all other addresses on the same interface must be tagged with a VLAN ID as well.
- You can use the same VLAN ID to tag only one IPv4 and one IPv6 address on the same interface. You cannot use the same VLAN ID to tag multiple IPv4 and IPv6 addresses on the same interface.
- You can assign one untagged IPv4 and one untagged IPv6 address to the same interface. These addresses are designated as the primary address for the interface.
- For IPv6, you must have a primary IPv6 address (either tagged or untagged) before you can add other tagged IPv6 addresses on the same interface.
- If you have multiple VLANs assigned to the LAN1 interface and the primary VLAN is untagged, DHCP listens on all VLAN interfaces and thus DHCP lease requests will succeed for the additional VLANs assigned to the LAN1 interface, but the request will actually be handled by the primary untagged VLAN interface.

Configuring VLANs

When you first set up a NIOS appliance, you can assign VLANs through the Grid Setup Wizard. For more information, see [Using the Setup Wizard](#) on page 242. After the initial setup, you can assign VLANs to the LAN1 port in the Required Ports and Addresses table, as described in [Modifying Ethernet Port Settings](#) on page 354.

On a Grid member, you can also assign up to 10 VLANs for each protocol (IPv4 or IPv6) on the LAN1 and LAN2 ports. You can assign up to 10 IPv4 VLAN addresses and 10 IPv6 VLAN addresses for each interface.

To assign additional VLANs to the LAN1 or LAN2 port, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. Select the **Network** -> **Basic** tab in the *Grid Member Properties* editor.
3. In the Additional IPv4 Ports and Addresses or Additional IPv6 Ports and Addresses table, click the Add icon and select **LAN1 (VLAN)** or **LAN2 (VLAN)** from the drop-down list. You can add up to 10 IPv4 and 10 IPv6 VLANs for each interface. Enter the following:

- **Interface:** Displays the name of the VLAN interface. This can be **LAN1 (VLAN)** or **LAN2 (VLAN)** depending on your selection. You cannot modify this.
- **Address:** Type the IP address for the VLAN port.
- **Subnet Mask:** For IPv4 only, specify an appropriate subnet mask.
- **Prefix Length:** For IPv6 only, choose the CIDR netmask for the subnet to which the VLAN address connects. CIDR is an alternative to classful subnet masking that organizes IP addresses into subnetworks. Also known as supernetting, CIDR allows multiple subnets to be grouped together for network routing. The prefix length ranges from 0 to 128, with common-sense values ranging from /48 to /128 due to the larger number of bits in the IPv6 address. Note that the IB-4030 supports the same netmask as the LAN1 interface or a /128 prefix.
- **Gateway:** Type the default gateway for the VLAN port.
You can now define a link-local address as the default IPv6 gateway and isolate the LAN segment so the local router can provide global addressing and access to the network and Internet. This is supported for both LAN1 and LAN2 interfaces as well as LAN1 and LAN2 in the failover mode.
- **VLAN Tag:** Enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly. For information about VLANs, see [About Virtual LANs](#) on page 346.
- **Port Settings:** For IPv4 only. From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
- **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Implementing Quality of Service Using DSCP

You can implement DiffServ (Differentiated Services) on the appliance by configuring the DSCP (Differentiated Services Code Point) value. DiffServ is a scalable and class-based mechanism that provides relative priorities to the type of services on your network. It can provide low latency for critical network traffic while providing simple best-effort service for non-critical services. The Infoblox DSCP implementation fully conforms to RFC 2475. For more information about DiffServ, refer to RFC 2475, *An Architecture for Differentiated Services*.

In IPv4 and IPv6 headers, DiffServ uses the DS (Differentiated Services) field for packet classification purposes. The DS field defines the layout of the ToS (Type of Services) octet in IPv4 and the Traffic Class octet in IPv6. The first six bits of the DS field are used as the DSCP value, which determines the PHBs (per-hop behaviors) on DiffServ compliant nodes and enables priorities of services to be assigned to network traffic. For more information about the DS field, refer to RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*.

When you configure the DSCP value for DiffServ, the appliance sets priorities for all outgoing IP traffic. It implements QoS (quality of service) rules so you can effectively classify and manage your critical network traffic. To ensure that core network services, such as DNS services, continue to operate in the event of network traffic congestion, you can set the DSCP value for the entire Grid and override it at the member level. Note that on an appliance, all outgoing IP traffic on all interfaces uses the same DSCP value.

DSCP is supported on both IPv4 and IPv6 transports. This feature is currently supported on the following Infoblox appliances: Trinzic 2210, Trinzic 2220, and Infoblox-4010. For information about these appliances, refer to the respective installation guides on the Infoblox Support web site at <http://www.infoblox.com/support>.

Note: You can set the DSCP value of the primary LAN using the `set network` CLI command. For information about the CLI command, refer to the *Infoblox CLI Guide*. DSCP values for all other interfaces and VLANs must be set through Grid Manager.

Configuring the DSCP Value

The DSCP value is set to zero (lowest priority) by default. You can change this value for the Grid and override the value at the member level. When you configure the DSCP value, all outgoing IP traffic on all interfaces uses the same value. Valid DSCP values are from 0 to 63. You can also set the DSCP value using the Infoblox CLI. For more information, refer to the *Infoblox CLI Guide*.

To configure the DSCP value for the Grid:

1. From the **Grid** tab -> **Grid Manager** tab, click **Grid Properties** -> **Edit** from the toolbar.
2. In the **General** -> **Advanced** tab of the *Grid Properties* editor, complete the following:
 - **DSCP Value:** Enter a value from 0 to 63. The default is 0 and it represents the lowest priority.
3. Save the configuration.

To override the DSCP value for a member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the **Network** tab -> **Basic** tab of the *Grid Member Properties* editor, complete the following:
 - **DSCP Value:** Click **Override**, and then enter a value from 0 to 63. The default is 0 and it represents the lowest priority.
3. Save the configuration.

You can override the Grid and member DSCP value at the interface level. For more information, see the following:

- For the LAN1 port, see [Modifying Ethernet Port Settings](#) on page 354.
- For the LAN2 port, see [Configuring the LAN2 Port](#) on page 357.
- For the MGMT port, see [Using the MGMT Port](#) on page 359.

Ethernet Port Usage

This section provides tables that detail the port usage and source and destination ports for different services, depending on your Grid configuration.

[Table 7.3](#) displays the type of traffic per port for both Grid and independent deployments. For a more detailed list of the different types of traffic, see [Table 7.5](#) on page 351.

Table 7.3 Appliance Roles and Configuration, Communication Types, and Port Usage

Appliance Role	HA Pair	HA Status	MGMT Port	Database Synchronization	Core Network Services	Management Services	GUI Access
HA Grid Master	Yes	Active	Disabled	VIP on HA	VIP on HA	LAN1	VIP on HA
HA Grid Master	Yes	Passive	Disabled	LAN1	—	LAN1	—
Single Grid Master	No	—	Disabled	LAN1	LAN1	LAN1	LAN1
HA Grid Member	Yes	Active	Disabled	LAN1	VIP on HA	LAN1	—
HA Grid Member	Yes	Passive	Disabled	LAN1	—	LAN1	—
Single Grid Member	No	—	Disabled	LAN1	LAN1	LAN1	—
Independent HA Pair	Yes	Active	Disabled	VIP on HA	VIP on HA	LAN1	VIP on HA
Independent HA Pair	Yes	Passive	Disabled	LAN1	—	LAN1	—
Single Independent	No	—	Disabled	—	LAN1	LAN1	LAN1
HA Grid Master	Yes	Active	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
HA Grid Master	Yes	Passive	Enabled	LAN1	—	MGMT	—

Appliance Role	HA Pair	HA Status	MGMT Port	Database Synchronization	Core Network Services	Management Services	GUI Access
Single Grid Master	No	–	Enabled	LAN1	LAN1 or MGMT	MGMT	MGMT
HA Grid Member	Yes	Active	Enabled	LAN1 or MGMT	VIP on HA	MGMT	–
HA Grid Member	Yes	Passive	Enabled	LAN1 or MGMT	–	MGMT	–
Single Grid Member	No	–	Enabled	LAN1 or MGMT	LAN1 or MGMT	MGMT	–
Independent HA Pair	Yes	Active	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
Independent HA Pair	Yes	Passive	Enabled	LAN1	–	MGMT	–
Single Independent	No	–	Enabled	–	LAN1 or MGMT	MGMT	MGMT
Reporting Member	No	–	Enabled	LAN1 or MGMT	LAN1 or MGMT	MGMT	MGMT

Table 7.4 Appliance Roles and Configuration, Communication Types, and Port Usage for Appliances with LAN2 Ports

Appliance Role	HA Status	MGMT Port	LAN2 Port	Database Synchronization	Core Network Services	Management Services	GUI Access
HA Grid Master	Active	Disabled	Enabled	VIP on HA	VIP on HA	LAN1 or LAN2	VIP on HA
HA Grid Master	Passive	Disabled	Enabled	LAN1	–	LAN1 or LAN2	–
Single Grid Master	–	Disabled	Enabled	LAN1	LAN1 and/or LAN2	LAN1 or LAN2	LAN1
HA Grid Member	Active	Disabled	Enabled	LAN1	VIP on HA	LAN1 or LAN2	–
HA Grid Member	Passive	Disabled	Enabled	LAN1	–	LAN1 or LAN2	–
Single Grid Member	–	Disabled	Enabled	LAN1	LAN1 and/or LAN2	LAN1 or LAN2	–
Independent HA Pair	Active	Disabled	Enabled	VIP on HA	VIP on HA	LAN1 or LAN2	VIP on HA
Independent HA Pair	Passive	Disabled	Enabled	LAN1	–	LAN1 or LAN2	–
Single Independent	–	Disabled	Enabled	–	LAN1 and/or LAN2	LAN1 or LAN2	LAN1
HA Grid Master	Active	Enabled	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
HA Grid Master	Passive	Enabled	Enabled	LAN1	–	MGMT	–
Single Grid Master	–	Enabled	Enabled	LAN1	LAN1, LAN2 and/or MGMT	MGMT	MGMT
HA Grid Member	Active	Enabled	Enabled	LAN1 or MGMT	VIP on HA	MGMT	–
HA Grid Member	Passive	Enabled	Enabled	LAN1 or MGMT	–	MGMT	–
Single Grid Member	–	Enabled	Enabled	LAN1 or MGMT	LAN1, LAN2 and/or MGMT	MGMT	–
Independent HA Pair	Active	Enabled	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
Independent HA Pair	Passive	Enabled	Enabled	LAN1	–	MGMT	–

Appliance Role	HA Status	MGMT Port	LAN2 Port	Database Synchronization	Core Network Services	Management Services	GUI Access
Single Independent	–	Enabled	Enabled	–	LAN1, LAN2 and/or MGMT	MGMT	MGMT
Reporting Member	–	Enabled	Enabled	LAN1 or MGMT	LAN1, LAN2, and/or MGMT	MGMT	MGMT

To see the service port numbers and the source and destination locations for traffic that can go to and from a NIOS appliance, see [Table 7.5](#). This information is particularly useful for firewall administrators so that they can set policies to allow traffic to pass through the firewall as required.

Note: The colors in both tables represent a particular type of traffic and correlate with each other.

Table 7.5 Sources and Destinations for Services

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
Key Exchange	LAN1 or MGMT on Grid member	VIP on HA Grid Master, or LAN1 on single master	17 UDP	2114	2114	Initial key exchange for establishing VPN tunnels Required for Grid
VPN	LAN1 or MGMT on Grid member	VIP on HA Grid Master, or LAN1 on single master	17 UDP	1194 or 5002, or 1024 -> 63999	1194 or 5002, or 1024 -> 63999	Default VPN port 1194 for Grids with new DNSone 3.2 installations and 5002 for Grids upgraded to DNSone 3.2; the port number is configurable Required for Grid
Discovery	LAN1 or LAN2 on Probe appliance		UDP		161	SNMP
Discovery	LAN1 or LAN2 on Probe appliance		UDP		260	SNMP - Needed for full discovery of some older Check Point models
Discovery	LAN1 or LAN2 on Probe appliance		ICMP		n/a	Ping Sweep
Discovery	LAN1 or LAN2 on Probe appliance		UDP, TCP		53	DNS
Discovery	LAN1 or LAN2 on Probe appliance		ICMP			Path Collection, for IPv4 addresses
Discovery	LAN1 or LAN2 on Probe appliance		UDP		33434 +1 per probe packet	Path Collection. Standard traceroute, for IPv6 addresses

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
Discovery	LAN1 or LAN2 on Probe appliance		ICMP, UDP, TCP			Port scan - all configured by user
Discovery	LAN1 or LAN2 on Probe appliance		UDP		137	NetBIOS
DHCP	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	17 UDP	68	67	Required for DHCP service
DHCP	LAN1, LAN2 or VIP on NIOS appliance	Client	17 UDP	67	68	Required for DHCP service
DHCP Failover	LAN1, LAN2 or VIP on Infoblox DHCP failover peer	LAN1, LAN2 or VIP on Infoblox DHCP failover peer	6 TCP	1024 -> 65535	519, or 647	Required for DHCP failover
DHCP Failover	VIP on HA Grid Master or LAN1 or LAN2 on single master	LAN1, LAN2 or VIP on Grid member in a DHCP failover pair	6 TCP	1024 -> 65535	7911	Informs functioning Grid member in a DHCP failover pair that its partner is down Required for DHCP failover
DDNS Updates	LAN1, LAN2, or VIP	LAN1, LAN2, or VIP	17 UDP	1024 -> 65535	53	Required for DHCP to send DNS dynamic updates
DNS Transfers	LAN1, LAN2, VIP, or MGMT, or client	LAN1, LAN2, VIP, or MGMT	6 TCP	53, or 1024 -> 65535	53	For DNS zone transfers, large client queries, and for Grid members to communicate with external name servers Required for DNS
DNS Queries	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	17 UDP	53, or 1024 -> 65535	53	For DNS queries Required for DNS
DNS Queries	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	6 TCP	53, or 1024 -> 65535	53	For DNS queries Required for DNS
NTP	NTP client	LAN1 or LAN2	17 UDP	1024 -> 65535	123	Required if the NIOS appliance is an NTP server
RADIUS Authentication	NAS (network access server)	LAN1 or VIP	17 UDP	1024 - 65535	1812	For proxying RADIUS Authentication-Requests. The default destination port number is 1812, and can be changed to 1024 - 63997. When configuring an HA pair, ensure that you provision both LAN IP addresses on the RADIUS server.

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
RADIUS Accounting	NAS (network access server)	LAN1 or VIP	17 UDP	1024 – 65535	1813	For proxying RADIUS Accounting-Requests. The default destination port number is 1813, and can be changed to 1024 – 63998.
RADIUS Proxy	LAN1 or VIP	RADIUS home server	17 UDP	1814	1024 -> 63997 (auth), or 1024 -> 63998 (acct)	Required to proxy requests from RADIUS clients to servers. The default source port number is 1814, and although it is not configurable, it is always two greater than the port number for RADIUS authentication.
ICMP Dst Port Unreachable	VIP, LAN1, LAN2, or MGMT, or UNIX-based client	LAN1, LAN2, or UNIX-based client	1 ICMP Type 3	–	–	Required to respond to the UNIX-based traceroute tool to determine if a destination has been reached
ICMP Echo Reply	VIP, LAN1, LAN2, or MGMT, or client	VIP, LAN1, LAN2, or MGMT, or client	1 ICMP Type 0	–	–	Required for response from ICMP echo request (ping)
ICMP Echo Request	VIP, LAN1, LAN2, or MGMT, or client	VIP, LAN1, LAN2, or MGMT, or client	1 ICMP Type 8	–	–	Required to send pings and respond to the Windows-based traceroute tool
ICMP TTL Exceeded	Gateway device (router or firewall)	Windows client	1 ICMP Type 11	–	–	Gateway sends an ICMP TTL exceeded message to a Windows client, which then records router hops along a data path
NTP	LAN1 on active node of Grid Master or LAN1 of independent appliance	NTP server	17 UDP	1024 -> 65535	123	Required to synchronize Grid, TSIG authentication, and DHCP failover Optional for synchronizing logs among multiple appliances
SMTP	LAN1, LAN2, or VIP	Mail server	6 TCP	1024 -> 65535	25	Required if SMTP alerts are enabled
SNMP	NMS (network management system) server	VIP, LAN1, LAN2, or MGMT	17 UDP	1024 -> 65535	161	Required for SNMP management
SNMP Traps	MGMT or LAN1 on Grid Master or HA pair, or LAN1 on independent appliance	NMS server	17 UDP	1024 -> 65535	162	Required for SNMP trap management. Uses MGMT (when enabled) or LAN1 on Grid Master or HA pair, or LAN1 on independent appliance for the source address, depending on the destination IP address.

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
SSHv2	Client	LAN1, LAN2, VIP, or MGMT on NIOS appliance	6 TCP	1024 -> 65535	22	Administrators can make an SSHv2 connection to the LAN1, LAN2, VIP, or MGMT port Optional for management
Syslog	LAN1, LAN2, or MGMT of NIOS appliance	syslog server	17 UDP	1024 -> 65535	514	Required for remote syslog logging
Traceroute	LAN1, LAN2, or UNIX-based appliance	VIP, LAN1, LAN2, or MGMT, or client	17 UDP	1024 -> 65535	33000 -> 65535	NIOS appliance responds with ICMP type code 3 (port unreachable)
TFTP Data	LAN1 or MGMT	TFTP server	17 UDP	1024 -> 65535	69, then 1024 -> 63999	For contacting a TFTP server during database and configuration backup and restore operations
VRRP	HA IP on the active node of HA pair	Multicast address 224.0.0.18	112 VRRP	802		For periodic announcements of the availability of the HA node that is linked to the VIP. The nodes in the HA pair must be in the same subnet.
HTTP	Management System	VIP, LAN1, or MGMT	6 TCP	1024 -> 65535	80	Required if the HTTP-redirect option is set on the Grid properties security page
HTTPS/SSL	Management System	VIP, LAN1, or MGMT	6 TCP	1024 -> 65535	443	Required for administration through the GUI
Reporting	Reporting Forwarders	LAN1, LAN2, or MGMT on the indexer	6 TCP	1024 - 65535	9997	Required for the reporting service. Communication is single directional from forwarders to the indexer. For example, a forwarder detects events and forwards them to the indexer.
Threat Protection	MGMT (with threat protection service running)	N/A (using FQDN)	HTTPS	N/A	443	For threat protection rule updates.

Modifying Ethernet Port Settings

By default, the NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between the 10/100Base-T and 10/100/1000Base-T ports on the NIOS appliance and the Ethernet ports on a connecting switch. It is usually unnecessary to change the default auto-negotiation setting; however, you can manually configure connection settings for a port if necessary.

Occasionally, for example, even though both the NIOS appliance and the connecting switch support 1000-Mbps (megabits per second) full-duplex connections, they might fail to auto-negotiate that speed and type, and instead connect at lower speeds of either 100 or 10 Mbps using potentially mismatched full- and half-duplex transmissions. If this occurs, first determine if there is a firmware upgrade available for the switch. If so, apply the firmware upgrade and test the connection. If that does not resolve the issue, manually set the ports on the NIOS appliance and on the switch to make 1000-Mbps full-duplex connections.

To change Ethernet port settings:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.

Note: You must enable the MGMT port before modifying its port settings. See [Using the MGMT Port](#) on page 359.

2. In the **Network** tab of the *Grid Member Properties* editor, the Required Ports and Addresses table lists the network settings that were configured. Complete the following to modify port settings:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Click the field and modify the IP address for the LAN1 port, which must be in a different subnet from that of the LAN2 and HA ports.
 - **Subnet Mask:** For IPv4 only, click the field and specify an appropriate subnet mask.
 - **Prefix Length:** For IPv6 only, click the field and specify the prefix length.
 - **Gateway:** Click the field and modify the default gateway for the LAN1 port.
 - **VLAN Tag:** Click the field and enter the VLAN tag ID if the port is configured for VLANs. You can enter a number from 1 to 4095. For information about VLAN, see [About Virtual LANs](#) on page 346.
 - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: The port settings on the connecting switch must be identical to those you set on the NIOS appliance.

USING THE LAN2 PORT

Note: This feature is not supported on vNIOS Grid members for Riverbed.

The LAN2 port is a 10/100/1000Base-T Ethernet connector on the front panel of Infoblox-250-A, -550-A, -1050-A, -1550-A, -1552-A, -1852-A, -2000-A, and -4010 appliances, and on the front panel of Trinzic 810, Trinzic 820, Trinzic 1410 and Trinzic 1420, Trinzic 2210, and Trinzic 2220 appliances. By default, the LAN2 port is disabled and the appliance uses the LAN1 port (and HA port when deployed in an HA pair). Before you can enable and configure the LAN2 port on a Grid member, you must first configure the member and join it to the Grid. You must also have read/write permission to the Grid member on which you want to enable the port. When you enable the LAN2 port and SNMP, the appliance sends traps from this port for LAN2 related events.

You can configure the LAN2 port in different ways. You can enable the port redundancy or port failover feature, which groups the LAN1 and LAN2 ports into one logical interface. The LAN1/LAN2 grouping can be activated for both IPv4 and IPv6. Alternatively, you can configure the LAN2 port on a different IP network than LAN1, and enable the LAN2 port to provide DNS and DHCP services. For information about these features, see the following sections:

- For information about the LAN2 failover feature, see [About Port Redundancy](#) on page 356.
- For information about configuring the LAN2 port, see [Configuring the LAN2 Port](#) on page 357.
- For information about enabling the LAN2 port to provide DHCP services, see [Enabling DHCP on LAN2](#) on page 358.
- For information about enabling the LAN2 port to provide DNS services, see [Enabling DNS on LAN2](#) on page 358.

Note that you cannot use the LAN2 port to access the GUI and the API, or to connect to the Grid. This can impact the ability of other appliances, such as the Network Automation and PortIQ appliances, to communicate with the Grid Master.

Any IPv6 services enabled for the LAN2 port also require provisioning of an IP address on the LAN2 port.

About Port Redundancy

You can configure the LAN2 or LAN2 (VLAN) port to provide redundancy and additional fault tolerance in your network. Port redundancy is transparently supported for both IPv4 and IPv6. When you enable port redundancy, the LAN1 or LAN1 (VLAN) and LAN2 or LAN2 (VLAN) ports are grouped into one logical interface. They share one IP address and appear as one interface to the network. Then, if a link to one of the ports fails or is disabled, the appliance fails over to the other port, avoiding a service disruption.

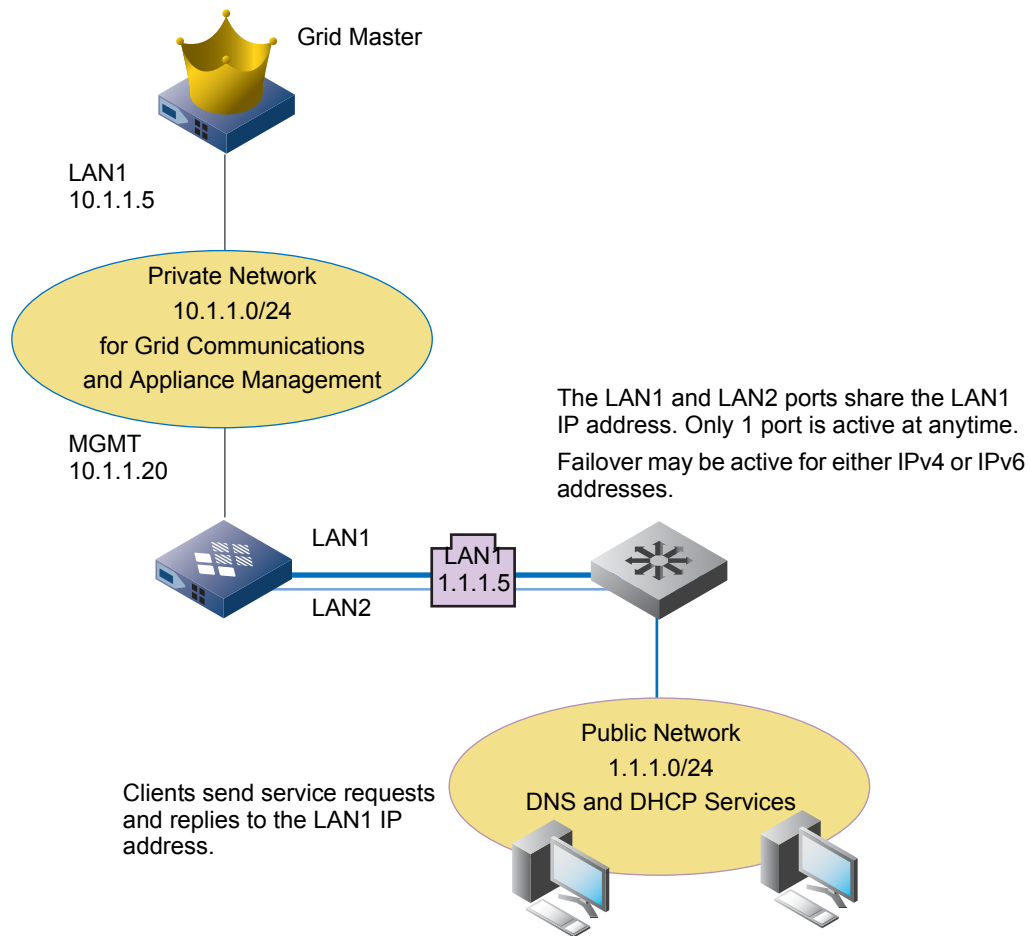
You can connect the LAN1 or LAN1 (VLAN) and LAN2 or LAN2 (VLAN) ports to the same switch or to different switches, but they must be on the same VLAN. One port is active and the other port is idle at all times. In case of failure in the LAN1 or LAN1 (VLAN) port, the LAN2 or LAN2 (VLAN) port becomes active and once the LAN1 or LAN1 (VLAN) port is active again, the LAN2 or LAN2 (VLAN) port becomes passive.

The LAN1 or LAN1 (VLAN) and LAN2 or LAN2 (VLAN) ports share the IP address of the LAN1 or LAN1 (VLAN) port; the port that is currently active owns the IP address. When you enable services on the appliance, such as DNS and DHCP, clients send their service requests to the LAN1 or LAN1 (VLAN) port IP address and receive replies from it as well. The port supports the services and features supported on the LAN1 or LAN1 (VLAN) port as listed in [Table 7.4](#) and [Table 7.5](#). You cannot enable the port redundancy feature if the LAN2 or LAN2 (VLAN) port is serving DNS or DHCP.

For example, you can use the MGMT port for Grid communications, as shown in [Figure 7.6](#), and the LAN1 and LAN2 ports are connected to the same switch. The LAN1 and LAN2 port share the IP address of the LAN1 port, which is 1.1.1.5. In the illustration, LAN1 is the active port.

You can also have the MGMT port disabled and configure LAN1 and LAN2 for port redundancy. You can enable port redundancy on single or HA independent appliances and Grid members.

Figure 7.6 Using the LAN2 Failover Feature



Before you enable port redundancy, ensure that both LAN1 and LAN2 are enabled. To enable port redundancy:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, select the **Enable port redundancy on LAN/LAN2** check box.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

The *Detailed Status* panel displays the status of both the LAN1 and LAN2 ports. In an HA pair, both nodes display the port information when port redundancy is enabled.

Configuring the LAN2 Port

Before you enable the LAN2 port to provide DHCP and DNS services, you must specify its IP address and other properties.

To configure the LAN2 port:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, click the Add icon of the Additional IPv4 Ports and Addresses or Additional IPv6 Ports and Addresses table and select **LAN2** from the drop-down list. Enter the following:
 - **Interface:** Displays the name of the interface. You cannot modify this.

- **Address:** Type the IP address for the LAN2 port, which must be in a different subnet from that of the LAN1 and HA ports.
 - **Subnet Mask:** For IPv4 only, specify an appropriate subnet mask.
 - **Prefix Length:** For IPv6 only, specify the prefix length.
 - **Gateway:** Type the default gateway for the LAN2 port.
 - **VLAN Tag:** Enter the VLAN tag ID if the port is configured for VLANs. You can enter a number from 1 to 4095. For information about VLAN, see [About Virtual LANs](#) on page 346.
 - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and then enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
 - **LAN2 Virtual Router ID (if HA):** If the appliance is in an HA pair, enter a VRID number.
3. Save the configuration and click **Restart** if it appears at the top of the screen.
The *Detailed Status* panel displays the status of the LAN2 port. In an HA pair, only the active node displays the LAN2 information.

Enabling DHCP on LAN2

You can configure an appliance to provide DHCP service through the LAN1 port, LAN2 port, or both the LAN1 and LAN2 ports. Note that when you enable both ports, they must be connected to different subnets. You can also start and stop DHCP service for IPv4 or IPv6 on the LAN1 or LAN2 port after you have enabled the service.

After you configure the LAN2 port, you can enable DHCP services on the LAN2 port as follows:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → *Grid_member* check box, and then click the Edit icon.
2. *If you are running DHCP for IPv4:* In the **General** → **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv4** check box for **LAN2** under DHCP Interfaces.
If you are running DHCP for IPv6: In the **General** → **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv6** check box for **LAN2** under DHCP Interfaces. (An IPv6 address must also be provisioned for the port.)
You can run either or both protocols for DHCP depending on your network deployment.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note that you can start or stop DHCP service on the LAN1 or LAN2 port.

Enabling DNS on LAN2

If you enable DNS on an appliance, it always serves DNS on the LAN1 port. Optionally, you can configure the appliance to provide DNS services through the LAN2 port as well. For example, the appliance can provide DNS services through the LAN1 port for internal clients on a private network, and DNS services through the LAN2 port for external clients on a public network.

After you configure the LAN2 port, you can enable DNS services on the LAN2 port as follows:

1. From the **Data Management** tab, select the **DNS** tab → **Members** tab → *Grid_member* check box, and then click the Edit icon.
2. In the **General** → **Basic** tab of the *Member DNS Configuration* editor, do the following:
If you are running DNS for IPv4: In the **General** → **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv4** check box for **LAN2** under DNS Interfaces.

If you are running DNS for IPv6: In the **General** -> **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv6** check box for **LAN2** under DNS Interfaces. (An IPv6 address must also be provisioned for the port.)

You can run either or both protocols for DNS depending on your network deployment.

- **Automatically create glue A and PTR records for LAN2's address:** The NIOS appliance can automatically generate A (address) and PTR records for a primary name server whose host name belongs to the name space of the zone. Select this check box to enable the appliance to automatically generate an A and PTR record.
 - **Automatically create IPv6 glue AAAA and PTR records for LAN2's address:** automatically generate AAAA and PTR records for the LAN2 IPv6 address. A glue record is the IP address of a name server held at the domain name registry. They are needed to set a domain's name server to a host name within the domain. Example: to set the name servers of ns1.corp100.com and ns2.corp100.com, provide the glue records, which are in effect the IP addresses, for ns1.corp100.com and ns2.corp100.com, within specific DNS record types. Without the glue records, DNS requests never resolve to the correct IP address because the domain registry does not associate the IP with the correct records.
3. In the **General** -> **Advanced** tab (click **Toggle Advanced Mode** if necessary), select one of the following from the **Send queries from** and the **Send notify messages and zone transfer request from** drop-down lists:
- **VIP:** The appliance uses the IP address of the HA port as the source for queries, notifies, and zone transfer requests.
 - **MGMT:** The appliance uses the IP address of the MGMT port as the source for queries, notifies, and zone transfer requests.
 - **LAN2:** The appliance uses the IP address of the LAN2 port as the source for queries, notifies, and zone transfer requests.
 - **Any:** The appliance chooses which port to use as the source for queries, notifies, and zone transfer requests.

The **Send queries from** drop-down list also includes loopback IP addresses that you configured. You can select a loopback address as the source for queries.

4. Save the configuration and click **Restart** if it appears at the top of the screen.
5. Click **Restart** to restart services.

USING THE MGMT PORT

Note: This feature is not supported on vNIOS Grid members for Riverbed.

The MGMT (Management) port is a 10/100/1000Base-T Ethernet connector on the front panel of an Infoblox-250-A, -550-A, -1050-A, -1550-A, -1552-A, -2000-A, Trinzic 810, Trinzic 820, Trinzic 1410, Trinzic 1420, Trinzic 2210, and Trinzic 2220 appliances. It allows you to isolate the following types of traffic from other types of traffic on the LAN and HA ports:

- [Appliance Management](#) on page 360
- [Grid Communications](#) on page 362
- [DNS Services](#) on page 364

For information about what types of traffic qualify as appliance management, Grid communications, and DNS services, see [Table 7.5](#) on page 351.

Note: The MGMT port currently does not support DHCP, NAT, or TFTP. IPv6 addressing may be applied to the MGMT port.

Some NIOS appliance deployment scenarios support more than one concurrent use of the MGMT port. The following table depicts MGMT port uses for various appliance configurations.

Table 7.6 Supported MGMT Port Uses for Various appliance Configurations

Appliance Configuration	Appliance Management	Grid Communications	DNS Services
Single Independent Appliance	✓	Not Applicable	✓
Independent HA Pair	✓	Not Applicable	◆
Grid Master	✓	✗	◆
Grid Master Candidate	✓	✗	◆
HA Grid Member	*	✓	◆
Single Grid Member	*	✓	✓

* Although you manage all Grid members through the Grid Master, if you enable the MGMT port on common Grid members, they can send syslog events, SNMP traps, and e-mail notifications, and receive SSH connections on that port.

Infoblox does not support MGMT port usage for some appliance configurations (indicated by the symbol ✗ in [Table 7.6](#)) because it cannot provide redundancy through the use of a VIP. A Grid Master that is an HA pair needs the redundancy that a VIP interface on the HA port provides for Grid communications. Similarly, DNS servers in an HA pair need that redundancy to answer DNS queries. Because the MGMT port does not support a VIP and thus cannot provide redundancy, Grid Masters (and potential Grid Masters) do not support Grid communications on the MGMT port.

In addition, NIOS appliances in an HA pair support DNS services on the active node only (indicated by the symbol ◆ in [Table 7.6](#)). Only the active node can respond to queries that it receives. If a DNS client sends a query to the MGMT port of the node that happens to be the passive node, the query can eventually time out and fail.

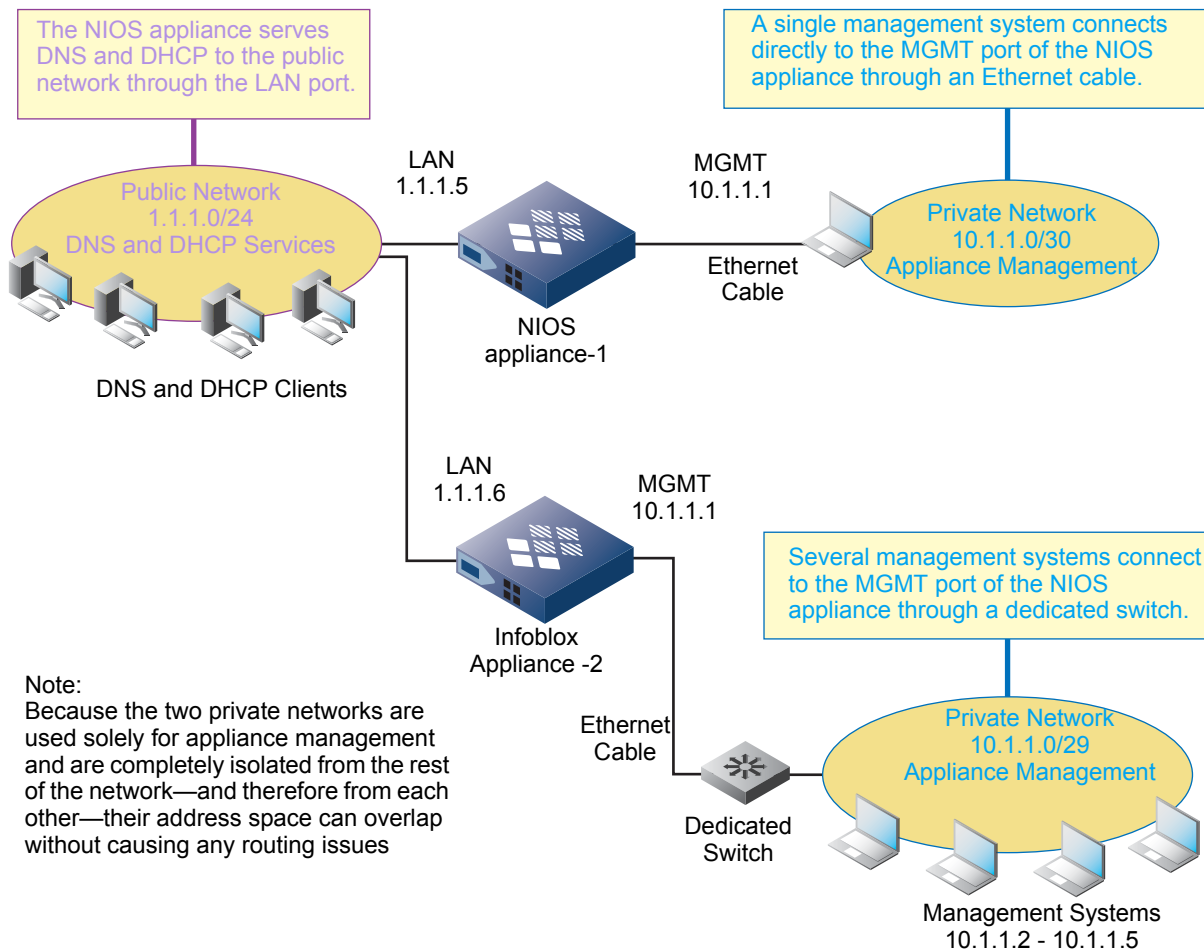
The MGMT port is not enabled by default. By default, a NIOS appliance uses the LAN port (and HA port when deployed in an HA pair). You must log in using a superuser account to enable and configure the MGMT port. You can enable the MGMT port through the Infoblox GUI, as explained in the following sections.

Appliance Management

You can restrict administrative access to a NIOS appliance by connecting the MGMT port to a subnet containing only management systems. This approach ensures that only appliances on that subnet can access the Infoblox GUI and receive appliance management communications such as syslog events, SNMP traps, and e-mail notifications from the appliance.

If you are the only administrator, you can connect your management system directly to the MGMT port. If there are several administrators, you can define a small subnet—such as 10.1.1.0/29, which provides six host IP addresses (10.1.1.1–10.1.1.6) plus the network address 10.1.1.0 and the broadcast address 10.1.1.7—and connect to the NIOS appliance through a dedicated switch (which is not connected to the rest of the network). [Figure 7.7](#) shows how an independent appliance separates appliance management traffic from network protocol services. Note that the LAN port is on a different subnet from the MGMT port.

Figure 7.7 Appliance Management from One or More Management Systems



Similarly, you can restrict management access to a Grid Master to only those appliances connected to the MGMT ports of the active and passive nodes of the Grid Master.

To enable the MGMT port on an independent appliance or Grid Master for appliance management and then cable the MGMT port directly to your management system or to a network forwarding appliance such as a switch or router:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, add the MGMT port to the Additional Ports and Addresses table as follows:
3. Click the Add icon and select **MGMT**.
Grid Manager adds a row for the MGMT port. For an HA pair, it adds two rows, one for each node.
4. Enter the following in the row of the MGMT port for a single Grid Master or independent appliance, and in the rows of the two nodes for an HA Grid Master or independent HA pair:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IP address for the MGMT port, which must be in a different subnet from that of the LAN and HA ports.
 - **Subnet Mask:** For IPv4 only, specify an appropriate subnet mask for the number of management systems that you want to access the appliance through the MGMT port.
 - **Prefix Length:** For IPv6 only, specify the prefix length.

- **Gateway:** Type the default gateway for the MGMT port. If you need to define any static routes for traffic originating from the MGMT port—such as SNMP traps, syslog events, and email notifications—destined for remote subnets beyond the immediate subnet, specify the IP address of this gateway in the route.
 - **Port Settings:** Choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and then enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
5. In the **Network** -> **Advanced** tab, make sure that the **Enable VPN on MGMT Port** check box is not selected.
 6. Save the configuration and click **Restart** if it appears at the top of the screen.
 7. Log out of Grid Manager.
 8. Cable the MGMT port to your management system or to a switch or router to which your management system can also connect.
 9. If your management system is in a subnet from which it cannot reach the MGMT port, move it to a subnet from which it can.
The Infoblox Grid Manager GUI is now accessible through the MGMT port on the NIOS appliance from your management system.
 10. Open an Internet browser window and enter the IP address of the MGMT port as follows: *https://<IP address of MGMT port>*.
 11. Log in to Grid Manager.
 12. Check the *Detailed Status* panel of the Grid member to make sure the status icons are green.

Grid Communications

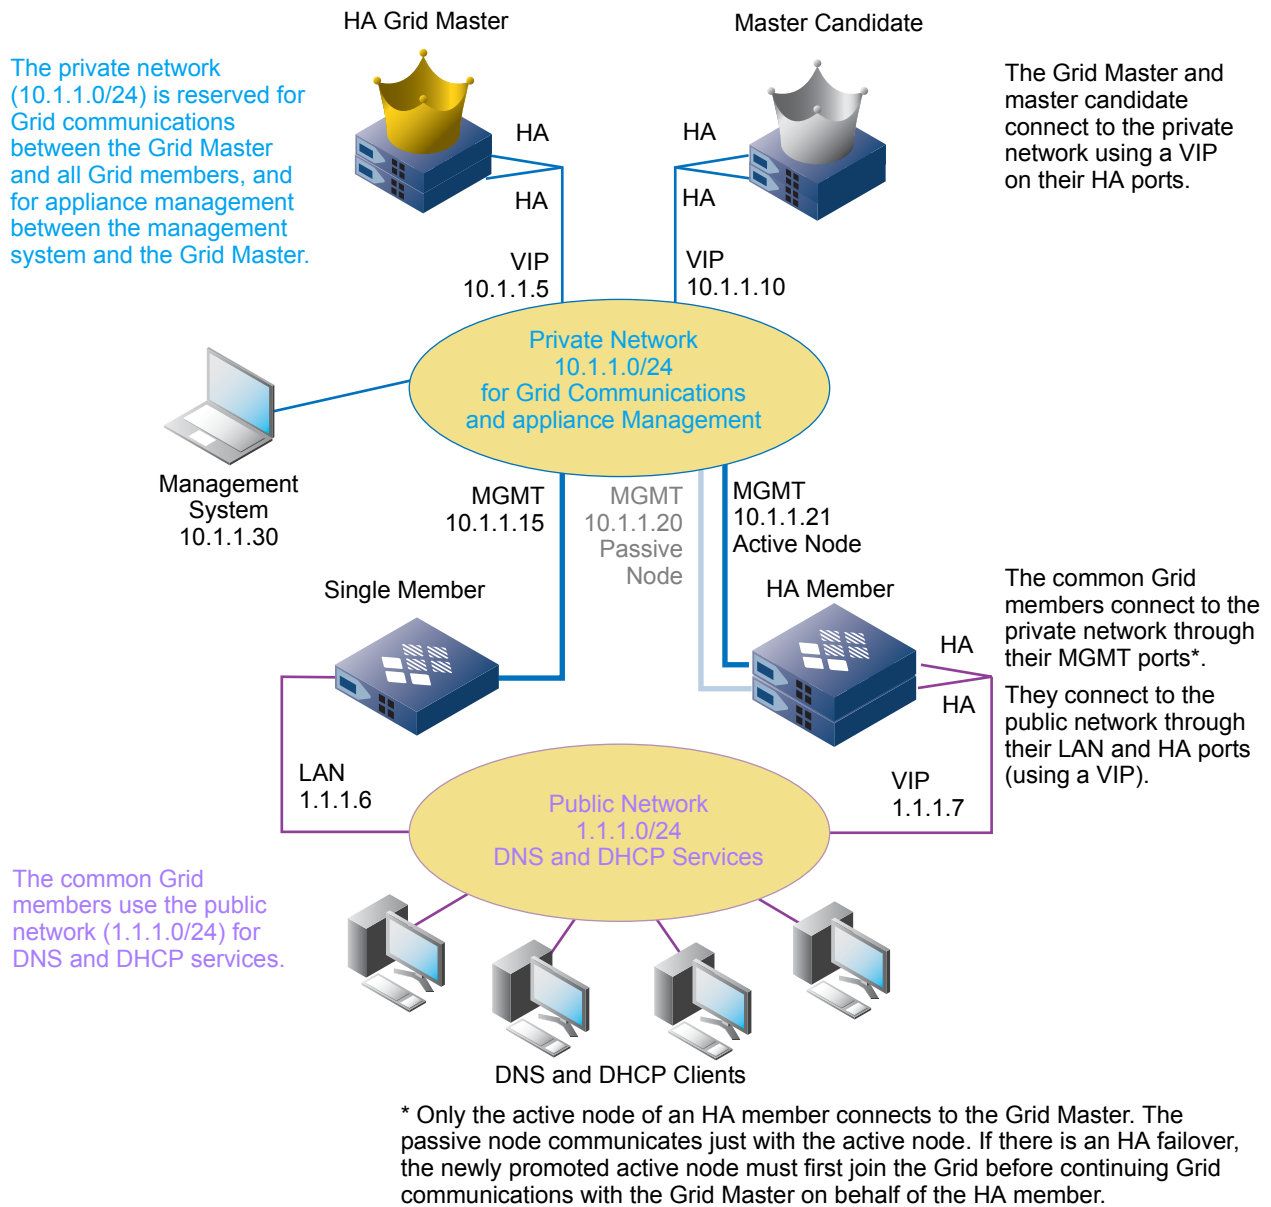
You can isolate all Grid communications to a dedicated subnet as follows:

- For Grid communications from the Grid Master, which can be an HA pair or a single appliance, the master uses either the VIP interface on the HA port of its active node (HA master) or its LAN port (single master). Neither a single nor HA Grid Master can use its MGMT port for Grid communications. (This restriction applies equally to master candidates.)
- Common Grid members connect to the Grid Master through their MGMT port.

This ensures that all database synchronization and Grid maintenance operations are inaccessible from other network elements while the common Grid members provide network protocol services on their LAN ports.

[Figure 7.8](#) shows how Grid members communicate to the master over a dedicated subnet.

Figure 7.8 Grid Communications



Enabling Grid Communications over the MGMT Port for Existing Grid Members

To enable the MGMT port for Grid communications on an existing single or HA Grid member:

1. Log in to the Grid Master with a superuser account.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.

Note: You must enable the MGMT port before modifying its port settings. See [Using the MGMT Port](#) on page 359.


3. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, add the MGMT port to the Additional Ports and Addresses table as follows:
4. Click the Add icon and select **MGMT**.
Grid Manager adds a row for the MGMT port. For an HA pair, it adds two rows, one for each node.

5. Enter the following in the row of the MGMT port for a single Grid Master or independent appliance, and in the rows of the two nodes for an HA Grid Master or independent HA pair:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IP address for the MGMT port, which must be in a different subnet from that of the LAN and HA ports.
 - **Subnet Mask:** For IPv4 only, specify an appropriate subnet mask for the number of management systems that you want to access the appliance through the MGMT port.
 - **Prefix Length:** For IPv6 only, specify the prefix length.
 - **Gateway:** Type the default gateway for the MGMT port. If you need to define any static routes for traffic originating from the MGMT port—such as SNMP traps, syslog events, and email notifications—destined for remote subnets beyond the immediate subnet, specify the IP address of this gateway in the route.
 - **Port Settings:** Choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
6. In the **Network** -> **Advanced** tab, select the **Enable VPN on MGMT Port** check box.
7. In the **Security** tab, do the following:
 - **Restrict Remote Console and Support Access to MGMT Port:** Select this check box to restrict SSH (Secure Shell) v2 access to the MGMT port only. This restricts Infoblox Technical Support and remote console connections—both of which use SSH v2—to just the MGMT port. For an HA pair, you can make an SSH v2 connection to the MGMT port on both the active and passive nodes.
Clear the check box to allow SSH v2 access to both the MGMT and LAN ports. For an HA pair, you can make an SSH v2 connection to the MGMT and LAN ports on both the active and passive nodes.
8. Save the configuration and click **Restart** if it appears at the top of the screen.
The master communicates the new port settings to the member, which immediately begins using them. The member stops using its LAN port for Grid communications and begins using the MGMT port.
9. To confirm that the member still has Grid connectivity, check that the status icons for that member are green on the *Detailed Status* and *Grid* panels.

DNS Services

You can configure a single independent appliance or single Grid member to provide DNS services through the MGMT port in addition to the LAN port. For example, the appliance can provide DNS services through the MGMT port for internal clients on a private network, and DNS services through the LAN port for external clients on a public network. While providing DNS services on the MGMT port, you can still use that port simultaneously for appliance management. [Figure 7.9](#) shows a management system communicating with a single independent appliance through its MGMT port while the appliance also provides DNS services on that port to a private network. Additionally, the appliance provides DNS services to an external network through its LAN port.

Figure 7.9 DNS Services on the LAN and MGMT Ports, and appliance Management on the MGMT Port



Like a single independent appliance, a single Grid member can also support concurrent DNS traffic on its MGMT and LAN ports. However, because you manage all Grid members through the Grid Master, a Grid member only uses an enabled MGMT port to send SNMP traps, syslog events, and email notifications, and to receive SSH connections.

In addition, the active node of an HA pair can provide DNS services through its MGMT port. To use this feature, you must enable DNS services on the MGMT ports of both nodes in the HA pair and specify the MGMT port IP addresses of both nodes on the DNS client as well, in case there is a failover and the passive node becomes active. Note that only the active node can respond to queries that it receives. If a DNS client sends a query to the MGMT port of the node that happens to be the passive node, the query can eventually time out and fail.

To enable DNS services on the MGMT port of an appliance:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.

Note: You must enable the MGMT port before modifying its port settings. See [Using the MGMT Port](#) on page 359.

2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, add the MGMT port to the Additional Ports and Addresses table as follows:
3. Click the Add icon and select **MGMT**.
Grid Manager adds a row for the MGMT port. For an HA pair, it adds two rows, one for each node.
4. Enter the following in the row of the MGMT port for a single Grid Master or independent appliance, and in the rows of the two nodes for an HA Grid Master or independent HA pair:
 - **Interface:** Displays the name of the interface. You cannot modify this.
 - **Address:** Type the IP address for the MGMT port, which must be in a different subnet from that of the LAN and HA ports.
 - **Subnet Mask:** For IPv4 only, specify an appropriate subnet mask for the number of management systems that you want to access the appliance through the MGMT port.
 - **Prefix Length:** For IPv6 only, specify the prefix length.
 - **Gateway:** Type the default gateway for the MGMT port. If you need to define any static routes for traffic originating from the MGMT port—such as SNMP traps, syslog events, and email notifications—destined for remote subnets beyond the immediate subnet, specify the IP address of this gateway in the route.

- **Port Settings:** Choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
 - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#) on page 348.
5. Click **Save & Close** to save your settings for the MGMT port.
 6. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
 7. In the **General** -> **Basic** tab of the *Member DNS Configuration* editor, do the following:
 - **Enable DNS service on MGMT port:** Select this check box.
 8. In the **General** -> **Advanced** tab, select one of the following from the **Send queries from** and the **Send notify messages and zone transfer requests from** drop-down lists:
 - **VIP:** The appliance uses the IP address of the HA port as the source for queries, notifies, and zone transfer requests.
 - **MGMT:** The appliance uses the IP address of the MGMT port as the source for queries, notifies, and zone transfer requests.
 - **LAN2:** The appliance uses the IP address of the LAN2 port as the source for queries, notifies, and zone transfer requests.
 - **Any:** The appliance chooses which port to use as the source for queries, notifies, and zone transfer requests.

The **Send queries from** drop-down list also includes loopback IP addresses that you configured. You can select a loopback address as the source for queries.
 9. Save the configuration and click **Restart** if it appears at the top of the screen.
- To see that the appliance now also serves DNS on the MGMT port:
1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid_member* check box.
 2. Expand the Toolbar and click **View** -> **View DNS Configuration**.
 3. Check that the IP address of the MGMT port appears in the address match list in the listen-on substatement.

ABOUT LIGHTS OUT MANAGEMENT

Infoblox LOM (Lights Out Management) is an implementation of the remote management and monitoring of Infoblox appliances that are LOM ready, such as the Trinzic 1410 and 2010 appliances.

The LOM feature is useful when you want to monitor your platforms remotely or consolidate your data centers. When you monitor your systems remotely, you can avoid issues such as overheating of a problematic system by remotely powering down the system. To conserve energy, you can also power up and down any systems based on service requirements.

You can enable LOM for the entire Grid and override the Grid settings for specific members. You can also configure LOM on independent appliances and HA pairs.

Note: You can configure LOM only on appliances that support LOM. This port automatically negotiates a speed of 100 Mbps. Devices connected to the LOM port should be configured to auto-negotiate and not have a fixed speed of 1000 Mbps.

LOM is disabled by default. Before you can configure LOM and remotely manage the appliance, ensure that the IPMI (Intelligent Platform Management Interface) port on your appliance is properly connected to the network. By default, IPMI uses UDP port 623. You can then enable LOM and add LOM users through the Infoblox GUI. When you add LOM users, you can assign them specific roles so they can perform only certain functions. When you add a LOM user, you can configure the user to be an “operator” or a “user” depending on the functions you want the user to perform. An operator can access an appliance remotely and perform the following functions:

- Access the serial console
- Reset the appliance
- Power up and down the appliance
- Monitor system status, such as CPU usage and system temperature

A user role can only monitor system status. Users with this role cannot perform any other functions remotely.

After you set up and configure your appliance, perform the following tasks through Grid Manager to enable LOM and set up LOM users:

1. Enable LOM for the Grid or members that support IPMI, as described in [Enabling LOM](#) on page 367.
2. Add LOM users based on your organizational needs, as described in [Adding LOM User Accounts](#) on page 368.
3. Configure the IPMI network interface on the appliance, as described in [Configuring the IPMI Network Interface](#) on page 368.
4. After you have configured LOM and set up the IPMI interface, install a utility such as IPMITool on your Linux management system. For information about IPMITool, visit the IPMITool web site at <http://ipmitool.sourceforge.net>. For the most commonly used commands and examples, see [IPMI Commands and Examples](#) on page 369.

You can also do the following from Grid Manager after you configure LOM:

- Enable and disable LOM for the Grid or members, as described in [Enabling LOM](#) on page 367.
- Modify LOM settings, as described in [Modifying LOM Settings](#) on page 369.
- View LOM users, as described in [Modifying LOM Settings](#) on page 369.

Enabling LOM

Before you can add LOM users and manage Infoblox appliances remotely, you must enable LOM. When LOM is configured for the entire Grid, all members inherit the Grid settings. You can also override the Grid settings for specific members. For an HA pair, you can configure LOM on the node that supports IPMI.

To enable and disable LOM:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
Independent appliance: From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **LOM** tab, complete the following:
 - **Enable Lights Out Management:** LOM is disabled by default. Select this check box to enable LOM. When LOM is enabled or disabled for the Grid, all members inherit the same setting.
3. Save the configuration.

Adding LOM User Accounts

You can add up to eight LOM user accounts. Admins must use the configured user name and password to remotely log in to the appliance.

Note that when you add LOM user accounts at the Grid level, all members inherit them. You can configure user accounts specific to a member by overriding the Grid accounts. When you click **Override** to modify the inherited Grid accounts, the appliance creates copies of the Grid level user accounts and saves them at the member level. These are new accounts at the member level and do not affect the Grid accounts or any accounts configured on other Grid members. You can also reset member accounts to the Grid accounts by clicking **Inherit**. When you do that however, all changes you previously made to the member accounts are lost.

To add a LOM user account:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.

Independent appliance: From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.

Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the **LOM** tab, complete the following:
 - **User Accounts:** Click the Add icon and complete the following:
 - **Name:** Enter the name of the LOM user account.
 - **Password:** Enter the password for the LOM user account.
 - **Confirm Password:** Enter the password again.
 - **Role:** From the drop-down list, select the role for the LOM user account. **Operator** allows users to perform all supported LOM related functions. **User** allows admins to only monitor system sensors such as temperature and CPU usage.
 - **Disable:** Select this to deactivate the user account but keep a user profile.
 - Click **Add** to add the new user account.
3. Save the configuration.

Configuring the IPMI Network Interface

You must configure the IPMI network interface before you can access the appliance remotely. To configure the IPMI network interface:

1. **Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the **LOM** tab, complete the following in the Network Configuration table:
 - **Address:** Enter the IPMI interface address here.
 - **Subnet Mask:** Enter the subnet mask for the IPMI interface.
 - **Gateway:** Enter the gateway address for the IPMI interface.
3. Save the configuration.

Modifying LOM Settings

To modify LOM settings:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
Independent appliance: From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Modify the following:
 - **Enable Lights Out Management:** LOM is disabled by default. Select this check box to enable LOM. When you enable or disable this for the Grid, all members inherit the same setting.
 - **Network Configuration:** Click the fields in the table to modify the IPv4 address, subnet mask, and gateway address for the IPMI interface. For an HA pair, the appliance displays information only for the nodes that support IPMI. Enter the information for the following fields: **Address**, **Subnet Mask**, and **Gateway**. The **Node** and **LAN Address** fields are read-only, and you cannot modify them. The LAN address is the IPMI interface address.
 - **User Accounts:** Click the Add icon to add new LOM users. You can also select an existing LOM user and click the Edit icon to modify the user settings, as described in [Adding LOM User Accounts](#).
3. Save the configuration.

Viewing LOM Users

To view information about LOM users:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
Independent appliance: From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.
Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **LOM** tab, Grid Manager displays the following information for each LOM user:
 - **Name:** The name of the LOM user.
 - **Role:** The user role to which the LOM user was assigned. This can be **Operator** or **User**.
 - **Disabled:** Indicates whether the LOM user account is disabled or not. When a LOM user account is disabled, the user cannot access the appliance remotely.

IPMI Commands and Examples

This section describes some of the most commonly used IPMITool commands and examples. For more information about the IPMI commands and usage, visit the IPMITool web site at <http://ipmitool.sourceforge.net>.

To use IPMI commands, complete the following:

1. Ensure that you have properly enabled and configured LOM and the IPMI network interface.
2. Install IPMITool on a Linux management system. For information, visit the IPMITool web site at <http://ipmitool.sourceforge.net>.
3. Access IPMITool and enter an IPMI command to perform a specific task.

The appliance displays the corresponding output.

Following are some of the most commonly used IPMI commands and their sample outputs. Note that command outputs vary by appliances. The following sample commands were performed on a Trinzic 1410 appliance. All sample commands in this section use the following syntax:

```
ipmitool -H <LOMIPAddress> -U username -P password -L [OPERATOR/USER] -I lanplus
<supported commands>
```

Command to be Used with Caution

```
power reset variant
```

Caution: Using this command has the same effect as pulling the power cord off the appliance.

Checking Power Status with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus power status
```

Command output:

```
Chassis Power is on
```

Checking Various Sensors [Temperature, Voltage, FANS, Physical Security, Power supply, OEM] with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus sensor
```

Command output:

```
System Temp | 23.000 | degrees C | ok | -9.000 | -7.000 | -5.000 | 75.000 | 77.000 |
79.000
CPU Temp | 0x0 | discrete | 0x0000| na | na | na | na | na | na
FAN 1 | 10390.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 | 29815.000 |
30370.000
FAN 2 | na | RPM | na | na | na | na | na | na | na
FAN 3 | 9835.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 | 29815.000 |
30370.000
FAN 4 | 11870.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 | 29815.000 |
30370.000
FAN 5 | 10390.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 | 29815.000 |
30370.000
CPU Vcore | 0.832 | Volts | ok | 0.640 | 0.664 | 0.688 | 1.344 | 1.408 | 1.472
+3.3VCC | 3.264 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
+12 V | 11.978 | Volts | ok | 10.494 | 10.600 | 10.706 | 13.091 | 13.197 | 13.303
CPU DIMM | 1.528 | Volts | ok | 1.152 | 1.216 | 1.280 | 1.760 | 1.776 | 1.792
+5 V | 5.088 | Volts | ok | 4.096 | 4.320 | 4.576 | 5.344 | 5.600 | 5.632
-12 V | -12.486 | Volts | ok | -13.844 | -13.650 | -13.456 | -10.934 | -10.740 | -10.546
VBAT | 3.120 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
+3.3VSB | 3.264 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
AVCC | 3.264 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
Chassis Intru | 0x0 | discrete | 0x0000| na | na | na | na | na | na
PS Status | 0x1 | discrete | 0x01ff| na | na | na | na | na | na
```

Printing System Event Log with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus sel list
```

Command output: The appliance displays all event log entries (if any)

Getting FRU Information with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus fru
```

Command output:

```
FRU Device Description : Builtin FRU Device (ID 0)
Board Mfg Date : Sun Dec 31 15:00:00 1995
Board Mfg : Supermicro
Board Serial :
Product Serial :
```

Powering Off the Appliance with Operator Role

Command:

```
ipmitool -H 10.37.2.70 -U operator -P infoblox -L OPERATOR -I lanplus power off
```

Command output:

```
Chassis Power Control: Down/Off
```

Powering On the Appliance with Operator Role

Command:

```
ipmitool -H 10.37.2.70 -U operator -P infoblox -L OPERATOR -I lanplus power on
```

Command output:

```
Chassis Power Control: Up/On
```

Activating the Serial Console Port using Operator role

Command:

```
ipmitool -H 10.37.2.70 -U operator -P infoblox -L OPERATOR -I lanplus sol activate
```

Command output:

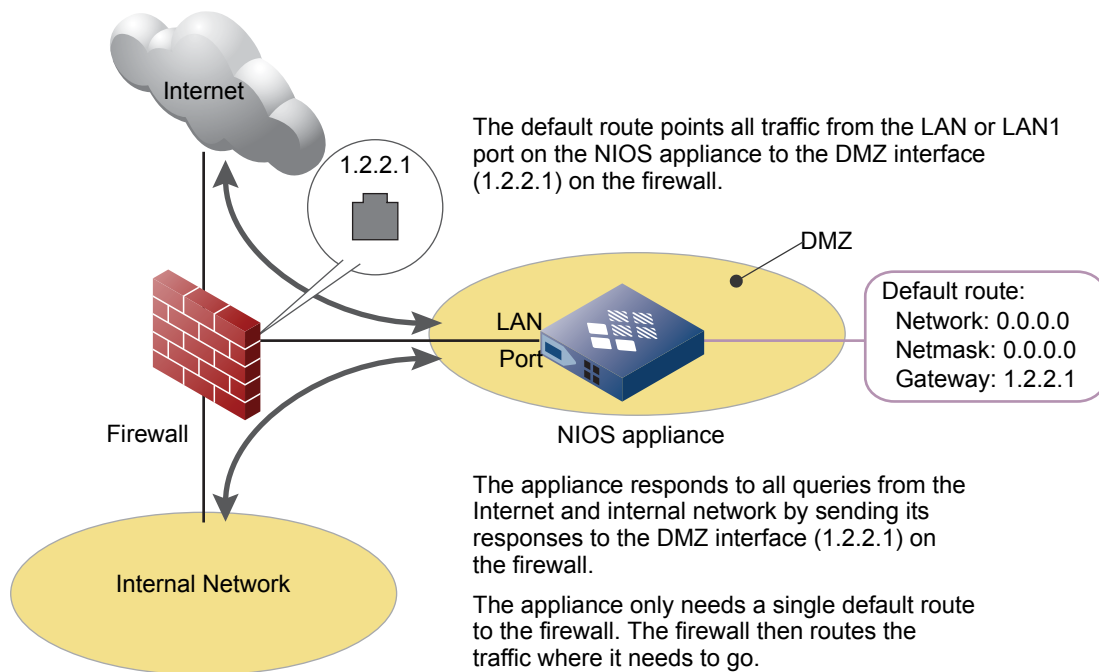
```
[SOL Session operational. Use ~? for help]
login: admin
password:
Infoblox NIOS Release 6.4.0-163715 (64bit)
Copyright (c) 1999-2012 Infoblox Inc. All Rights Reserved.
type 'help' for more information
Infoblox >
```

SETTING STATIC ROUTES

When you put the NIOS appliance on a segment of the network where there is a single path to and from it, a single default route is sufficient. For example, in [Figure 7.10](#) on page 372, the appliance is in the DMZ behind a firewall and connects to the rest of the network through the DMZ interface on the firewall. For example, when hosts send DNS queries from the Internet and the internal network to the appliance and when the appliance replies to those hosts, the firewall takes care of all the routing.

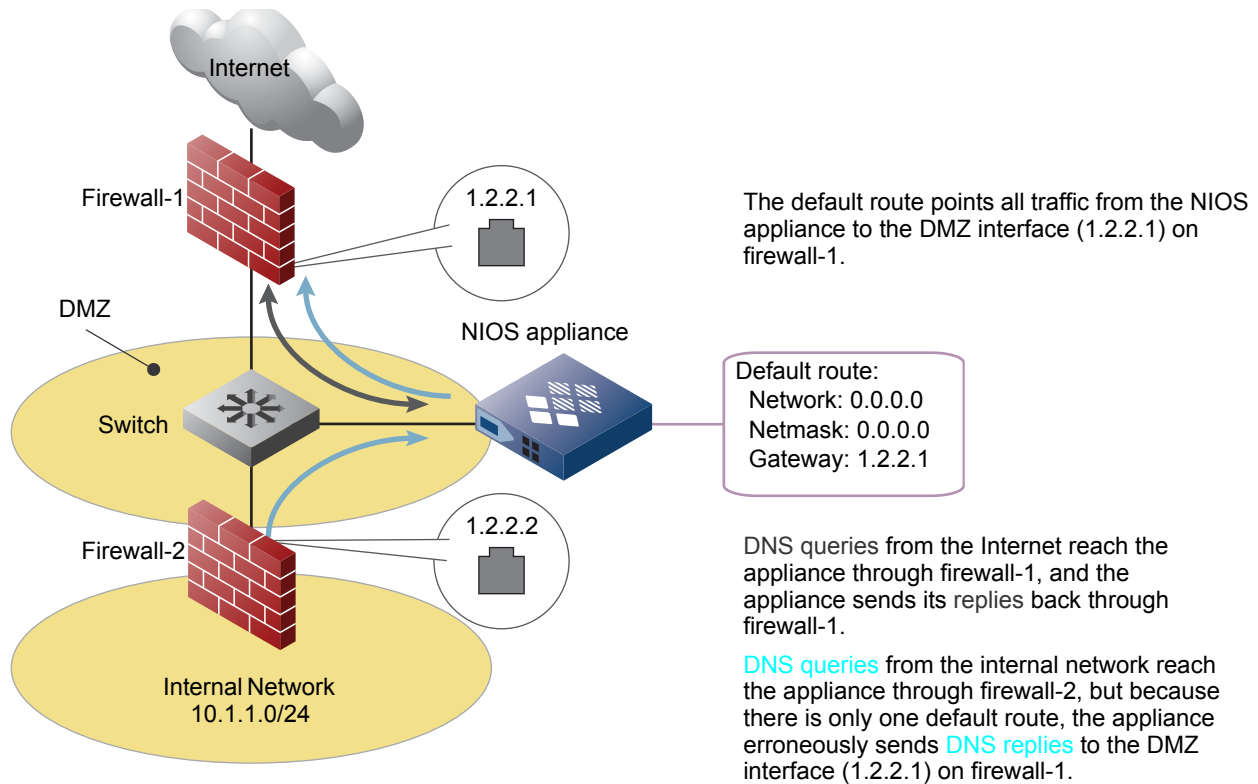
Note: This feature is not supported on vNIOS Grid members for Riverbed.

Figure 7.10 Single Default Route



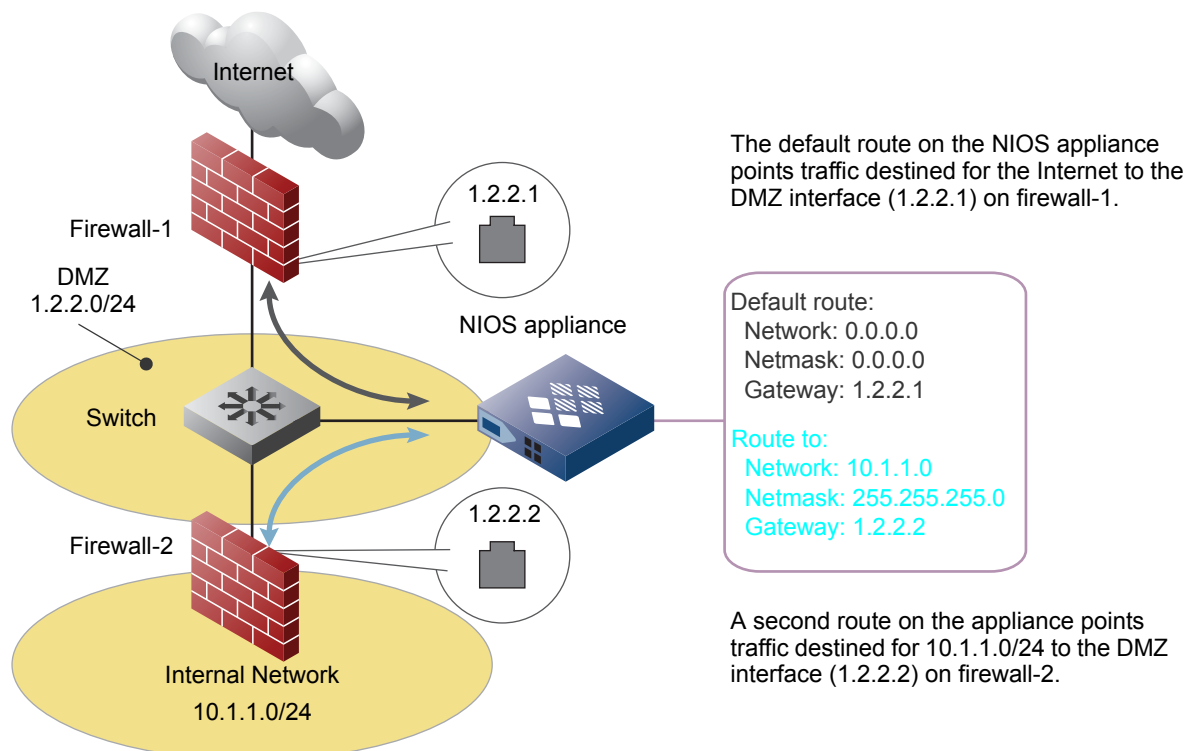
When the NIOS appliance is on a segment of the network where there are multiple gateways through which traffic to and from the appliance can flow, a single default route is insufficient. For an example, see [Figure 7.11](#).

Figure 7.11 Erroneously Routed DNS Replies



To resolve the problem illustrated in [Figure 7.11](#) on page 373, add a second route pointing traffic destined for 10.1.1.0/24 to use the gateway with IP address 1.2.2.2 on firewall-2. This is shown in [Figure 7.12](#).

Figure 7.12 Properly Routed DNS Replies

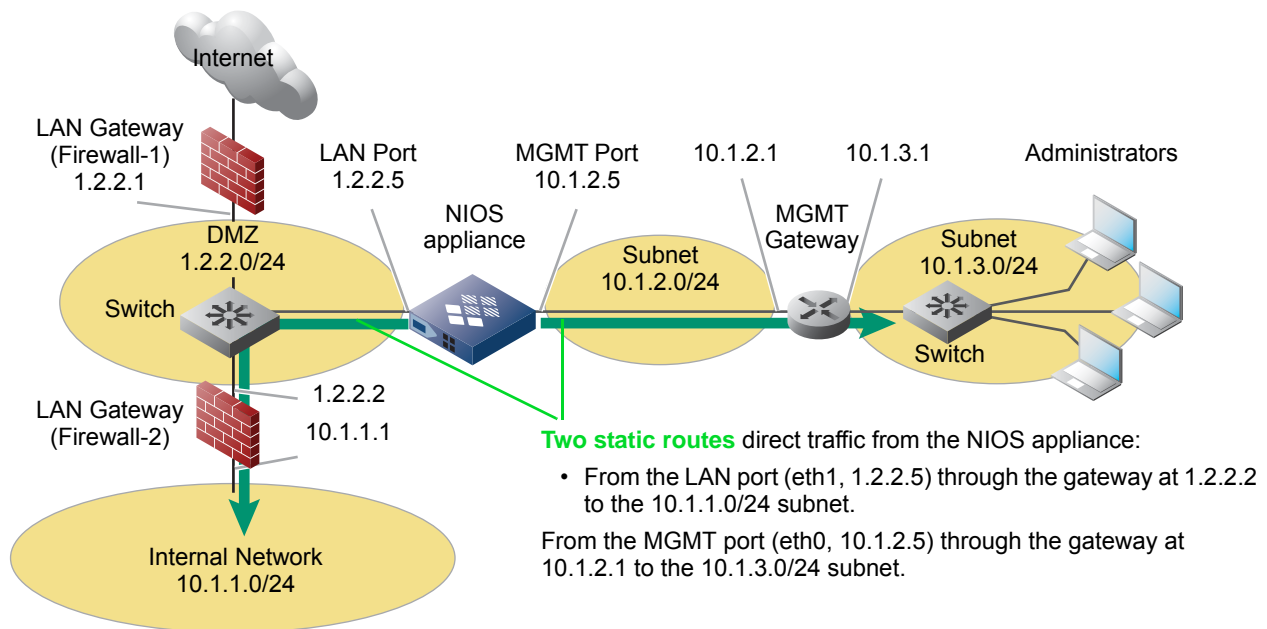


Whenever you want the NIOS appliance to send traffic through a gateway other than the default gateway, you need to define a separate route. Then, when the appliance performs a route lookup, it chooses the route that most completely matches the destination IP address in the packet header.

When you enable the MGMT port, the gateway you reference in a static route determines which port the NIOS appliance uses when directing traffic to a specified destination.

- If a route definition references a gateway that is in the same subnet as the IP and VIP addresses of the LAN (or LAN1) and HA ports, the NIOS appliance uses the LAN (or LAN1) or HA port when directing traffic to that gateway.
- If a route definition references a gateway that is in the same subnet as the MGMT port, the NIOS appliance uses the MGMT port when directing traffic to that gateway.

Figure 7.13 Static Routes for the LAN and MGMT Ports



Route Tables on the NIOS appliance

```
From LAN:
1.2.2.0/24 dev eth1 scope link
10.1.1.0/24 via 1.2.2.2 dev eth1
default via 1.2.2.1 dev eth1

From MGMT:
10.1.2.0/24 dev eth0 scope link
10.1.3.0/24 via 10.1.2.1 dev eth0
default via 10.1.2.1 dev eth0

From all:
10.1.1.0/24 via 1.2.2.2 dev eth1
10.1.3.0/24 via 10.1.2.1 dev eth0
1.2.2.0/24 dev eth1 proto kernel scope link src 1.2.2.5
10.1.2.0/24 dev eth0 proto kernel scope link src 10.1.2.5
default via 1.2.2.1 dev eth1
```

Note: There is a route table for each port as well as a comprehensive route table. For an HA pair, the LAN port route table is duplicated for the HA port.

In this illustration, the static routes are shown in green.

The need for routes can apply to any type of traffic that originates from the appliance, such as DNS replies, DHCP messages, SNMP traps, ICMP echo replies, Infoblox GUI management, and Grid communications.

To set a static route, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.

2. In the **Network** -> **Advanced** tab of the *Grid Member Properties* editor, click the Add icon for the **IPv4 Static Routes** table, and then enter the following:
 - **Network Address:** Type the address and netmask of the remote network to which the NIOS appliance routes traffic.
 - **Gateway Address:** Type the IP address of the gateway on the local subnet through which the NIOS appliance directs traffic to reach the remote network. The gateway address must meet the following requirements:
 - It must belong to a working gateway router or gateway switch.
 - It must be in the same subnet as the NIOS appliance.

Note: Consult your network administrator before specifying the gateway address for a static route on the appliance. Specifying an invalid gateway address can cause problems, such as packets being dropped or sent to an incorrect address.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Defining IPv6 Static Routes

Note: IPv6 addressing and connectivity requires an IPv4 address configuration on the same interface. Ensure that IPv4 configuration is in place when attempting to define a default route through any interface on the NIOS appliance.

Principles and applications related to IPv4 static routing in this section apply equally to IPv6. In [Figure 7.14](#), a NIOS appliance supports both IPv4 and IPv6 on its LAN1 port. IPv6 is routed to the internal network while the default IPv4 route remains to the outbound 10.2.2.1 address.

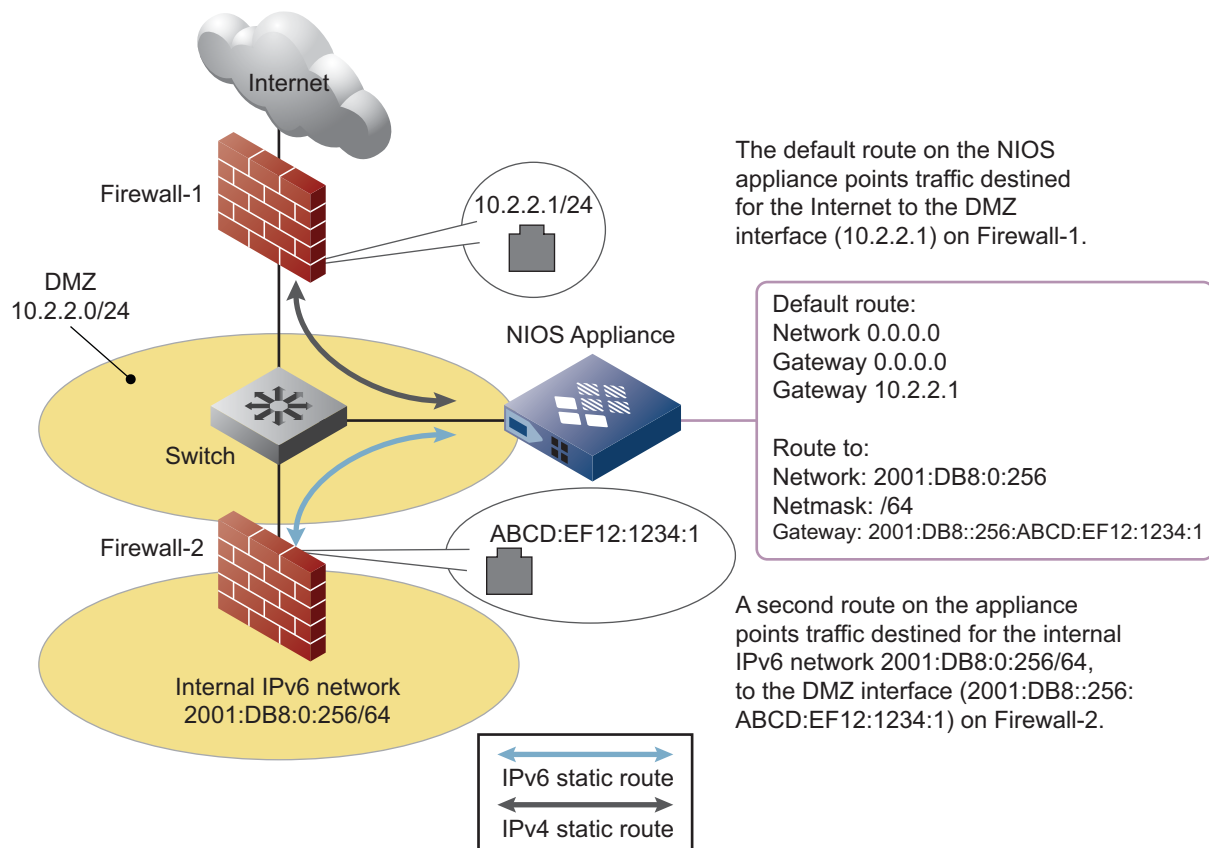
You can use prefix notation to enter an IPv6 network address; the full 128-bit gateway value must be entered. To set an IPv6 static route, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the **Network** -> **Advanced** tab of the *Grid Member Properties* editor, click the Add icon for the **IPv6 Static Routes** table, and then enter the following:
 - **Network Address:** Type the prefix and prefix length of the remote network to which the NIOS appliance routes traffic. As an example: 2001:DB8::256:/64. The double colon is required at the end of the prefix. NIOS performs validity checks on the address while it is being entered.
 - **Gateway Address:** Type the IP address of the gateway on the local subnet through which the NIOS appliance directs traffic to reach the remote network. As an example: 2001:DB8::256:ABCD:EF12:1234:1. The gateway address must meet the following requirements:
 - It must belong to a working gateway router or gateway switch.
 - It must be in the same subnet as one of the interfaces of the NIOS appliance.
 - The gateway address cannot be the same value as that for the VIP.

Note: Consult your network administrator before specifying the gateway address for a static route on the appliance. Specifying an invalid gateway address can cause problems, such as packets being dropped or sent to an incorrect address.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Figure 7.14 Static Routes for IPv6 Traffic



ENABLING DNS RESOLUTION

You can specify a network server to perform domain name queries and specify up to two name servers for resolving a DNS name. You can specify the IP address of a preferred name server and that of an alternate name server, plus use a search list for performing partial name resolution.

To enable DNS resolution for a Grid or for an independent appliance or HA pair:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**.

Member: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the *Grid Properties* or *Member Properties* editor, select the **DNS Resolver** tab, and then enter the following:

- **Enable DNS Resolver:** Select the check box to enable the NIOS appliance to send DNS queries to the preferred or alternate name servers whose IP addresses you specify in the following fields.

Click the Add icon and enter the IP addresses (IPv4 or IPv6) of the servers to which the appliance sends queries. The appliance attempts to send queries to the servers in the order they are listed if it does not receive a response from a listed name server. To move a server up or down on the list, select it and drag it to its new location or click the up and down arrows.

- **Search List:** You can define a group of domain names that the NIOS appliance can add to partial queries that do not specify a domain name. For example, if you define a RADIUS authentication home server as "as1", and you list "corp100.com" and "hq.corp100.com" in the domain group list, then the NIOS appliance sends a query for "as1.corp100.com" and another query for "as1.hq.corp100.com" to the preferred or alternate name server. To specify domain names containing IDNs, manually convert it into punycode and specify domain names in punycode.

To add a domain name, click the Add icon and type a domain name in the Search List field. To remove a domain name from the group, select it, and then click **Delete**.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

MANAGING LICENSES

Licenses come pre-installed on a NIOS appliance according to the software packages you ordered at the time of purchase. If you wish to upgrade an existing appliance with the Grid license, you must contact Infoblox Technical Support and follow the procedures in [Viewing Licenses](#) on page 378.

On a vNIOS virtual appliance, you can install licenses when you deploy the appliance. You must install both the Grid and vNIOS licenses on a vNIOS appliance for it to join a Grid. You can transfer the valid licenses of a vNIOS appliance on VMware from one ESX/ESXi server to another. For more information, refer to the *Infoblox Installation Guide for vNIOS Software on VMware*.

There are three types of licenses:

- **Maintenance licenses** – Examples: NIOS and Grid (or Keystone) maintenance licenses. The duration of maintenance licenses are one, two, or three years. You can obtain these licenses from your Infoblox sales representative.
- **Service licenses** – Examples: DNS and DHCP licenses. These are permanent licenses. You can obtain these licenses from your Infoblox sales representative. Note that the DNS and DHCP services are disabled by default. Once you have obtained licenses for these services, start the services after you complete the configuration.
- **Temporary licenses** – You can enable one of several sets of temporary service licenses through the CLI command `set temp_license`.

Before a maintenance license or a temporary license expires, an expiration warning appears during the GUI login process. The warning reappears during each login until you renew the license. To renew a license, contact Infoblox Technical Support.

Obtaining and Adding Licenses

A valid Grid license is required for deploying a Grid with NIOS and vNIOS appliances. You can upgrade existing independent NIOS and vNIOS appliances to use a Grid license and then add them to a Grid. To upgrade your licenses, contact Infoblox Technical Support.

When you receive a new license key, it is in CSV (comma separated values) format with the following information: serial number, hardware ID, license type, end date, and license string. You can either upload the file to the appliance or copy the information and paste it in the text field of the **Licenses** tab of the Infoblox GUI. Note that you must copy the entire string—serial number, hardware ID, license type, end date, and license string—and save it to the text field.

To add a license:

1. From the **Grid** tab, select the **Licenses** tab and click the Add icon.
2. Do one of the following:
 - **Upload License File:** Click **Select File** and navigate to the license file.
 - **Paste License(s):** Paste the license key in this text field. You must paste the entire string in CSV format: serial number, hardware ID, license type, end date, and license string. If you are pasting multiple licenses, start each string on a new line.

3. Click **Save License(s)**.

The appliance validates the license and adds it to the table. Close the browser window and log in to the Infoblox GUI. If you are activating licenses for an HA pair, you must follow this procedure for both nodes.

Note: To transfer licenses between vNIOs on VMware appliances, refer to the *Infoblox Installation Guide for vNIOs Software on VMware*.

Obtaining Temporary Licenses

You can use the CLI command `set temp_license` to generate and install temporary licenses. This can provide licensed features and functionality for the interim, while you wait for your permanent license to arrive.

To generate a temporary license:

1. Log in to the NIOS appliance through a remote console window. For more information on how to open a remote console window, refer to the *Infoblox CLI Guide*.
2. At the Infoblox command prompt, enter `set temp_license`. The appliance lists the available licenses, and you select those you need.
3. Enter the number of the licenses you want to install.
4. Confirm the selection when prompted, and the following message appears:

```
Temporary license is installed.
```

Viewing Licenses

If the appliance is part of a Grid, you must log in to the Grid Master to view license information from Grid Manager. If the appliance is an independent appliance, log in to System Manager on the appliance. If you have transferred licenses from one vNIOs on VMware appliance to another, you can view information about the new and replaced licenses.

In this panel, you can use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

To view license information on a NIOS or vNIOs appliance:

1. Log in to Grid Manager on the Grid Master or System Manager on the independent appliance.
2. Select the **Grid** or **System** tab -> **Licenses** tab. The appliance displays the following information:
 - **Name:** The name of the appliance.
 - **HA:** Indicates whether the appliance is an HA pair.
 - **Address:** The IP address of the appliance.
 - **Hardware ID:** The unique hardware ID of the appliance. The ID is highlighted in red if the license on the appliance was removed.
 - **Type:** The type of license installed.
 - **Type Context:** Depending on the license type, this field displays the attribute, such as **Model** and **Connector**, that the license controls. For example, when you purchase an IF-MAP Federation license with a limited number of federation connectors, this field displays **Connector**. This field is blank if the license does not control any attribute type. This field can display one of the following:
 - **Leases:** Indicates that this DHCP license supports a specific number of DHCP leases. The number of leases supported is displayed in the **Type Details** field.
 - **Model:** Indicates that this vNIOs license supports a specific vNIOs virtual appliance model. This includes the supported virtual appliance for the IF-MAP Starter Kit. The model supported is displayed in the **Type Details** field.

- **Connector:** Indicates that this IF-MAP Federation license supports a specific number of IF-MAP federation connectors. The number of federation connectors supported is displayed in the **Type Details** field.
- **Tier:** Indicates various levels of performance limits on the DNS cache acceleration license of the IB-4030 appliance. This is only applicable to the IB-4030 appliance.
- **Type Details:** Information about the attribute type that the license monitors. This field can display the following information for each attribute:
 - **Leases:** The number of DHCP leases that the DHCP license supports.
 - **Model:** The model of the vNIOS virtual appliance, such as IB-VM-550 or IB-VM-1050.
 - **Connector:** The number of federation connectors that the IF-MAP Federation license supports.
 - **Tier:** The performance limit value of an IB-4030 appliance with DNS Cache Acceleration, such as Tier-1 for full capacity (up to 1M qps), Tier-2 for high (up to 600K qps), and Tier-3 for base (up to 300K qps) performance limits. This is only applicable to the IB-4030 appliance.
- **Expiration:** The expiration date of the license.
- **Replaced Hardware ID:** The hardware ID of the appliance whose license was removed.

Backing Up Licenses

You can back up the licenses installed on the appliance, in case you need to re-install them at a later time. Infoblox recommends backing up the licenses before removing any of them.

When you back up the licenses, Grid Manager creates a CSV file that lists the following information for each license: serial number, hardware ID, license type, end date, license string.

To back up licenses:

1. From the **Grid** tab, select the **Licenses** tab.
2. Click the Backup Licenses icon in the toolbar.

Grid Manager generates a CSV file that contains all the licenses. You can then open the file or save it to a specified location.

Removing Licenses

You can remove licenses and reset a NIOS appliance to its factory default settings. For example, if you have a NIOS appliance running the DNSone package with the Grid upgrade, but you want to use it as an independent appliance, you can remove the Grid license. Infoblox recommends that you back up licenses before removing them, in case you decide to re-install them at a future time.

Note: Exercise caution when removing licenses; you may render an appliance unusable by removing the wrong license. Other feature sets may be affected once you remove a license; for example, removing licensing for DNS and DHCP services will also disable task packs in the **NIOS Dashboards → Tasks** page.

To remove a license:

1. From the **Grid** tab, select the **Licenses** tab.
2. Select the license and click the Delete icon.
Check the license that you are about to remove. Note that removing the wrong license can render an appliance unusable.
3. Click **Yes** when the confirmation dialog appears.
4. Close the browser window and log in to the Infoblox GUI.

MANAGING THE ORDER OF MATCH LISTS

When you configure certain DNS and DHCP functions, you can create match lists that the appliance uses to filter specific IP addresses for specific operations. For example, you can create a DNS blackhole list for including and excluding DNS traffic to certain IP addresses, configure a list of IP addresses for allowing and denying DDNS updates, or define a Match Destinations list that identifies destination addresses and TSIG keys that are allowed access to a DNS view.

The appliance matches rules in these lists from top to bottom. Rules at the top always take precedence over those at the bottom. Therefore, ensure that you put the most specific rules at the top of the list, and then put the more general rules at the bottom. For example, when you add network 10.10.0.0/24 to a DNS blackhole list, all 256 IP addresses in the network are put on the blackhole list. To allow DNS traffic to the specific IP addresses 10.10.0.55 and 10.10.0.88, you must add these two addresses at the top of the blackhole list before the network address 10.10.0.0/24, and then set their permissions to “Exclude.” The same applies when you set up the list of clients for DDNS updates. If you want to deny DDNS updates from a specific client (10.0.0.99) and allow DDNS updates from all other clients in the 10.0.0.0/24 network, you must put 10.0.0.99 at the top of the list and configure the appliance to deny DDNS updates from this client. You then add network 10.0.0.0/24 for allowing DDNS updates from all other clients at the bottom of the list.

SHUTTING DOWN, REBOOTING, AND RESETTNG A NIOS APPLIANCE

To reboot and shut down a NIOS appliance, you can use Grid Manager or the Infoblox CLI. To reset a NIOS appliance, you must use the Infoblox CLI.

Rebooting a NIOS Appliance

You can reboot a single NIOS appliance, a single node in an HA pair, or both nodes in an HA pair.

To reboot a single NIOS appliance or one or both nodes in an HA pair:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box.
2. Expand the Toolbar and click **Control** -> **Reboot**.
 - For an HA pair, choose whether to boot one node (and which one) or both nodes, and then click **OK**.

Depending on the browser you use, Grid Manager may display a dialog box that indicates the system is unavailable during a restart or reboot.

To reboot a single NIOS appliance using the CLI:

1. Log in to the Infoblox CLI using a superuser account for the NIOS appliance that you intend to reboot.
2. Enter the following CLI command: **reboot**

Shutting Down a NIOS Appliance

Under normal circumstances, you do not need to turn off or shut down a NIOS appliance. It is designed to operate continuously. However, if you want to turn off a NIOS appliance, use the GUI or the CLI to shut down the appliance, instead of just turning off the power switch. Before shutting down a remote appliance, make sure you can restart it. You cannot restart the system using the GUI.

Note: If there is a disruption in power when the NIOS appliance is operating, the NIOS appliance automatically reboots itself when power is restored.

To shut down a NIOS appliance:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box.
2. Expand the Toolbar and click **Control** -> **Shutdown**.

- For an HA pair, choose whether to shut down one node (and which one) or both nodes, and then click **OK**.

The NIOS appliance shuts down. The fans might continue to operate until the appliance cools down.

To shut down a NIOS appliance using the CLI:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: **shutdown**

Resetting a NIOS Appliance

There are three ways to reset a NIOS appliance:

- [Resetting the Database](#) on page 381
- [Resetting a NIOS Appliance to Factory Settings](#) on page 381
- [Resetting the NIOS Appliance to Factory Settings and Removing Licenses](#) on page 382

You can perform these functions only through the CLI.

Resetting the Database

You can reset the database if you lose the administrator account and password or if you want to clear the database but preserve the log files to diagnose a problem. This function removes the configuration files, and the DNS and DHCP data from the appliance database. During this procedure, you are given the option to preserve the network settings of the appliance, which are the IP address and subnet mask, the IP address of the gateway, the host name, and the remote access setting.

To reset the database:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: **reset database**

The appliance then displays a message similar to the following:

The following network settings can be restored after reset:

IP Address: 10.1.1.10

Subnet Mask: 255.255.255.0

Gateway: 10.1.1.1

Host Name: ns1.corp100.com

Remote Console Access: true

The entire database will be erased.

Do you wish to preserve basic network settings? (y or n)

3. Press the **Y** key to preserve the network settings or the **N** key to return the network settings to their default values (192.168.1.2, 255.255.255.0, 192.168.1.1).

Resetting a NIOS Appliance to Factory Settings

You can reset a NIOS appliance to its original factory settings. This removes the database, network settings, logs, and configuration files. Then, it reboots with its factory settings, which are the default user name and password, and default network settings. When you perform this procedure, the appliance does not give you the option to preserve your network settings.

Note: If you have previously imported HTTPS certificates, the appliance regenerates the certificates and replaces them.

To reset the NIOS appliance to its factory settings:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: **reset all**

Resetting the NIOS Appliance to Factory Settings and Removing Licenses

You can also reset a NIOS appliance to its original factory settings and remove all the licenses installed on the appliance. This removes the database, network settings, logs, configuration files, and licenses. The appliance then reboots with its factory settings, which are the default user name and password, and default network settings.

Note: If you have previously imported HTTPS certificates, the NIOS appliance regenerates the certificates and replaces them.

To reset the NIOS appliance to its factory settings and remove all its licenses:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: `reset all licenses`.

MANAGING THE DISK SUBSYSTEM ON THE INFOBLOX-2000-A AND -4010


Among its many features, the Infoblox-2000-A and Infoblox-4010 use a RAID (Redundant Array of Independent Disks) 10 array to provide the optimum mix of high database performance and redundant data storage with recovery features in the event of disk failures. The disk array is completely self managed. There are no maintenance or special procedures required to service the disk subsystem.

Caution: Never remove more than one disk at a time from the array. Removing two or more disks at once can cause an array failure and result in an unrecoverable condition. You should replace only one disk at a time, using a replacement disk from Infoblox. For information, see [Replacing a Failed Disk Drive](#) on page 384.

About RAID 10

RAID 10 (or sometimes called RAID 1+0) uses a minimum of four disk drives to create a RAID 0 array from two RAID 1 arrays, as shown in [Figure 7.15](#). It uses mirroring and striping to form a stripe of mirrored subsets. This means that the array combines—or stripes—multiple disk drives, creating a single logical volume (RAID 0). RAID 10 combines the high performance of RAID 0 and the high fault tolerance of RAID 1. Striping disk drives improves database write performance over a single disk drive for large databases. The disks are also mirrored (RAID 1), so that each disk in the logical volume is fully redundant.

Figure 7.15 RAID 10 Array Configuration



When evaluating a fault on the Infoblox-2000-A or -4010, it is best to think of the disk subsystem as a single, integrated unit with four components, rather than four independent disk drives. For information, see [Evaluating the Status of the Disk Subsystem](#) on page 383.

Evaluating the Status of the Disk Subsystem

You can monitor the disk subsystem through the Infoblox Grid Manager GUI, the scrolling front panel LCD display, and four front panel LEDs next to the disk drives. In addition, you can monitor the disk status by using the CLI command `show hardware_status`. The following example displays the status of an Infoblox-2000-A or 4010 using the command:

```
Infoblox > show hardware_status
POWER:      Power OK
Fan1:       7258 RPM
Fan2:       6887 RPM
Fan3:       7258 RPM
CPU1_TEMP:  +20.0 C
CPU2_TEMP:  +24.0 C
SYS_TEMP:   +35 C




RAID_ARRAY: OPTIMAL
RAID_BATTERY: OK READY Yes 103 HOURS
```

The *Detailed Status* panel provides a detailed status report on the appliance and service operations. To see a detailed status report:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box -> Detailed Status icon in the table toolbar.

After displaying the *Detailed Status* panel, you can view the status of the selected Grid member. For more information on the *Detailed Status* panel, see [Viewing Status](#) on page 1004.

The RAID icons indicate the status of the RAID array on the Infoblox-2000-A and 4010.

Icon	Color	Meaning
	Green	The RAID array is in an optimal state.
	Yellow	A new disk was inserted and the RAID array is rebuilding.
	Red	The RAID array is degraded. At least one disk is not functioning properly. The GUI lists the disks that are online. Replace only the disks that are offline.

The appliance also displays the type of each disk. In the event of a disk failure, you must replace the failed disk with one that is qualified and shipped from Infoblox and has the same disk type as the rest of the disks in the array. The disk type of the Infoblox-2000-A can be one of the following:

- IB-Type 1: Infoblox supported disk type
- IB-Type 2: Infoblox supported disk type
- Unk: Unknown disk type that Infoblox does not support

Infoblox-4010 uses only the IB-Type 3 disk type. All disk drives in the array must have the same disk type for the array to function properly. You can have either IB-Type 1, IB-Type 2, or IB-Type-3, but you cannot mix both in the array. When you have a mismatched disk in the array, you must promptly replace the disk with a replacement disk from Infoblox to avoid operational issues.

Disk Drive Front Panel LEDs

The disk drives of the Infoblox-2000-A are located on the right side of the appliance front panel. To the right of each drive, there is an LED that displays the status of each drive. The front panel LCD scrolls and displays the disk array status every 20 seconds. [Table 7.7](#) lists the disk drive LEDs.

Table 7.7 Infoblox-2000-A Disk Drive LEDs

LED Color	Condition	Action
Green	Disk operating normally	None
Yellow	Disk read/write activity	Disk is functioning normally or is synchronizing if recently inserted.
Dark	Disk has failed or not inserted	Verify the failure in the GUI or CLI. Remove the disk and replace with a functional disk drive. Note that the drive rebuilds with its twin.

The disk drives of the Infoblox-4010 are located on the appliance front panel. To the right of each drive, two LEDs display connection and activity status. [Table 7.8](#) lists the disk drive LED combinations and the states they represent.

Table 7.8 Infoblox-4010 Disk Drive LED Combinations

Online/Activity LED (Green)	Fault/UID LED (Amber/Blue)	Description
On, off, or blinking	Alternating amber and blue	The drive has failed, or it has received a predictive failure alert; it also has been selected by a management application.
On, off, or blinking	Steadily blue	The drive is operating normally.
On	Amber, blinking regularly (1 Hz)	The drive has received a predictive failure alert. Replace the drive as soon as possible.
On	Off	The drive is online but it is not currently active.
Blinking regularly (1Hz)	Off	Do not remove the drive. The drive is rebuilding. Removing the drive may terminate the current operation and cause data loss.
Blinking irregularly	Amber, blinking regularly (1 Hz)	The drive is active, but it has received a predictive failure alert. Replace the drive as soon as possible.
Blinking irregularly	Off	The drive is active and operating normally.
Off	Steadily amber	A critical fault condition has been identified for this drive, and the controller has placed it offline. Replace the drive as soon as possible.
Off	Amber, blinking regularly (1 Hz)	The drive has received a predictive failure alert. Replace the drive as soon as possible.
Off	Off	The drive is offline, a spare, or not configured as part of an array.

Replacing a Failed Disk Drive

The Infoblox-2000-A and -4010 are designed to provide continuous operation in the event of a failed disk. Replace an original RAID disk only when there is a disk failure. Hot-swapping a disk drive is a simple process that does not require issuing commands or a GUI operation.

When you replace a failed disk, you must replace it with an Infoblox supplied disk. To ensure that you receive the correct replacement disk, report the disk type or part number of the failed disk. The appliance displays the disk type in the *Detailed Status* panel, and the Infoblox part number is printed on the disk. Installing disks that are not qualified and shipped from Infoblox could cause failures in the appliance.

To replace a disk drive, follow this procedure:

1. Identify and verify the failed drive via the Grid Manager, front panel LCD, or CLI.
2. Make sure you have identified the correct drive.

Note: Do not remove a correctly functioning drive.

3. Push in the latch for the drive and pull the release lever out towards you.
4. When the drive disengages, wait about 30 seconds for the disk to completely stop spinning.
5. Slide it out of the slot.

Replacement drives are shipped as a complete unit, ready to insert into the appliance. There is no preparation required. To install a replacement drive, follow this procedure:

1. Insert the replacement drive into the drive bay slot.
2. Gently slide the drive into place. When you feel the release lever engage, continue applying gentle pressure to the drive while pushing the release lever towards the appliance.
3. The release lever locks into place and the LED next to the disk drive lights up. Note that if the alarm buzzer is sounding, it automatically turns off about 20 seconds after the drive is inserted.
4. The disk drive automatically goes into rebuild mode.

Disk Array Guidelines

Infoblox has designed the disk array to be completely self managed. There are no maintenance procedures required for a normally functioning disk array. Mishandling the disk array can cause an unrecoverable error and result in a failed appliance. Infoblox highly recommends that you observe the following guidelines:

- Remove only one disk at a time. Do not remove two or more disks from the appliance at the same time. Removing two or more disks at the same time might result in an appliance failure and require an RMA of the appliance. This rule applies to both powered and powered down appliances.
- If the status of the array is degraded, remove the failed or failing disk drive only. Do not remove an optimally functioning drive.
- If your acceptance procedure requires a test of the RAID hot swap feature, remove only one disk drive at a time. You can remove a second disk only after you replace the first disk and the array completes its rebuilding process.
- Do not remove a disk drive if the array is rebuilding. This could result in an appliance failure. Verify the status of the array before removing a disk drive.
- Use the following procedure to remove a spinning disk:
 1. Unlatch and pull the disk about two cm (one inch) to disengage contact.
 2. Wait about 30 seconds for the disk to completely stop spinning.
 3. Remove the disk and handle it with care. Do not drop the disk or ship it loosely in a carton.
- You can hot swap a drive while the appliance remains in production.
- There are some conditions that may require powering down the appliance to replace a failed unit. This normally happens if the RAID controller detects an error that could damage the array. If you insert a replacement drive into a live array and the controller doesn't recognize the drive, power down the appliance.
- If you inadvertently remove the wrong disk drive, do not immediately remove the disk drive that you originally intended to remove. Verify the status of the array and replace the disk drive that you removed earlier before removing another drive. Removing a second drive could render the appliance inoperable.

- Older appliances have an audio alarm buzzer that sounds if a drive fails. The alarm automatically stops about 20 seconds after a functional disk has been inserted into the array.
- All disks in the RAID array should have the same disk type for the array to function properly.
- In the unlikely event that two disk drives fail simultaneously and the appliance is still operational, remove and replace the failed disk drives one at a time.
- Rebuild time depends on a number of factors, such as the system load and Grid replication activities. On very busy appliances (over 90% utilization), the disk rebuild process can take as long as 40 hours. On a Grid Master serving a very large Grid, expect the rebuild process to take at least 24 hours.
- Replace a failed or mismatched disk only with a replacement disk shipped from Infoblox. When you request a replacement disk, report the disk type displayed in the *Detailed Status* panel of the GUI or the Infoblox part number on the disk.

RESTARTING SERVICES

Whenever you make a change (such as add a zone, a network, or a range), you click the **Restart** icon to restart services. You can restart the DNS and DHCP services after you make configuration changes. You can also specify a future restart time.

You can restart services at the Grid level or at the member level as described in:

- [Restarting Grid Services](#) on page 386
- [Restarting Member Services](#) on page 387

The following rules apply to superusers and limited-access users:

- You can cancel a schedule that you create to restart services. A superuser can cancel any scheduled restarts.
- Only superusers and administrators with read/write permission to all Grid members can schedule a Grid restart.
- When a superuser schedules a Grid restart, a limited-access user cannot schedule a member-level restart.
- Limited-access users cannot cancel a superuser's scheduled changes.
- Limited-access users cannot create or modify a schedule for a Grid member if a schedule for the member (created by another user) already exists.

The system writes every scheduled change action to the audit log as follows:

```
USER logon_id action service restart schedule 'schedule' on Grid (or member) Grid name
or member node id
```

For example:

```
USER jdoe insert service restart schedule '02/20/2007 01:30:00' on Grid Infoblox
USER jdoe deleted service restart schedule '02/22/2007 01:30:00' on node id 3
```

For more information on the audit log, see [Using the Audit Log](#) on page 1018.

Restarting Grid Services

Only superusers and administrators with read/write permission to all Grid members can schedule a Grid restart. You can restart services at the Grid level either simultaneously or sequentially, and also specify the restart time.

After you enter a specific date and time, the system schedules the restart at the specified time on each Grid member. To restart services at the Grid level:

1. From the **Data Management** tab, select the **DHCP**, **DNS**, or **Grid** tab, or select the **Administration** tab, and then click **Restart Services** from the Toolbar.
The *Restart Grid Services* wizard appears.
2. You can specify whether the member restarts services when necessary or you can force it to restart services. Select one of the following in the *Restart Grid Services* section:
 - **If needed:** Select this to restart all active DNS and DHCP services if there are any changes requiring a service restart.

- **Force restart services:** Select this to force all active services to restart, regardless of their state.
3. Select one of the following in the *Restart Services on all Members* section:
 - **Simultaneously:** Restarts the services on all of the members in a Grid at the same time.
 - **Sequentially:** Restarts the services on each Grid member according to the number of seconds you enter in the **Sequential every (seconds)** field. NIOS sets this option by default with a delay of 10 seconds. For example, if you enter every 10 seconds, the system restarts services on the first member, and 10 seconds later on the second member.

Impacted Members and Services: Click the Poll Members icon to display the affected members and services when the system restarts. Grid Manager displays the member names and one of the following for each service:

 - **YES:** The service is active and the system will restart the service upon execution of this task.
 - **NO:** The service will not restart unless the **Force restart services** option is selected.
 - **DISABLED:** The service is currently disabled.
 4. To schedule a service restart, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, complete the following:
 - **Now:** Restarts services upon clicking **Restart**.
 - **Later:** Enter the following information to schedule all Grid members to restart services at a certain date and time:
 - **Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. When you enter the time in a 24-hour format such as 23:00, Grid Manager displays 11:00:00 PM. You can also select a time from the drop-down list by clicking the time icon.
 - **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
 5. Click **Restart** to restart services immediately or click **Schedule Restart** to schedule the restart.

Restarting Member Services

The member restart time always supersedes the Grid restart time. If the member restart time is later than the Grid restart time, then the member restarts services at its scheduled time. If the member restart time is before the Grid restart time, then the member restarts services at its scheduled restart time, and again during the Grid restart time. To restart member services:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click **Restart Services** from the Toolbar.
or
From the **Grid** tab, select the **Grid Manager** tab, and then select a member check box.
2. You can specify whether the member should restart services when necessary or you can force it to restart services. Select one of the following in the *Restart Member Services* section:
 - **If needed:** Select this to restart all active DNS and DHCP services, if there are any changes requiring a service restart.
 - **Force restart services:** Select this to force all active services to restart, regardless of their state.

Impacted Services: This table displays the affected services when the system restarts. It can display one of the following for each service:

 - **YES:** The service is active and the system will restart the service upon execution of this task.
 - **NO:** The service will not restart unless the **Force restart services** option is selected.
 - **DISABLED:** The service is currently disabled.
3. To schedule a service restart, click the Schedule icon at the top of the editor. In the *Schedule Change* panel, complete the following:
 - **Now:** Restarts services immediately.

- **Later:** Enter the following information to schedule the member to restart services at a certain date and time:
 - **Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
 - **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
4. Click **Restart** to restart services immediately or click **Schedule Restart** to schedule the restart.

Canceling a Scheduled Restart

Limited-access users can only cancel a schedule that they created. Superusers can cancel a schedule that any user created. You can cancel scheduled restarts from **Task Manager**. For information, see [Viewing Tasks](#) on page 72.

When you delete a scheduled restart, the system cancels the schedule to restart services on the member or Grid and does not restart services. To cancel a scheduled restart, see [Canceling Scheduled Tasks](#) on page 79.



Chapter 8 File Distribution Services

This chapter describes the file distribution services on the NIOS appliance. It contains the following sections:

- [File Distribution Overview](#) on page 391
 - [Staged Upgrade Limitations](#) on page 391
- [File Distribution Storage](#) on page 392
 - [Usage Threshold Alerts](#) on page 392
 - [Modifying File Distribution Storage Limits](#) on page 392
- [Managing File Distribution Services](#) on page 393
 - [Configuring the TFTP Service](#) on page 393
 - [Configuring the FTP Service](#) on page 393
 - [Configuring the HTTP Service](#) on page 394
 - [Configuring Access Control for File Distribution](#) on page 394
 - [Modifying Access Control Lists](#) on page 395
 - [Starting and Stopping File Distribution Services](#) on page 396
 - [Monitoring File Distribution Services](#) on page 396
- [Managing Directories](#) on page 397
 - [Adding Directories](#) on page 397
 - [Modifying Directories](#) on page 397
 - [Creating a Virtual TFTP Root Directory](#) on page 397
 - [Viewing Directories From the Files Tab](#) on page 398
- [Managing Files](#) on page 399
 - [Uploading Files](#) on page 399
 - [Enabling Upload to Grid Members](#) on page 399
 - [Uploading Files using Grid Manager](#) on page 399
 - [Uploading Files Using TFTP, FTP, or HTTP File Transfer Client](#) on page 400
 - [Deleting Files From the Grid Master](#) on page 400
 - [Deleting Files From a Member](#) on page 400
- [Viewing Files](#) on page 401
 - [Viewing Files from the Files Tab](#) on page 401
 - [Viewing Files from the Members Tab](#) on page 401
- [Managing Users](#) on page 402
 - [Users Default Home Directory](#) on page 402
 - [Enabling FTP Anonymous User](#) on page 394

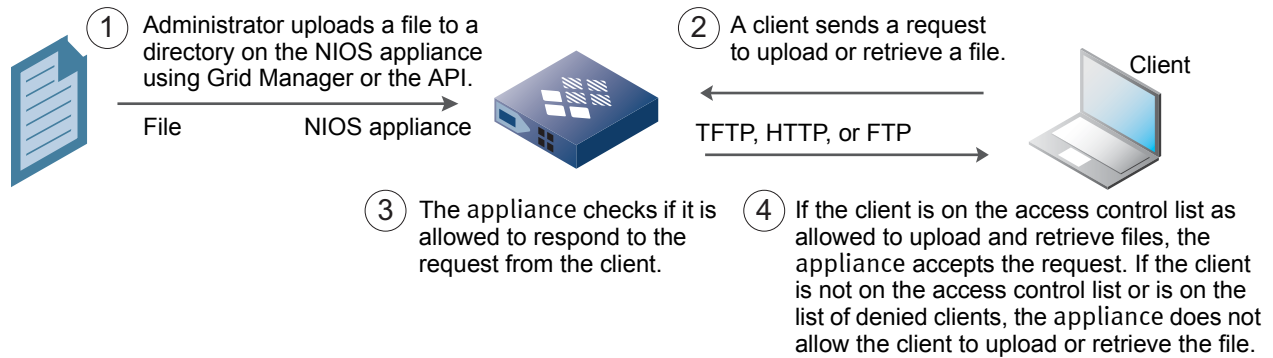
- [*Adding FTP Users through Grid Manager*](#) on page 402
- [*Adding FTP Users through CSV Import*](#) on page 403
- [*Modifying FTP Users*](#) on page 403

FILE DISTRIBUTION OVERVIEW

You can upload files to the NIOS appliance and to individual Grid members using TFTP, HTTP, and FTP, clients. You can also upload files using the Grid Manager web interface or the API. Using access control lists, you can specify which network devices can upload files or retrieve files.

Network devices, such as VoIP phones, can use the DHCP services on the appliance for IP address assignments and use the file distribution services for IP device configuration downloads. Downloads can be accomplished with TFTP, HTTP, or FTP.

Figure 8.1 Uploading and retrieving files



All file uploads and downloads by file distribution clients are logged in the SysLog under the **Administration -> Logs** tabs. The logs store the following information:

- Client IP
- Date and Time
- Event type
- File(s) downloaded and/or uploaded

Staged Upgrade Limitations

All Grid members must be running the same version for the synchronization to work properly among members. It is best to restrict enabling the new enhancements until the entire Grid upgrades to NIOS 6.4 or later. Until the Grid completes staged upgrades and all Grid members are at the Grid master version, the following issue may occur:

- New functionalities such as file uploads from clients, FTP users, and virtual TFTP root directory will fail on Grid members running earlier NIOS versions.

FILE DISTRIBUTION STORAGE

This section describes the storage capacity for file distribution, and how to manage file distribution storage settings. Maximum storage space allowed for all file distribution services on a Grid is equal to the storage space allowed on the Grid member with the smallest amount of space allowed. For example, if a Grid has a Riverbed member, which has a maximum limit of 1GB file distribution storage, then the maximum storage capacity you can set at the Grid level is 1GB, even if the Grid includes appliances with maximum limits of 10GB.

Maximum storage space is shown in table [Table 8.1](#).

Table 8.1 Maximum Storage Space by Platform Type

Member Type	Description	Max Limit
INFOBLOX_MEMBER	All Infoblox appliances	10GB
VNIOS_MEMBER	All virtual appliances (VMWare)	5GB
CISCO_MEMBER	Cisco members	5GB
RIVERBED_MEMBER	Riverbed members	1GB
VM_MEMBER	Virtual IPAM member (I PAM Free Ware)	1GB

Usage Threshold Alerts

An SNMP trap generates an alarm message when a member nears storage capacity. The default threshold value is 90%, and the default reset value is 70%. If email notification is enabled, NIOS sends an email when either of these thresholds are reached.

- When the Grid member storage capacity reaches 100%, the SNMP trap generates a “High Usage” message. For information on how to modify the threshold values, see [Defining Thresholds for Traps](#) on page 1043.
- File distribution clients will fail to PUT files if the file is large enough that it will put the member over the storage limit.

Modifying File Distribution Storage Limits

1. From the **Data Management** tab, select the **File Distribution** tab, and then click **Grid File Distribution Properties** from the Toolbar.
2. In the *Grid File Distribution Properties* editor, complete the following:
 - **Storage Limit (MB):** Enter the maximum storage space in megabytes.
 - **Include files and directories in system backup:** This is selected by default to ensure that the appliance includes the uploaded files in the backup. You can clear this check box to improve the backup performance if you have stored these files separately.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: To avoid data loss, after you change the storage limit FD services will be disrupted briefly and will take some time to resume. Wait until the File Distribution services are running again on the members before you upload any files.

MANAGING FILE DISTRIBUTION SERVICES

This section describes how to configure file distribution services such as TFTP, FTP and HTTP. This section also describes how to configure access control lists which determine which clients are granted access to the service, and which clients are denied access to the service.

Configuring the TFTP Service

The TFTP file distribution service is disabled on the appliance by default. To allow file distribution access using TFTP, you must specify the clients that are allowed to use the service and then enable the service on the appliance. If you do not specify this information or enable the service, the appliance denies access to all clients. The appliance provides read-only access to the files.

To configure the TFTP file distribution service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the **TFTP** tab, and then complete the following:
 - **Listen on Port:** Enter the number of the port on which the appliance receives TFTP file distribution requests. The default is port 69.
 - **Allow file transfers from:** Configure the appliance to grant or deny permissions to TFTP file distribution requests from clients, as described in [Configuring Access Control for File Distribution](#) on page 394.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

After you configure the TFTP service, you must enable the service to allow file distribution access. For information, see [Starting and Stopping File Distribution Services](#) on page 396.

Configuring the FTP Service

The FTP file distribution service is disabled on the appliance by default. To allow file distribution access using FTP, you must create at least one user (see [Managing Users](#) on page 402), specify the clients that are allowed to use the service, and then enable the FTP service on the appliance. If you do not specify this information or enable the service, the appliance denies access to all clients. User creation is not necessary to access the FTP service if anonymous is enabled at Grid level. The appliance provides read-only access to the files.

To configure the FTP file distribution service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the **FTP** tab, and then complete the following:
 - **Listen on Port:** Enter the number of the port on which the appliance receives FTP file distribution requests. The default is port 21.
 - **Login Banner:** Enter your own login banner text that appears after you establish an FTP connection or use the default (**Restricted Access Only**).
 - **FTP Passive Mode:** By default, this is selected to enable FTP in passive mode; otherwise, it is in active mode. An FTP connection between a client and server can be in active or passive mode. In active mode, the server initiates the data connection. In passive mode, the client initiates the data connection. Depending on your firewall policy, firewalls can block active mode connections. There is no firewall filtering in passive mode.
 - **FTP File Listing:** Select this to allow users to list files and subdirectories on the appliance.
 - **Allow file transfers from:** Configure the appliance to grant or deny permissions to FTP file distribution requests from clients, as described in [Configuring Access Control for File Distribution](#) on page 394.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Enabling FTP Anonymous User

The 'anonymous' FTP login is disabled by default, except when upgrading an earlier version in which case anonymous FTP is automatically enabled.

When you enable anonymous FTP at Grid level, you enable anonymous FTP on all Grid members running the FTP service. Anonymous user is only allowed to download files, even if the member is enabled to allow uploads.

1. From the **Data Management** tab, select Grid **File Distribution** Properties on the toolbar.
2. In the *Grid File Distribution Properties* dialog box, select the **Enable Anonymous FTP** check box.
3. Click **Save & Close**.

Configuring the HTTP Service

To allow file distribution access using HTTP, you must specify clients that can request the service and then enable the HTTP service on the appliance.

Before you enable the HTTP service, however, be aware of the following configuration rules:

- HTTP only runs on the active member of an HA pair.
- HTTP can run on the master or any member.
- HTTP always runs on the LAN port, never the MGMT port.
- HTTP to HTTPS redirect becomes non-functional if the file distribution service is enabled and all administrative access is run on the LAN port. For more information on HTTP redirect, see [Enabling HTTP Redirection](#) on page 344. For information on how to specify the MGMT port for HTTP, see [Using the MGMT Port](#) on page 359.

To configure the HTTP file distribution service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the **HTTP** tab, and then complete the following:
 - **Allow Any:** This is selected by default to allow HTTP file distribution requests from any client.
 - **Only these addresses:** Select this to configure the access control list for allowing HTTP file distribution requests from clients, as described in [Configuring Access Control for File Distribution](#) on page 394.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring Access Control for File Distribution

You can select a named access control list (ACL) or create individual access control entries (ACEs) for each file distribution service (TFTP, FTP, HTTP) to control access to file distribution requests from specific clients. You can grant or deny access from specific IPv4 addresses and IPv4 networks, but you cannot do so for IPv6 addresses and IPv6 networks as well as TSIG key based ACEs.

Note: For HTTP service, you can grant permissions to all clients or specific clients, but you can deny permissions only to all clients, not specific clients.

When you grant access to a network for a specific file distribution service, all clients in the network are allowed to request file distribution service. You can deny services to specific IP addresses within the network by adding these addresses to an access control list and denying access to the service. Ensure that you list these IP addresses before the network address in the list because the appliance applies permissions to the addresses in the order they are listed. You can use the arrow keys to move the addresses up and down the list after you add them. For information about how to create a named ACL, see [Configuring Access Control](#) on page 306.

To configure an access control list for a file distribution service:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select a service tab: **TFTP**, **FTP**, or **HTTP**.

3. In the **Allow these clients to perform file transfers** section, select one of the following:
 - For TFTP and FTP: **None**: Select this to deny any clients from using the TFTP and FTP file distribution services. This is selected by default.
 - For HTTP: **Any**: Select this to allow any clients to use the HTTP file distribution service. This is selected by default.
 - **Named ACL**: Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. File distribution does not support IPv6 addresses/networks and TSIG key based ACEs. When you select this, the appliance allows clients that have the **Allow** permission in the named ACL to use the file distribution service. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs**: Select this to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
 - **IPv4 Address**: Select this to add an IPv4 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network**: In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address**: Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission**: Select **Allow** or **Deny** from the drop-down list.
 - **Any Address/Network**: For TFTP and FTP only. Select this to allow or deny access to all IPv4 addresses and networks. The default permission is **Allow**, which means that the appliance allows access to and from all IPv4 clients. You can change this to **Deny** to block access.

After you have added access control entries, you can do the following:

 - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying Access Control Lists

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the tab of the service to which the list belongs.
3. In the **Allow file transfers from** section, modify the fields as described in [Configuring Access Control for File Distribution](#) on page 394.

You can also do the following:

- Add a new permission. For information, see [Configuring Access Control for File Distribution](#) on page 394.
- Delete a permission by selecting it and clicking the Delete icon.
- Reorder the list by selecting a permission and clicking an arrow next to the list to move the permission up or down the list.

Starting and Stopping File Distribution Services

You can enable and disable a file distribution service on a specific Grid member or on multiple members. You must have read/write permission to the Grid members to start and stop a service on them.

Starting a service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Start icon from the Toolbar. You can select multiple members by selecting their check boxes.
2. From the Start drop-down menu, select the service you want to start.
3. In the *Start Service* dialog box, click **Yes**.
Grid Manager enables the selected service on the selected member and displays the service status in the Status column in the panel.

Stopping a service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* check box, and then click the Stop icon from the Toolbar. You can select multiple members by selecting their check boxes.
2. From the Stop drop-down menu, select the service you want to stop.
3. In the *Stop Service* dialog box, click **Yes**.
Grid Manager disables the selected service on the selected member and displays the service status in the Status column in the panel.

Note: When you start or stop a service, there may be a short delay before Grid Manager displays the correct status.

Monitoring File Distribution Services

To view the current status of the file distribution services:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab.
2. Grid Manager displays the following information:
 - **Name:** The name of the Grid member.
 - **Address:** The IP address of the Grid member.
 - **Status:** The overall status of the file distribution services running on the member. You can mouse over on the field to view the status of each service. This field can display one of the following:
 - **Not Running:** All the file distribution services are disabled.
 - **Running:** One or more of the file distribution services are running properly.
 - **Warning:** The services are functioning properly. However, there are some issues, such as storage space has reached 90%, about the services.
 - **Error:** One or more of the services have service issues.
 - **Comment:** Information about the member.
 - **Site:** The location to which the member belongs. This is one of the pre-defined extensible attributes.

You can sort the information in ascending or descending order by columns. You can also print and export the information in this panel.

MANAGING DIRECTORIES

You can create directories on the Grid Master and on Grid members, in which you can store your files. You can manage the directories in the following ways:

- Create a directory structure for file distribution, as described in [Adding Directories](#).
- Modify the directory name and permissions, as described in [Modifying Directories](#) on page 397.
- Create a Virtual TFTP root directory, as described in [Creating a Virtual TFTP Root Directory](#) on page 397.
- View the directories, as described in [Viewing Directories From the Files Tab](#) on page 398.

Adding Directories

To add a directory:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Click the parent directory link, and then click **Add** -> **Directory** from the Toolbar.
3. Grid Manager adds a new directory to the parent directory and gives it the default name **NewDirectory**.

You can modify the directory name and permissions, as described in [Modifying Directories](#) on page 397.

Modifying Directories

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Select a directory check box and click the Edit icon.
3. The *Directory* editor provides the following tabs from which you can modify data:
 - **General:** You can modify the directory name here, except for the Root directory.
 - **Virtual TFTP Root:** You can add an IP Address, a Network or a Range of IP addresses to support VMware ESX hosts who need different PXE boot images based on where they are in the network.
 - **Permissions:** You can add or delete admin permissions in this tab. For information, see [About Administrative Permissions](#) on page 160.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

You can also select a directory and click the Delete icon to delete it.

Note: When you delete a directory, the appliance automatically removes all its contents in that directory.

Creating a Virtual TFTP Root Directory

This section describes how to create a Virtual TFTP root directory for a specific IP address, network, or range of IP addresses.

Note: Root directory can not be a virtual TFTP root.

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Select the directory check box and click the Edit icon.
3. In the *Directory* editor select **Virtual TFTP Root**. Click the Add icon and select one of the following:
 - **IP Address:** This creates a virtual TFTP directory that the clients from a specified IP address will see as the root directory.
 - **Network:** This creates a virtual TFTP directory that the clients on a specified network will see as the root directory.

- **Range:** This creates a virtual TFTP directory that the clients in a specified range of IP addresses will see as the root directory.
- 4. From the drop-down in the **Member** column, select the member on which to make the virtual TFTP root directory.
- 5. In the **Address/Network** column, enter a value:
 - **IP Address:** Enter the IP address of the client that will have access to the virtual TFTP root directory. This IP address must be on the allow list in the TFTP access control list.
 - **Network:** Enter a network address using the format 10.0.0.0/24. This allows all clients in that network to access the virtual TFTP root directory. This network address must be on the allow list in the TFTP access control list.
 - **Range:** Enter the first IP address in the range Address/Network column, and the last IP address in the range in the End column. This allows all clients in that range to access the virtual TFTP root directory. This range must be on the allow list in the TFTP access control list.
- 6. Click Save & Close. Click **Restart** if it appears at the top of the screen.
- 7. To create more virtual TFTP root directories, repeat Steps 3 through 5.

Viewing Directories From the Files Tab

1. From the **Data Management** tab, select the **File Distribution** tab → **Files** tab.
2. Grid Manager displays the following information in the Root directory.
 - **Name:** The name of the directory or file.
 - **Type:** Depending on the file type, this can be **Directory** or **File**.
 - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
 - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.
3. Click the directory link to view files and directories in a specific directory.

You can also do the following in this panel:

- Sort the information in ascending or descending order by columns.
- Use the breadcrumb to go to a specific directory.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Print and export the information in this panel.
- Add a directory or a file. For information, see [Adding Directories](#) on page 397 and [Managing Files](#) on page 399.
- Open and edit a directory. For information, see [Modifying Directories](#) on page 397.

MANAGING FILES

This section describes how to upload files using the Grid Manager or a file transfer client. You can upload files to the Grid Master or to individual members.

Uploading Files

Some things to keep in mind when you upload files:

- When you use the Grid Manager to upload files, you can upload files only to the Grid Master, not to individual members of the Grid.
- To upload files to a member, you must use an FTP client, TFTP client or HTTP client. Files uploaded by file transfer clients to any member, will be synchronized back to Grid Master.
- The logs for file transfers using third party clients can be found in syslog.
- You can use a third party file transfer client to upload and retrieve files:
 - If the ‘anonymous’ login is enabled, you can retrieve files but this ‘anonymous’ user can not upload files even if the “Allow uploads” option is enabled.
 - If you create a user to use with a third party transfer client, this must be an FTP user with read/write permissions in their directory.
- You can upload a maximum of 10,000 files.
- If uploading the file exceeds the storage limit, the appliance logs a message and does not upload the file. For information about file distribution storage, see [Modifying File Distribution Storage Limits](#) on page 392.
- If you upload a file that has the same name and path as an existing file, the appliance automatically replaces the old file.

Note: Administrators with superuser privileges can manage uploading files. Limited-access admins with read/write permissions to specific directories can upload files to the directories. For information, see [Administrative Permissions for File Distribution Services](#) on page 214.

Enabling Upload to Grid Members

1. From the **Data Management** tab, select **Grid File Distribution Properties** on the toolbar.
2. In the *Grid File Distribution Properties* dialog box, select the **Allow Upload to Grid Members** check box.
3. Click **Save & Close**.

Uploading Files using Grid Manager

The Grid Manager uploads files only to the Grid Master. The Grid Master propagates the files to the members. You must use a third party file transfer client to upload files directly to an individual member:

- If the ‘anonymous’ login is enabled, you can retrieve files but this ‘anonymous’ user can not upload files even if the “Allow uploads” option is enabled.
 - If you create a user to use with a third party transfer client, this must be an FTP user with read/write permissions in their directory.
1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
 2. Select the destination directory link.
 3. Click the **Add** icon -> **File** from the Toolbar.
 4. Select the **Extract files after upload (.zip, .tar, .gz, .tgz)** check box in the *Upload* dialog box if you are uploading .zip, .tar, .gz, or .tgz files and you want to automatically extract the files upon upload.

Note: The directory structure in the compressed file is restored when the files are extracted. A directory that already exists it will be replaced by an extracted directory with the same name.

5. Click **Select** to navigate to the file you want to upload.
 6. Select the file you want to upload, and then click **Open**.
 7. If you want to upload more than one file, repeat Steps 4 and 6 until you have selected all the files you want to upload. You can upload a maximum of ten files at one time.
-

Note: You can delete an incorrect file selection by clicking the red icon next to the filename before you click Upload

8. To verify the upload was successful, roll the mouse cursor over the green check mark next to the file name. If the upload was successful, the message “Upload succeeded.” appears.

Uploading Files Using TFTP, FTP, or HTTP File Transfer Client

You can upload files to the Grid Master or to individual members using a third party FTP client. Files uploaded by file transfer clients to any member, will be synchronized back to Grid Master.

To upload files to a member, you must first enable the **Allow Upload to Grid Members** check box in the *Grid File Distribution Properties* dialog box. See [Enabling Upload to Grid Members](#) on page 399.

You must add an FTP user before you can upload files using a third party FTP client. This must be an FTP user. It is not the NIOS admin. For information see [Adding FTP Users through Grid Manager](#) on page 402.

Deleting Files From the Grid Master

If the FTP user has read/write permissions, then that user can delete files from the Grid member wherever that FTP user is connected. Only files can be deleted but not directories.

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the files and folders in the Root directory. Click the directory link to see the files in a specific directory.
3. To delete a file, select the check box and then click the **Delete** icon.

Deleting Files From a Member

You can delete files from a member only if “**No**” appears in the *Synchronized with Grid Master* column.

If the FTP user has Read/Write permissions, then that user can delete files from the Grid member wherever that FTP user is connected. Only files can be deleted but not directories.

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the files and folders in the Root directory. Click the directory link to see the files in a specific directory.
3. If “**No**” appears in the *Synchronized with Grid Master* column, select the check box, then click the **Delete** icon.

VIEWING FILES

You can view files from the Files Tab and from the Members Tab.

Viewing Files from the Files Tab

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the following information in the Root directory.
 - **Name:** The name of the file.
 - **Type:** Depending on the file type, this can be **Directory** or **File**.
 - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
 - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.

You can view files and directories in a specific directory by clicking the directory link.

Viewing Files from the Members Tab

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab.
2. Grid Manager displays the following information in the Root directory.
 - **Name:** Member name.
 - **IPv4 Address:** Member's IP address.
 - **Status:** State of the member, running or not running.
 - **Comment:** Additional comments about the member.
 - **Site:** User defined information about the site.
3. To see the files on the Grid Master, click on the name of the Grid Master. Grid Manager displays the following information:
 - **Name:** The name of the file.
 - **Type:** Depending on the file type, this can be **Directory** or **File**.
 - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
 - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.

Tip: When you drill down on the Grid Master from the Members tab, the Add icon is activated.

4. To see the files on a member, click on the name of the member. Grid Manager displays the following information:
 - **Name:** The name of the file.
 - **Type:** Depending on the file type, this can be **Directory** or **File**.
 - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
 - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.
 - **Synced with Grid Master:** You cannot delete files with a value other than “No”. If this value is “No”, you must delete the file.

Note: You cannot upload, modify, or delete a file or a directory when you drill down from the Members tab.

You can also do the following in this panel:

- Sort the information in ascending or descending order by columns.
- Use the breadcrumb to go to a specific directory.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Print and export the information in this panel.
- Add a directory or a file. For information, see [Adding Directories](#) on page 397 and [Managing Files](#) on page 399.
- Open and edit a directory. For information, see [Modifying Directories](#) on page 397.

MANAGING USERS

This section describes how to add and modify user accounts for use with an FTP client.

You must be a NIOS admin user with super user privileges to add, modify, or delete FTP users. FTP users are created at Grid level, so the same users will be available to access FTP service on all members.

- Each user must have unique Username.
- By default, the home directory with the user name is under the `/ftputers` directory. However, the user can also choose to use an existing directory outside of `/ftputers` as his home directory. If the admin specified home directory is not available then it will raise an error.
- Permission: Read-write or Read-only are assigned for each FTP user. Users with read-write permissions are allowed to upload files, delete files and list the files and directories under his home directory.
- You can have multiple users to use same home directory. One user may have read-only permissions while others have read-write permissions on same home directory.
- FTP users are not allowed to add, modify, or delete the directories, even with read-write permissions.
- If the "Allow uploads on the member" is disabled, then users with read-write permission also can not upload files to his home directory.

Users Default Home Directory

- FTP users default home directory is `/ftputers`.
- The `/ftputers` directory is created by default and listed in the 'Files' viewer under the root directory. By default, home directories for FTP users are under this directory.
- NIOS admin is allowed to upload and delete files to and from users home directories.
- Files uploaded by FTP users are visible in the Grid Manager.

Adding FTP Users through Grid Manager

1. From the **Data Management** tab, select the **File Distribution** tab → FTP Users tab, and then click the Add icon.
2. In the *Add FTP User* dialog box, complete the following:
 - **Username:** Enter a name for the user. This is the username that the user uses to log in.
 - **Password:** Enter a password for the user to use when logging in.
 - **Confirm Password:** Enter the same password.
 - **Permissions:** From the drop down choose from the following:

- **Read Only:** This allows the user to display files, but not to upload, or delete files using a third party FTP client.
- **Read/Write:** This allows the user to upload, delete and list files using a third party FTP client.
- Choose a directory for this user. This is the directory where files uploaded with this Username will go:
 - **Create Home Directory:** This creates a directory using the Username.
 - **Choose Specified Directory:** This allows the user to choose an existing directory.

3. Click **Save & Close**.

Adding FTP Users through CSV Import

You can add an FTP User by importing a CSV file with the headers in the following format:

```
version 1.0,,,,
header-ftpuser,username*,password*,create_home_dir,home_dir,permission
ftpuser,user1,passwd1,True,/ftpusers/user1,RO
```

Modifying FTP Users

1. From the **Data Management** tab, select the **File Distribution** tab -> FTP Users tab.
2. Select the check box for the user you want to modify and click the Edit icon.
3. In the *FTP User Editory* you can modify the following:
 - **Password:** Enter a password for the user to use when logging in.
 - **Confirm Password:** Enter the same password.
 - **Permissions:** From the drop down choose from the following:
 - **Read Only:** This allows the user to display the files and their properties, but not to edit them.
 - **Read/Write:** This allows the user to display and edit the files and their properties.
4. Click **Save & Close**.



Chapter 9 Managing NIOS Software and Configuration Files

This chapter explains how to manage upgrade groups and perform software upgrades and downgrades for NIOS appliances. It also describes how to back up and restore configuration files. It includes the following sections:

- [About Upgrades](#) on page 406
 - [Lite Upgrades](#) on page 406
 - [Full Upgrades](#) on page 407
 - [Guidelines for Scheduling Full Upgrades](#) on page 407
- [Managing Upgrade Groups](#) on page 408
 - [Adding Upgrade Groups](#) on page 409
 - [Modifying Upgrade Groups](#) on page 410
 - [Viewing Upgrade Groups](#) on page 410
 - [Deleting Upgrade Groups](#) on page 411
- [Viewing Software Versions](#) on page 411
- [Upgrading NIOS Software](#) on page 411
 - [Uploading NIOS Software](#) on page 412
 - [Distributing Software Upgrade Files](#) on page 412
 - [Managing Distributions](#) on page 414
 - [Testing Software Upgrades](#) on page 415
 - [Performing Software Upgrades](#) on page 416
 - [Managing Upgrades](#) on page 420
 - [Monitoring Distribution and Upgrade Status](#) on page 421
- [Downgrading Software](#) on page 422
- [Reverting the Grid to the Previously Running Software](#) on page 423
- [Backing Up and Restoring Configuration Files](#) on page 423
 - [Backing Up Files](#) on page 423
 - [Automatically Backing Up Data Files](#) on page 424
 - [Manually Backing Up Data Files](#) on page 426
 - [Downloading Backup Files](#) on page 427
 - [Restoring Backup Files](#) on page 428
 - [Downloading Backup Files from a Different Appliance](#) on page 429
- [Downloading Support Bundles](#) on page 429

ABOUT UPGRADES

Infoblox frequently releases updated NIOS software. Contact Infoblox Technical Support to learn which file name to use when downloading a new upgrade file, or watch your email for periodic notifications that a new software upgrade is available. To get the latest upgrade, your local network must be capable of downloading a file from the Internet. For information about how to upgrade, see [Upgrading NIOS Software](#) on page 411.

You can upgrade an appliance to a specific release if the current release on your appliance supports the upgrade path. For information about the upgrade and revert paths of a specific release, refer to the latest release notes at <https://support.infoblox.com>. Depending on whether there are database schema changes between the existing and upgrade releases, the appliance can perform either a lite or full upgrade. For information, see [Lite Upgrades](#) and [Full Upgrades](#).

You can schedule certain upgrades for a Grid. Scheduling an upgrade can minimize network and operational outages, especially when Grid members are spanned across different time zones. You can also arrange the upgrade to happen during non-peak hours for specific members to avoid overloading the network traffic. When you schedule an upgrade, you can schedule to update all Grid members at the same time or at different times. Depending on the configuration of your Grid and the software version that is currently running in the Grid, you can also schedule your upgrades for different members over a period of time. For more information, see [Scheduling Upgrades](#) on page 416.

Based on your network requirements and topology, you can organize your members into upgrade groups so these members can be upgraded at the same time. For more information about upgrade groups, see [Managing Upgrade Groups](#) on page 408.

You can also import and export upgrade groups and their distribution and upgrade schedules in CSV format. For information about how to import and export in CSV format, see [About CSV Import](#) on page 86 and [Exporting Data to Files](#) on page 89.

Note: When you promote a Grid Master candidate to a Grid Master, you cannot revert to the previous release.

Lite Upgrades

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

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

Whenever possible, the appliance uses the lite upgrade mode to speed up the upgrade process. You can always schedule a lite upgrade. Note that the appliance disables the testing function for lite upgrades because you do not need to test a lite upgrade for any database translation. For information about how to schedule an upgrade, see [Scheduling Upgrades](#) on page 416.

Full Upgrades

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

Depending on the upgrade paths your current release supports, when you schedule a full upgrade, the Grid Master immediately replicates certain core network service tasks to Grid members while putting other tasks in queue until the members have been upgraded. For information about which data and tasks the Grid Master replicates to members immediately, see [Guidelines for Scheduling Full Upgrades](#) on page 407. For information about how to schedule an upgrade, see [Scheduling Upgrades](#) on page 416.

GUIDELINES FOR SCHEDULING FULL UPGRADES

When you schedule a full upgrade from NIOS 6.6.x to a later release, the Grid Master immediately replicates the following to the Grid members, including those that have not been upgraded:

DNS resource records, DNS zones, DNS views, name server groups, shared record groups, IPv4 and IPv6 host addresses, roaming hosts, IPv4 and IPv6 networks, IPv4 and IPv6 shared networks, fixed addresses, DHCP ranges, DHCP failover association, DHCP option spaces, DHCP options, DHCP filters, MAC filter items, blacklist & NXDOMAIN rules, DNSSEC key pairs, DNSSEC import keyset operation, sign and unsign zones, DNSSEC rollover KSK and ZSK operations.

You can also perform the following tasks:

- Upgrade a specific member during the scheduled Grid upgrade. For information about how to upgrade a single member during a scheduled Grid upgrade, see [Upgrading a Single Member Immediately](#) on page 418.
- Revert a single member that has already been upgraded so you can troubleshoot issues, such as service outages, on the specific member. You can then reschedule its upgrade. For more information, see [Reverting a Single Member](#) on page 418.
- During an upgrade, you can start, stop, or restart DNS and DHCP services through Grid Manager on members that have not been upgraded.
- Clear authentication cache and authentication records.
- Perform AD (Active Directory) configurations. Note that you must upload the keytab file before the upgrade starts.

The appliance also puts certain rules in place to ensure data integrity and controls data that can cause undesirable results during the upgrade process. When you schedule a full upgrade from NIOS 6.6.x to a later release, the following rules apply:

- You cannot modify member properties for the following: DNS, DHCP, TFTP/HTTP/FTP, bloxTools, Captive Portal, Reporting, and load balancing until the member has completed the upgrade and exited its revert time windows.
- You cannot delete DNS views.
- You cannot delete DNS zones and IPv4 and IPv6 networks that are under Microsoft Management until the managing member of the Microsoft servers has completed its upgrade and exited its revert time window. Certain Microsoft management restrictions also apply, as described in [Managing Upgrade Groups](#) on page 408.
- Synchronization between load balancers and the appliance is disabled until the load balancer managing member has completed its upgrade and exited its revert time window. You cannot change the managing member during the upgrade.
- You cannot add, modify, or delete network views, rulesets, and DNS64 synthesis groups.
- Replication of Grid and member DNS and DHCP properties is not supported.
- You can create named ACLs (access control lists) only after the entire Grid has been upgraded. For information about named ACLs, see [Configuring Access Control](#) on page 306.

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

- DHCP fingerprint detection is disabled
- You cannot add DHCP fingerprint filters
- You cannot apply DHCP fingerprint filters to any DHCP address range

Microsoft Management Rules

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

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

Note: Note that the deactivation of these functions does not affect any data on the Microsoft servers. After the upgrade, the member automatically restarts the synchronization of Microsoft data.

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

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

MANAGING UPGRADE GROUPS

To minimize the impact of Grid upgrades on your system operations, you can organize members into upgrade groups and schedule their software distributions. This is useful, for example, in a large Grid spanning multiple time zones where there are fluctuating network and downtime considerations at various locations. Note that you can also schedule their upgrades, depending on the existing releases and their upgrade paths. For information about the different upgrade methods, see [About Upgrades](#) on page 406.

You can also import and export upgrade groups and their schedules in CSV format. For more information, refer to the *Infoblox CSV Import Reference*.

Infoblox provides two default upgrade groups:

- **Grid Master**—After you configure the Grid Master, it automatically becomes the only member of this group. You cannot modify or delete this group.
- **Reporting Member**—After you configure a reporting member in a Grid, it automatically becomes the only member of this group. This group will be upgraded automatically after the Grid Master and before other upgrade groups. You cannot modify, delete, or schedule this upgrade group. For information about reporting, see [Infoblox Reporting Solution](#) on page 1113.
- **Default**—This is the default upgrade group to which the appliance automatically assigns Grid members. If you do not explicitly assign a member to an upgrade group, it remains in the Default group. You cannot delete or rename this group. For information, see [Modifying Upgrade Groups](#) on page 410.

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

- Add an upgrade group, as described in [Adding Upgrade Groups](#).
- Modify an upgrade group, as described in [Modifying Upgrade Groups](#) on page 410.
- View upgrade group information, as described in [Viewing Upgrade Groups](#) on page 410.
- Delete an upgrade group, as described in [Deleting Upgrade Groups](#) on page 411.

Adding Upgrade Groups

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

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

Note the following recommendations when you create an upgrade group:

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

To add an upgrade group:

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

Note: After you add a member, the appliance adds it to the group members list. The first Grid member in the list determines the time zone of the group when you schedule the distribution and upgrade. Therefore, Grid Manager displays the time zone of the first Grid member in the list. (For information about setting time zones, see [Managing Time Settings](#) on page 312.)

5. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying Upgrade Groups

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

To modify an upgrade group:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Group List View**.
2. Select an *upgrade_group* check box, and then click the Edit icon in the row. You can also click the Edit icon directly without selecting the check box.
3. The *Upgrade Group* editor provides the following tabs from which you can modify data:
 - **General:** Modify the fields as described in [Adding Upgrade Groups](#) on page 409.
 - **Member Assignment:** Add or delete members as described in [Adding Upgrade Groups](#) on page 409.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Viewing Upgrade Groups

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

Before a distribution or upgrade starts, you can move members from one group to another, reorder the members, or remove a member from an upgrade group. The member you remove automatically joins the Default group. (For information, see [Managing Distributions](#) on page 414.) You cannot add, delete, or reorder members in an upgrade group while a distribution or upgrade is in progress. You can skip a member in an upgrade group from a distribution only before the distribution starts, or after you pause it. For information, see [Pausing and Resuming Distributions](#) on page 414.

To view the upgrade groups in a Grid:

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

You can hide some of the default columns, but you cannot sort the information in this table. You can use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

Deleting Upgrade Groups

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

To delete an upgrade group:

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

VIEWING SOFTWARE VERSIONS

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

To view software information:

1. From the **Grid** tab, select the **Upgrade** tab.
2. Grid Manager displays the following in the Grid Version Information section:
 - **Running:** The NIOS software version that is currently running on the Grid.
 - **Uploaded:** The latest NIOS image file you have uploaded and is available for distribution.
 - **Distribution:** The NIOS software version used for distribution or is available for distribution.
 - **Revert:** The NIOS software version to which the appliance can revert.
 - **Distribution Schedule:** Displays the date and time of the next scheduled distribution.
 - **Upgrade Schedule:** Displays the date and time of the next scheduled upgrade.

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

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

UPGRADING NIOS SOFTWARE

Infoblox frequently releases updated NIOS software. Contact Infoblox Technical Support to learn which file name to use when downloading a new upgrade file, or watch your email for periodic notifications that a new software upgrade is available. To get the latest upgrade, your local network must be capable of downloading a file from the Internet.

After you download and store the new upgrade file on your local network, complete the following tasks to upgrade an Infoblox independent appliance or a Grid.

- Upload the new software to the Grid Master, as described in [Uploading NIOS Software](#).
- Distribute the software upgrade files, as described in [Distributing Software Upgrade Files](#) on page 412.
- Optionally, test the upgrade, as described in [Testing Software Upgrades](#) on page 415.
- Perform the software upgrade, as described in [Performing Software Upgrades](#) on page 416.

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

Note: You cannot upgrade directly to NIOS 5.x from NIOS releases less than 4.2r4. Refer to the release notes for the appropriate upgrade and revert paths.

Before you upgrade to a later NIOS release, use the `show upgrade_compatible` command to check if your Grid is compatible with the release. For information about using this command, refer to the *Infoblox CLI Guide*.

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

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

Uploading NIOS Software

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

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Upload** in the panel or from the Toolbar.
2. Navigate to the directory where you have stored the NIOS software upgrade, and then click **Open** or **Upload**.
The appliance uploads the file and displays the status of the upload in the status bar. You can click the Stop icon in the status bar to stop the upload. Ensure that you do not navigate away from the **Upgrade** tab until after the upload is complete. Otherwise, the upload process stops.

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

Distributing Software Upgrade Files

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

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

Distributing Software Immediately

The Grid Master distributes the software upgrade to each member in the Grid, including itself. As an alternative to scheduling the Grid distribution (see [Scheduling Distributions](#) on page 413), you can distribute the software upgrade throughout the Grid immediately, as follows:

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

The distribution starts and if there is an active distribution scheduled, the appliance changes its status to inactive. The appliance distributes the upgrade files and displays the status of the distribution in the status bar. You can pause, resume, or stop the distribution by clicking the corresponding icon in the status bar. For information, see [Managing Distributions](#) on page 414.

Note that starting a manual distribution cancels a scheduled distribution.

Scheduling Distributions

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

To schedule a software distribution:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Distribute -> Schedule Distribution** from the Toolbar.
 2. In the *Schedule Distribution* editor, complete the following:
 - **Activate Distribution Schedule:** Select this to enable the distribution schedule. Clear this if you are creating a distribution schedule you plan to activate at a later date. You can configure and save information in this editor even when you deactivate a scheduled distribution.
 - **Grid Master Distribution Start Information:** Enter a Grid Master distribution date, time, and time zone. The distribution date and time must be before those of the upgrade groups.
 - **Date:** Enter a start date of the Grid Master distribution in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter a start time of the Grid Master distribution in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can also select a time from the drop-down list.
 - **Time Zone:** Select a time zone that applies to the start time you enter. If this time zone is different from the Grid time zone, the appliance converts the time you enter here based on the Grid time zone, after you save this schedule. When you display this schedule again, it displays the converted time. Selecting the time zone here does not affect any time zone settings in the Grid. (For information about selecting the Grid and member time zones, see [Managing Time Settings](#) on page 312.)
 - **Admin Local Time:** Displays the Grid Master distribution start date and time in the time zone of the administrator, as explained in [Creating Local Admins](#) on page 169.
 - In the upgrade group table, specify the following for each upgrade group by clicking the corresponding field in each row:
 - **Start Distribution:** Specify when the distribution occurs. Select one of the following from the drop-down list:
 - **Date/Time:** Select this to configure the distribution start date, time, and time zone.
 - **After <group>:** Select **After Grid Master** to start the distribution immediately after the completion of the Grid Master distribution. Select an upgrade group that must complete its distribution before the group you are configuring. When you select this option, you cannot enter a date, time, and time zone.
- Date, Time, and Time Zone** are enabled only when you select **Date/Time** for **Start Distribution**.
- **Date:** Enter a distribution start date in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter a distribution start time in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.
 - **Time Zone:** By default, the appliance displays the time zone of the first Grid member in the Upgrade Group. You can change this time zone if you want to enter the time using a different time zone. After you save the schedule though, the appliance converts the time you entered to the time zone of the upgrade group, if it is different. (For information about setting the Grid and member time zones, see [Managing Time Settings](#) on page 312.) To change the default time zone of the upgrade group, change the time zone of the first group member, as explained in [Adding Upgrade Groups](#) on page 409.

- **Admin Local Time:** Displays the start date and time in the time zone of the administrator, as explained in [Creating Local Admins](#) on page 169.
 - **Distribute to Members:** Indicates whether the distribution within the group occurs simultaneously or sequentially. You cannot edit this field here. You define this when you create the upgrade group. To change this setting, see [Modifying Upgrade Groups](#) on page 410.
3. Save the configuration and click **Restart** if it appears at the top of the screen.
- Grid Manager confirms that the schedule is saved and indicates whether the distribution schedule is active. You can click the Refresh icon to refresh the information in this panel.
- Note that the appliance does not save the schedule and displays an error message if the schedule contains the following:
- Circular dependencies between upgrade groups. For example, the distribution of Group A is scheduled after Group B, and the distribution of Group B is scheduled after Group A.
 - The distribution time is in the past.

Software Distribution Process

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

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

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

Managing Distributions

After you start a distribution, you can pause, resume, or stop it. For information, see [Pausing and Resuming Distributions](#) on page 414 and [Stopping Distributions](#) on page 415. Grid Manager displays the status of the overall distribution as well as the status of individual members. You can view this information in the **Upgrade** tab.

Pausing and Resuming Distributions

The following are some operational guidelines for performing a distribution:

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

To pause a distribution:

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

The Grid Distribution Status bar indicates the distribution is paused. For information about the distribution status of each member, see [Monitoring Distribution and Upgrade Status](#) on page 421.

To skip a member from a distribution:

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

To resume a distribution:

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

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

Stopping Distributions

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

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

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

To stop a distribution:

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

Testing Software Upgrades

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

To start an upgrade test:

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

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

To stop an upgrade test:

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

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

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

After the test successfully finishes, the appliance displays a message confirming that the test upgrade is complete. You can then perform the actual upgrade as described in [Performing Software Upgrades](#) on page 416.

Performing Software Upgrades

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

Note: Before you upgrade the software, Infoblox recommends that you back up the current configuration and database. For information, see [Backing Up Files](#) on page 423.

Depending on your upgrade paths, you can upgrade to a new release immediately or you can schedule the upgrade. For information about how to upgrade immediately, see [Upgrading the Grid Immediately](#). Before you schedule an upgrade, ensure that you understand the limitations, as described in [Managing Upgrade Groups](#) on page 408. For information about how to schedule an upgrade, see [Scheduling Upgrades](#).

Upgrading the Grid Immediately

For unschedulable full upgrades, all the Grid members in the Grid must upgrade at the same time. For lite upgrades and schedulable full upgrades, you can schedule the upgrades as described in [Scheduling Upgrades](#) on page 416, or you can upgrade all the Grid members at the same time.

To upgrade a Grid immediately:

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

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

Scheduling Upgrades

You can schedule lite upgrades and full upgrades for certain NIOS versions. For limitations about scheduling a full upgrade, see [Managing Upgrade Groups](#) on page 408. When you schedule an upgrade, you schedule the upgrade for the Grid Master and the upgrade groups, including the Default group. The Grid Master must always upgrade before the upgrade groups. Depending on your upgrade paths, you can schedule the upgrade for the Grid Master and upgrade groups at different times over a period of nine days. If you schedule an upgrade that takes more than nine days, the appliance displays a warning.

To schedule an upgrade:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Upgrade** -> **Schedule Upgrade** from the Toolbar.
2. In the *Upgrade Schedule* editor, complete the following:
 - **Activate Upgrade Schedule:** Select this to enable the upgrade schedule. Clear it if you are creating an upgrade schedule that you plan to activate at a later date. You can configure and save information in this editor even when you deactivate a distribution.
 - **Grid Master Upgrade Start Information:** Enter a Grid Master upgrade date, time, and time zone. The date and time must be before those of the upgrade groups.
 - **Date:** Enter a start date of the Grid Master upgrade in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter a start time of the Grid Master upgrade in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.
 - **Time Zone:** Select a time zone that applies to the start time you enter. If this time zone is different from the Grid time zone, the appliance converts the time you enter here based on the Grid time zone, after you save this schedule. When you display this schedule again, it displays the converted time. Selecting the time zone here does not affect any time zone settings in the Grid. (For information about setting the Grid and member time zones, see [Managing Time Settings](#) on page 312.)
 - **Admin Local Time:** Displays the Grid Master upgrade date and start time in the time zone of the administrator, as explained in [Creating Local Admins](#) on page 169.

- In the upgrade member table, specify the following by clicking the corresponding field in each row:
 - **Group:** The name of the upgrade group. You can assign a different upgrade group by selecting the group from the drop-down list.
 - **Group Members:** When you expand an upgrade group, this field displays the group members.
 - **Warning:** This field turns yellow when there is a conflict among the upgrade groups. Hover your mouse over the field and the tooltip displays the member that contains the conflict. It also displays recommended upgrade groups in the **Group** column so you can change the group assignment to resolve the conflict. The tooltip can display one of the following: **GMC**, **DNS Primary**, **DHCP Logging Member**, or **DHCP Failover**. For information about how to resolve a conflict, see [Resolving Upgrade Warnings](#) on page 417. Select an upgrade group from the drop-down list in the Group column to assign a different upgrade group. Click **Validate and Refresh** to validate the new group assignment.
 - **Start Upgrade:** Specify when the upgrade occurs. Select one of the following from the drop-down list:
 - **Date/Time:** Select this to configure the upgrade start date, time, and time zone.
 - **After <group>:** Select **After Grid Master** to start the distribution immediately after the completion of the Grid Master distribution. Select an upgrade group that must complete its distribution before the group you are configuring. If you select this option, you cannot enter a date, time, and time zone.

Date, Time, and Time Zone are enabled only when you select **Date/Time** for **Start Upgrade**.
 - **Date:** Enter an upgrade start date in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
 - **Time:** Enter an upgrade start time in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.
 - **Time Zone:** By default, the appliance displays the time zone of the first Grid member in the Upgrade Group. You can change this time zone, if you want to enter the time using a different time zone. After you save the schedule though, the appliance converts the time you entered to the time zone of the upgrade group, if it is different. (For information about setting the Grid and member time zones, see [Managing Time Settings](#) on page 312.) To change the default time zone of an upgrade group, change the first group member in the Upgrade Group list, as explained in [Adding Upgrade Groups](#) on page 409.
 - **Admin Local Time:** Displays the data and time in the time zone of the administrator, as explained in [Creating Local Admins](#) on page 169.
 - **Upgrade Members:** Indicates whether the upgrade within the group occurs simultaneously or sequentially. You cannot edit this field here. You define this when you create the upgrade group. To change this setting, see [Modifying Upgrade Groups](#) on page 410.

3. Save the configuration.

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

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

The appliance also does not save the schedule and displays a warning when there is a group assignment conflict. For information about how to resolve these conflicts, see [Resolving Upgrade Warnings](#).

Otherwise, the appliance confirms that the schedule is saved and indicates whether the upgrade schedule is active.

Resolving Upgrade Warnings

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

- **GMC:** To resolve this warning, put all Grid Master candidates in the first upgrade group.
- **DNS Primary:** To resolve this warning, put all the members that are serving as DNS primaries in the first upgrade group.
- **DHCP Logging Member:** To resolve this warning, put the DHCP logging member in the first upgrade group.

- **DHCP Failover:** To resolve this warning, place the peers of a DHCP failover association in separate upgrade groups. Ensure that you schedule upgrades of the failover peers close to each other to minimize configuration rules. NIOS does not allow DHCP configuration changes that affect the communication between the peers until both peers are upgraded.

Upgrading Groups Immediately

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

To upgrade an upgrade group now:

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

Upgrading a Single Member Immediately

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

To upgrade a specific member now:

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

Reverting a Single Member

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

Once a member is upgraded, the appliance starts counting down and displays the time that is left for you to revert this member. You can revert the member before the revert time window expires. The default time window to revert a member is 24 hours. You can view the time that is left to revert the member in the Member List view, as described in [Grid and Member Status](#) on page 421. You can also use the CLI commands `set default_revert_window` to configure the default revert time window for the Grid. For information about this command, refer to the *Infoblox CLI Guide*. Once a member exits the revert time window, you must revert the entire Grid in order to revert the member.

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

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

1. From the **Grid** tab, select the **Upgrade** tab.
2. Pause the upgrade.
3. Click **Toggle Member List View**, and then select the member check box.
4. From the Toolbar, click **Revert** -> **Revert Single Member**.

Grid Manager displays a message indicating that the revert process disrupts Grid services. Read the message carefully, and then click **Yes** to confirm your decision to revert the member. Be aware that when you revert a member, some changes made since the member was last upgraded may get lost.

Upgrade Process

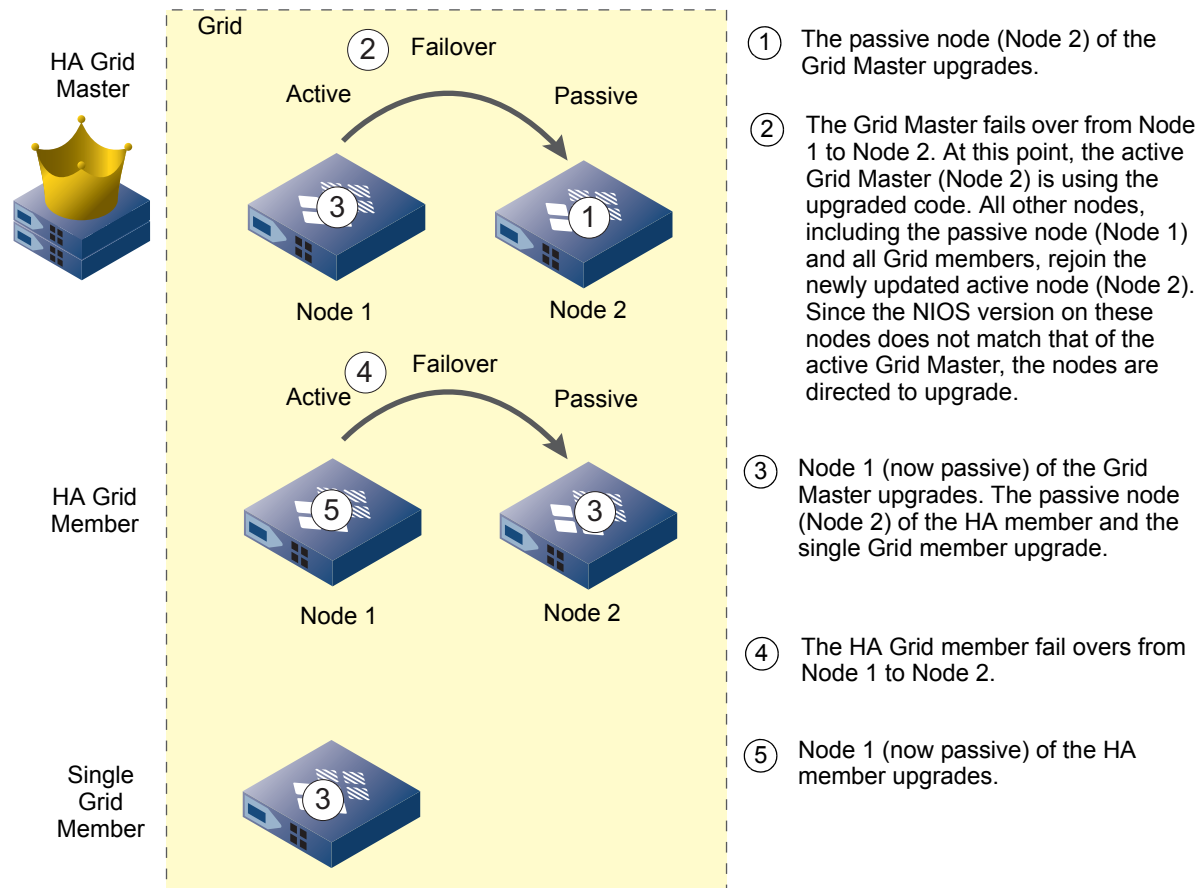
When an upgrade starts, Grid Manager checks if the nodes of an HA Grid Master have the same NIOS software version on their alternate partitions. If they do not have the same software version, the upgrade process stops. Grid Manager displays an error message and if it is a scheduled upgrade, Grid Manager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

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

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

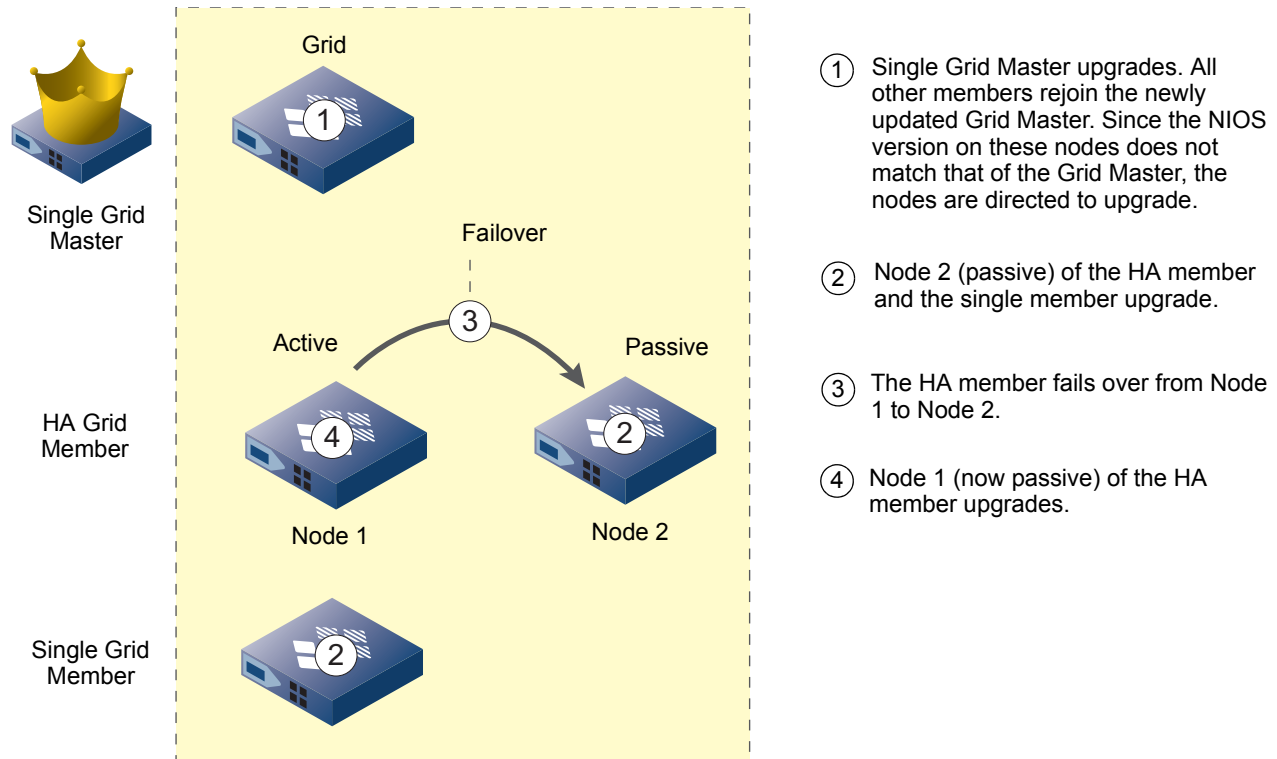
Due to the nature of the upgrade sequence, HA pairs fail over during the upgrade. Therefore, be aware that the active and passive nodes reverse roles. The order in which Grid members upgrade, including when HA pairs fail over, is shown in [Figure 9.2](#) (for an HA Grid Master) and [Figure 9.2](#) on page 420 (for a single Grid Master).

Figure 9.1 Upgrade Sequence for an HA Grid Master and Grid Members



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

Figure 9.2 Upgrade Sequence for a Single Grid Master and Grid Members



The Grid Manager session terminates when the HA Grid Master fails over from Node 1 to Node 2, or when the single Grid Master reboots and goes offline.

During a scheduled upgrade, the Grid members that have not upgraded yet can join the Grid and function normally until their scheduled upgrade time. When the upgrade finishes, the upgrade schedule is set to inactive.

Managing Upgrades

During an upgrade, Grid Manager displays a system message at the top of the screen indicating the Grid is being upgraded. After you start an upgrade, you can pause or resume it. For information, see [Pausing and Resuming Upgrades](#) and [Monitoring Distribution and Upgrade Status](#) on page 421.

Pausing and Resuming Upgrades

The following are some operational guidelines for performing an upgrade:

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

To pause an upgrade, from the Grid Upgrade Status bar, click the Pause icon. When you pause an upgrade, Grid Manager displays a system message at the top of the screen indicating the upgrade is paused, until you resume the upgrade. For information about the upgrade status of each member, see [Monitoring Distribution and Upgrade Status](#) on page 421.

To resume an upgrade:

1. From the Grid Upgrade Status bar, click the Resume icon.
2. When the appliance displays a dialog box confirming that you want to resume the upgrade, click **Yes** to continue. Members that have not completed or started upgrades that were scheduled at an earlier time resume the upgrade.

Monitoring Distribution and Upgrade Status

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

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

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

Grid and Member Status

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

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

Grid Manager displays the following information for each member:

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

The appliance automatically refreshes the information in this panel.

Upgrade Group Status

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

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

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

Detailed Status

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

To view detailed process information:

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

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

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

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

DOWNGRADING SOFTWARE

Each Infoblox appliance model has a minimum required release of Infoblox software. Before downgrading an appliance, refer to the document, *Minimum Required Release Software for Hardware Platforms*, that shipped with your product.

The downgrade procedure is for single independent appliances only. Infoblox does not support software downgrades for Grid members, but you can revert to the previous NIOS release (see the next section) on a Grid Master.

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

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

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

REVERTING THE GRID TO THE PREVIOUSLY RUNNING SOFTWARE

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

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

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

- From the **Grid** tab, select the **Upgrade** tab, and then click **Revert** -> **Revert Grid** from the Toolbar.
Grid Manager displays a warning indicating that the revert process disrupts Grid services. Read the warning carefully, and then click **Yes** to confirm your decision to revert.

BACKING UP AND RESTORING CONFIGURATION FILES

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

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

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

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

- [Backing Up Files](#)
- [Automatically Backing Up Data Files](#) on page 424
- [Manually Backing Up Data Files](#) on page 426
- [Restoring Backup Files](#) on page 428
- [Downloading Backup Files from a Different Appliance](#) on page 429

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

Backing Up Files

You can back up system files and discovery databases periodically and on demand. You can then restore the files on the same appliance or on a different appliance. For information about restoring files, see [Restoring Backup Files](#) on page 428. You can configure the appliance to automatically back up the files on a weekly, daily, or hourly basis.

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

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

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

Local Backup

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

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

Using TFTP

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

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

Using FTP

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

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

Using SCP

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

When you use SCP to back up the system files to an SSH server, you must specify the username and password the appliance uses to log on to the server. The user account must have write permission to the directory to which the appliance uploads the backup file. In addition, make sure that you enter the correct IP address of the SSH server; the appliance does not check the credentials of the SSH server to which it connects.

Automatically Backing Up Data Files

Infoblox recommends that you regularly back up your configuration files and/or discovery database files. The easiest way to accomplish this task is to configure the appliance for scheduled automatic backups of the NIOS configuration files. When you automatically back up a configuration file on the appliance, the file is named in the format `<GRIDNAME>_YYYY_MM_DD_HH.MM.tar.gz`. The default time for an automatic backup is 3:00 AM. Infoblox recommends scheduling configuration file backups to take place during the slowest period of network activity. You can choose a schedule for when and how often files are backed up: weekly, daily, or hourly.

You may also schedule automatic backups of the Discovery database, which consists of the complete discovery data for networks and network devices such as core, distribution and edge routers, enterprise switches, security devices, and end host devices. NIOS backs up the Discovery database in a .tar.gz file, with the raw Discovery data formatted as an XML file. For information on discovery features and requirements, see the chapter *Network Insight* on page 517.

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

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> Schedule Backup** from the Toolbar.
2. In the *Schedule Backup* dialog box, select the destination of the backup file from the **Backup to** drop-down list:

- **TFTP:** Back up system files to a TFTP server.
- **Keep local copy:** Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and Network Automation. Note that when you select this, the total backup time will increase.

- **IP Address of TFTP Server:** Enter the IP address of the TFTP server to which you want to back up the system files.
- **Directory Path:** Enter the directory path of the file. For example, you can enter `/archive/backups` on a Linux system, or `c:\archive\backups` on a Microsoft Windows system. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
- **Recurrence:** Select how often you want to back up the files. You can select **Weekly**, **Daily**, or **Hourly** from the drop-down list. When you select **Weekly**, complete the following:
 - **Every:** Choose a day of the week from the drop-down list.
 - **Time:** Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list. The Grid Master creates a backup file on the selected day and time every week.

When you select **Daily**, enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.

When you select **Hourly**, complete the following:

- **Minutes after the Hour:** Enter the minute after the hour when the Grid Master creates a backup file. For example, enter 5 if you want the Grid Master to create a backup file five minutes after the hour every hour.
- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now. You can still save the settings for future use.
- **FTP:** Back up system files to an FTP server.
- **Keep local copy:** Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and Network Automation. Note that when you select this, the total backup time will increase.
- **IP Address of FTP Server:** The IP address of the FTP server.
- **Directory Path:** Enter the directory path of the file. For example, you can enter `/archive/backups` on a Linux system, or `c:\archive\backups` on a Microsoft Windows system. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
- **Username:** Enter the username of your FTP account.
- **Password:** Enter the password of your FTP account.
- **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see TFTP.
- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now, but want to save the settings for future use.
- **SCP:** Back up system files to an SSH server that supports SCP.
- **Keep local copy:** Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and Network Automation. Note that when you select this, the total backup time will increase.

- **IP Address of SCP Server:** The IP address of the SCP server.
- **Directory Path:** Enter the directory path of the file. For example, you can enter `/archive/backups` on a Linux system, or `c:\archive\backups` on a Microsoft Windows system. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
- **Username:** Enter the username of your SCP account.
- **Password:** Enter the password of your SCP account.
- **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see the TFTP section.
- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now. You can still save the settings for future use.

Note: When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

- **Grid Master (Local):** Back up to a local directory on the Grid Master. This is the default.
By default, the Grid Master generates a backup file and saves it locally in its own storage at 3:00 AM daily. Be aware that backing up the Grid and saving it locally on an hourly basis increases the turnover of files stored on the Grid Master. Backing it up hourly to a remote server increases the overall amount of traffic on your network.

3. To back up NIOS configuration data for the Grid, select the **NIOS data** check box.
4. To back up Discovery data for the Grid, select the **Discovery data** check box.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Manually Backing Up Data Files

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

To back up manually:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> Manual Backup** from the Toolbar.
2. In the *Backup* wizard, select the destination of the backup file from the **Backup to** drop-down list:
 - **My Computer:** Back up system files to a local directory on your computer. This is the default.
 - **TFTP:** Back up system files to a TFTP server.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of TFTP Server:** Enter the IP address of the TFTP server to which you want to back up the system files.
 - **FTP:** Back up system files to an FTP server.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Username:** Enter the username of your FTP account.
 - **Password:** Enter the password of your FTP account.

- **SCP:** Back up system files to an SSH server that supports SCP.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of SCP Server:** The IP address of the SCP server.
 - **Username:** Enter the username of your SCP account.
 - **Password:** Enter the password of your SCP account.

Note: When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

3. To back up NIOS configuration data for the Grid, select the **NIOS data** check box.
4. To back up Discovery data for the Grid, select the **Discovery data** check box.
5. Click **Backup**.

Downloading Backup Files

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

To download an existing backup file:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup** → **Manage Local Backup** from the Toolbar. Grid Manager displays the current backup files in the *Manage Local Backups* dialog box.
2. To download a backup file, select the check box of a backup file, and then click the Transfer icon. You cannot select multiple files for downloading.
3. Select one of the following from the **Backup to** drop-down list:
 - **My Computer:** Backup to a local directory on your computer. This is the default.
 - **TFTP:** Save the backup file to a TFTP server.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of TFTP Server:** Enter the IP address of the TFTP server to which you want to save the backup file.
 - **FTP:** Save the backup file to an FTP server.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Username:** Enter the username of your FTP server account.
 - **Password:** Enter the password of your FTP server account.
 - **SCP:** Save the backup file to an SSH server that supports SCP.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of SCP Server:** The IP address of the SCP server.
 - **Username:** Enter the username of your SCP server account.
 - **Password:** Enter the password of your SCP server account.

Note: When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

4. Click **Transfer Copy**.

Restoring Backup Files

You can restore a backup file of a NIOS configuration or a Discovery database to an appliance running the same NIOS version as that of the appliance from which the backup file originates. You can also restore a backup file from an appliance running a NIOS version to an appliance running a later NIOS version as long as the upgrade from the earlier NIOS version to the later version is supported. For example, you can restore a backup file from an appliance running NIOS 4.3r6-1 to an appliance running NIOS 5.0r1-0 because upgrading from NIOS 4.3r6-1 to 5.0r1-0 is supported. However, you cannot restore a backup file from an appliance running NIOS 4.1r2-1 to an appliance running NIOS 5.0r1-0 because upgrading from NIOS 4.1r2-1 to 5.0r1-0 is not supported.

You can restore an existing backup file on the appliance from which it originates, or restore a backup file from a different appliance (referred to as a forced restore). To download a backup file from a different appliance, see [Downloading Backup Files from a Different Appliance](#) on page 429.

You must log in with a superuser account to back up and restore files.

NIOS provides three ways to restore a backup file:

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

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

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Restore** from the Toolbar.
 2. In the *Restore* dialog box, choose one of the following from the **Restore from** drop-down list:
 - **My Computer:** Restore a file from your local computer. This is the default.
 - **Filename:** Click **Select File** to navigate to the configuration file.
 - **TFTP:** Restore a file from a TFTP server.
 - **Filename:** Enter the directory path and the file name you want to restore. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of TFTP Server:** Enter the IP address of the TFTP server from which you restore the configuration file.
 - **FTP:** Restore a file from an FTP server.
 - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30` on a Linux server, or `c:\archive\backups\Infoblox_2009_10_20_15_30` on a Microsoft Windows server.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Username:** Enter the username of your FTP server account.
 - **Password:** Enter the password of your FTP server account.
 - **Grid Master (Local):** Restore from a local directory on the Grid Master. In the *Backup Set* table, select the file you want to restore.
3. To restore NIOS configuration data, select the **NIOS data** check box.
 4. To restore Discovery data, select the **Discovery data** check box. Discovery data should be restored to Consolidator appliances with the correct licensing.

5. To download a backup file from one appliance to a different appliance, select **Force Restore from Different Grid** to enable the feature, and then select one of the following:
 - **Retain Current Grid Master IP Settings** (this is the default)
 - **Overwrite Grid Master IP Settings**
6. Click **Restore**. In the *Confirm Restore* dialog box, click **Yes**.
After restoring the file, the appliance restarts. The restore process overwrites all existing data. All pending scheduled tasks are not restored or reverted.
7. Close your current browser window, wait a few minutes, and then reconnect to the NIOS appliance.

Downloading Backup Files from a Different Appliance

When you “force restore” a NIOS appliance, you download a backup file from one appliance to a different appliance. To restore a backup file to the same appliance or Grid Master, use the Restore function as described in [Restoring Backup Files](#) on page 428.

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

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Restore** from the Toolbar.
2. In the *Restore* wizard, do the following:
 - **Restore from:** Choose a source from which you restore the configuration file, as described in [Restoring Backup Files](#) on page 428.
3. Select **Force Restore from Different Grid** to enable the feature, and then select one of the following:
 - **Retain Current Grid Master IP Settings** (this is the default)
 - **Overwrite Grid Master IP Settings**
4. Click **Restore**. In the *Confirm Restore* dialog box, click **Yes**.
After restoring the file, the appliance reboots. The restore process overwrites all existing data. All pending scheduled tasks are not restored or reverted.
5. Close your current browser window, wait a few minutes, and then reconnect to the NIOS appliance.

DOWNLOADING SUPPORT BUNDLES

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

By default, the appliance includes the following files in the support bundle: core files, current logs, and rotated logs. Because core files can be quite large and take a significant amount of time to download, Infoblox recommends that you include core files in the support bundle only when requested by Infoblox Technical Support.

To download a support bundle:

1. From the **Grid** tab, select a *member* check box, and then click **Download** -> **Support Bundle** from the Toolbar.
2. In the *Download Support Bundle* dialog box, select the files you want to include in the support bundle, and then click **OK**:
 - **Core Files:** Infoblox recommends that you include these files only when requested by Infoblox Technical Support.
 - **Current Logs:** Infoblox recommends that you always include these files in the support bundle.
 - **Rotated Logs:** These are rotated logs that contain historical information.
 - **Discovery SNMP Logs:** Event logs related to device discovery SNMP probes of routers, switches and other network infrastructure devices.

3. Navigate to the location you want to save the file and change the file name. Do not change the `.tar.gz` file extension in the file name.
4. Send this file to Infoblox Technical Support.



Chapter 10 bloxTools Environment

The bloxTools environment provides a pre-installed environment for hosting custom web-based applications. This chapter includes the following sections:

- [*About the bloxTools Environment*](#) on page 432
 - [*System Requirements*](#) on page 432
- [*Using the bloxTools Environment*](#) on page 433
 - [*Configuring the Service*](#) on page 433
 - [*Allocating Memory*](#) on page 434
 - [*Uploading Files*](#) on page 434
 - [*Scheduling Tasks*](#) on page 435
 - [*Moving the bloxTools Service*](#) on page 435
- [*Monitoring the Service*](#) on page 435
 - [*Viewing the Logs*](#) on page 435
 - [*Viewing Detailed Status*](#) on page 436

ABOUT THE bloxTOOLS ENVIRONMENT

The bloxTools environment provides tools for creating custom applications that facilitate the administrative tasks in your organization. It provides a pre-installed environment for running applications using Perl, Python, PHP, CGI scripting, and Infoblox API libraries. Note that no direct external remote user (telnet and ssh, for example) or shell access is available in this environment.

The bloxTools environment “borrows” resources such as CPU, memory, disk space, and networking from the host Infoblox appliance, but is logically separated from the NIOS. The logical separation ensures that any failure in the bloxTools service does not affect the other services running on the appliance.

The bloxTools environment can only be configured to run on an independent appliance or a Grid member. You cannot run the bloxTools service on a Grid Master, a Grid Master candidate, or a virtual appliance, such as vNIOS for Riverbed, or VMware.

Note: In previous NIOS releases, you could run the bloxTools service only on a Grid Master. If bloxTools has been configured to run on a Grid Master before an upgrade, the bloxTools service continues to run on the Grid Master after an upgrade. This configuration is preserved mainly for migration purposes only. Infoblox strongly recommends that you move the bloxTools service to a Grid member after the upgrade. For information, see [Moving the bloxTools Service](#) on page 435.

In a Grid, you can run the bloxTools service only on one Grid member at a time, and you cannot configure this member as a Grid Master candidate. However, you can move the bloxTools service from one member to another. For information, see [Moving the bloxTools Service](#) on page 435.

On an HA member, the bloxTools service runs on the active node. If there is an HA failover, the bloxTools service is automatically launched after the passive node becomes active. For information, see [About HA Pairs](#) on page 233.

Note: When you run the bloxTools service on an independent appliance or a Grid member, the performance of other services running on the appliance may be affected. Infoblox recommends that you run the bloxTools environment on a member that does not host critical services.

After you enable the bloxTools service and configure its built-in file transfer services, you can upload content to the bloxTools portal using either an FTP (File Transfer Protocol) or SFTP (SSH File Transfer Protocol) client. The uploaded content is included in system backups and you can restore it from the backups.

For more information about the bloxTools environment and to access free applications, visit <https://www.bloxtools.com>.

Note: The bloxTools environment is not supported on vNIOS appliances on Riverbed, Cisco, VMware, and Microsoft Hyper-V.

System Requirements

[Table 10.1](#) shows which Infoblox appliances support the bloxTools service and the memory requirement for each. The service “borrows” host resources such as CPU, memory, and disk space from the host Infoblox appliance.

Table 10.1 Memory and Disk Space Requirements

Supported Infoblox Appliance	Memory Requirement
Infoblox-1550-A Infoblox-1552-A	128 MB to 2048 MB The default is 256 MB
Infoblox-1852-A Infoblox-2000-A	128 MB to 4096 MB The default is 256 MB

Table 10.1 Memory and Disk Space Requirements

Supported Infoblox Appliance	Memory Requirement
Trinzic 1410 Trinzic 1420	128 MB to 2048 MB The default is 256 MB
Trinzic 2210 Trinzic 2220	128 MB to 2048 MB The default is 256 MB
Infoblox-4010	128 MB to 4096 MB The default is 256 MB

USING THE BLOXTOOLS ENVIRONMENT

Complete the following tasks to upload custom applications to the bloxTools environment:

1. Log in to the appliance as a superuser and configure the bloxTools service, as described in [Configuring the Service](#).
2. Use an FTP or a SFTP client to upload content to the bloxTools environment.

In addition, you can schedule tasks as described in [Scheduling Tasks](#) on page 435, and monitor the bloxTools service as described in [Monitoring the Service](#) on page 435.

WARNING: RESETTNG THE GRID MEMBER USING EITHER THE RESET ALL OR RESET DATABASE CLI COMMANDS PERMANENTLY DELETES THE CONTENT YOU UPLOADED TO THE BLOXTOOLS ENVIRONMENT. INFOBLOX RECOMMENDS THAT YOU BACKUP THE APPLIANCE BEFORE USING ANY OF THESE COMMANDS.

Configuring the Service

When you configure the bloxTools service, you can enable FTP, SFTP, and HTTPS, and set their operational parameters. FTP and SFTP are the services you use to upload data. You can disable these services when they are not in use. HTTPS must remain enabled to allow the web based bloxTools applications to run. Note that the bloxTools service uses the same SSL certificate as the host Infoblox appliance. For information on certificates, see [Managing Certificates](#) on page 53.

To configure the bloxTools service:

1. Log in as a superuser.
2. From the **Grid** tab, select the **Grid Manager** tab, and then click **bloxTools**. In the **Services** tab, click **Edit -> Grid bloxTools Properties** from the Toolbar.
3. In the *Grid bloxTools Properties* editor, complete the following:
 - **Enable Web Service:** Select **HTTPS Port** to enable users to access the applications through an HTTPS connection. The default port is 444. You can change the port number to suit your environment.
 - **Enable FTP Service:** Select **FTP Port** to enable the FTP service. The default port is 26. You can change the port number to suit your environment.
 - **Enable SFTP Service:** Select **SFTP Port** to enable the SFTP service for secure file transfer. The default port is 28. You can change the port to a number between 1024 and 63999, provided that the port is not currently used for another purpose.
 - **Login:** Enter the username for the FTP and SFTP services. The username can contain lower case letters, numbers, underscores (_), and dollar signs (\$), and it must begin with a letter, not a number.
 - **Set Password:** Enter the password for the FTP and SFTP services in this field.
 - **Retype Password:** Enter the same password.

Note: The password is sent as clear text when you use the FTP service. To maintain security on the Infoblox appliance, this password should be different from the password set for the Infoblox appliance.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

When you configure the bloxTools service on an independent appliance, you can configure the allocated memory in the *System bloxTools Properties* editor. For information, see [Allocating Memory](#) on page 434.

Allocating Memory

You can configure the memory you want to allocate to the bloxTools service. You must configure this at the member level. If you run the bloxTools service on an independent appliance, you can configure the allocated memory in the *System bloxTools Properties* editor.

To configure the allocated memory:

1. Log in as a superuser.
2. From the **Grid** tab, select the **Grid Manager** tab, and then click **bloxTools**. In the **Services** tab, click **Edit -> Member bloxTools Properties** from the Toolbar.
3. In the *Member bloxTools Properties* editor, complete the following:
 - **Allocated Memory (MB):** The service “borrows” host resources such as CPU, memory, and disk space from the host Infoblox appliance. The default amount of memory the appliance allocates for the bloxTools environment is 256 MB. You can change this allocation, depending on the appliance platform. See [System Requirements](#) on page 432 for the requirements and allowed values of each appliance.

Uploading Files

Use an FTP or a SFTP client to upload content, such as Perl modules, JavaScript files, PHP files, CGI files, and image files, to the bloxTools environment. You can upload a maximum of 4 GB of data. After you have uploaded content to your bloxTools environment, you should disable the FTP and SFTP services to prevent unauthorized or accidental changes.

To upload files using the FTP service:

1. Open an Internet browser window and log in to the FTP service by entering:
`ftp://Grid_member_ip_addr:ftp_port`
 For example, if the IP address of the Grid member is 10.1.1.1 and the FTP port number is 26, enter:
`ftp://10.1.1.1:26`
2. In the *Authentication Required* dialog box, enter the username and password. This is the username and password you entered for the FTP service in the *bloxTools Environment* editor on the appliance.
3. Follow the instructions provided by your FTP client to upload the files.

To upload files using the SFTP service:

1. Open a terminal window and log in to the SFTP service by entering:
`sftp -oPort=sftp_port sftp_user@Grid_member_ip_addr`
 For example, if the IP address of the Grid member is 10.1.1.1, the login username for the SFTP service is jdoe, and the SFTP port number is 28, enter:
`sftp -oPort=28 jdoe@10.1.1.1`
2. Enter the password. This is the password you entered for the SFTP service in the *bloxTools Environment* editor on the appliance.
3. Follow the instructions provided by your SFTP client to upload the files.

Note: On a computer running Microsoft Windows, you can use WinSCP as the FTP or SFTP client for uploading files.

The bloxTools environment stores the uploaded data in the `/portal` directory.

Scheduling Tasks

bloxTools includes support for the Perl module `Config::Crontab` so you can manage scheduler services. You can use the scheduler to execute commands in the future. You can also schedule recurring commands. For example, you can schedule the creation of a host record or schedule recurring reports. The scheduler allows default “user level” crontab access and you can use the user account ‘nobody’ to submit commands. The Grid Master replicates the crontab data to the master candidates.

Moving the bloxTools Service

In a Grid, you can move the bloxTools service from one Grid member to another. When you move the bloxTools service, the source member synchronizes data with the Grid Master, and the Grid Master synchronizes data with the destination member. The time to resynchronize the bloxTools data on to the destination member depends on the amount of data to synchronize and the Grid configuration. If the migration takes longer than two minutes, it becomes a long running task. This allows the move of the bloxTools service to run in the background while you perform other tasks. For information, see [About Tasks](#) on page 72. Note that on an independent appliance, you cannot move the bloxTools service to another member.

After an upgrade from previous NIOS releases, Grid Manager displays a warning message in the system message panel if you have previously configured to run the bloxTools service on the Grid Master. You can click **Move** in this panel to launch the *Move bloxTools* dialog box to move the bloxTools service to a Grid member.

To move the bloxTools Service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **bloxTools** -> **Move** from the Toolbar.
2. In the *Move bloxTools* dialog box, complete the following:
 - **Source Member:** Displays the name of the Grid member that is currently running the bloxTools service. You cannot modify this field.
 - **Destination Member:** Click **Select**. In the *Member Selector* dialog box, select the member to which you want to move the bloxTools service. Grid Manager displays the name of the selected member here.
3. Click **Move**.

The appliance synchronizes data with the Grid Master, and the Grid Master synchronizes data with the destination member. This may take a while to complete depending on your Grid configuration and the amount of data.

MONITORING THE SERVICE

Infoblox provides several tools for monitoring the bloxTools Environment. The bloxTools Environment has its own syslog service which you can access to view logs generated by the bloxTools service and its processes. The *Detailed Status* panel also displays the status of the bloxTools Environment.

Viewing the Logs

The bloxTools Environment generates the following logs:

- `access.log`: The Apache access log
- `error.log`: The Apache error log
- `syslog.log`: The bloxTools Environment system log

These log files are included in the support bundle. You can download the log files using FTP. You can also connect to the CLI of the member running the bloxTools environment and use the following commands to view the logs:

- Use the `show file` command to view the list of log files.
- Use the `show bloxtools` command to view the status of the bloxTools Environment.
- Use the `show file bloxtools portal_access` command to view the web portal access log.
- Use the `show file bloxtools portal_error` command to view the web portal error log.

- Use the `show file bloxtools portal_log` command to view the web portal system log.





Viewing Detailed Status

You can view the status of the bloxTools Environment from the **Services** tab of the **Grid Manager** tab. To display the bloxTools service status, from the **Grid** tab, select the **Grid Manager** tab -> **Services** tab, and then click **bloxTools**. Grid Manager displays all Grid members that can host the bloxTools service. The name of the Grid Master is displayed only if you have completed an upgrade and previously configured the bloxTools service to run on the Grid Master. Though you can continue to run the bloxTools service on the Grid Master, Infoblox strongly recommends that you move the bloxTools service to a Grid member. For information, see [Moving the bloxTools Service](#) on page 435.

Grid Manager displays the following information about all Grid members:

- **Name:** The Grid member name.
- **Service Status:** Indicates the current operational status of the bloxTools service running on the member. This can include the migration status if you are moving the bloxTools service to another member.
- **IP Address:** The IP address of the member.
- **Comment:** Information about the bloxTools Environment.
- **Site:** The location to which the member belongs. This is one of the predefined extensible attributes.

The service status icon indicates the operational status of the bloxTools Environment and the usage percentages for the CPU, memory and disk resources. The status icon can be one of the following:

Icon	Color	Meaning
	Gray	The bloxTools Environment is disabled or offline.
	Green	The CPU, memory, and disk usage is below 80%.
	Yellow	Usage of at least one of the following resources is greater than or equal to 80%: CPU, memory or disk. The description indicates the percentage of each resource.
	Red	The bloxTools Environment is down, or an essential service within the bloxTools Environment has failed.



Chapter 11 RIR Registration Updates

This chapter explains how to configure the Infoblox Grid to manage RIR (Regional Internet Registries) allocated addresses and submit registration updates to the RIPE database. It includes the following sections:

- [*RIR Address Allocation and Registration Updates*](#) on page 438
 - [*About the RIPE Database*](#) on page 438
- [*Configuring RIR Registration Updates*](#) on page 439
 - [*Requirements and Permissions*](#) on page 438
 - [*Enabling Support for RIR Registration Updates*](#) on page 439
 - [*Configuring RIR Communication Settings*](#) on page 440
- [*Managing RIR Data*](#) on page 440
 - [*Adding RIR Organizations*](#) on page 441
 - [*Modifying RIR Organizations*](#) on page 442
 - [*Deleting RIR Organizations*](#) on page 442
 - [*Adding and Assigning RIR Networks*](#) on page 442
 - [*Deleting RIR Networks*](#) on page 444
- [*Managing RIR Attributes*](#) on page 444
 - [*RIR Organizational Attributes*](#) on page 445
 - [*RIR Network Attributes*](#) on page 447
- [*Monitoring RIR Data*](#) on page 451
 - [*Viewing RIR Organizations*](#) on page 451
 - [*Previewing Registration Updates*](#) on page 451

RIR ADDRESS ALLOCATION AND REGISTRATION UPDATES

You can configure the Infoblox Grid to manage allocated IP address blocks that ISPs (Internet Service Providers) receive from their RIRs (Regional Internet Registries). An RIR is an entity that manages the Internet number resources, which include IP addresses and autonomous system numbers, within a specific region of the world. RIRs use SWIP (Share WHOIS Project) or RWhois (Referral WHOIS) servers to provide address allocation information for IP address blocks. Typically, an RIR determines the address blocks to be allocated for specific organizations (typically ISPs), while an ISP manages the allocated address blocks, associated organizations and corresponding RIR registrations. An organization can determine when to request for more address blocks from its RIR. Most ISPs manage multiple organizations and synchronize network address data with their RIRs every few months.

To leverage IPAM (IP Address Management) on the NIOS appliance, you can enable the Infoblox Grid to manage RIR allocated addresses and send registration updates to the RIPE (Réseaux IP Européens) database as often as you update RIR data on NIOS. RIPE is one of the five RIRs in the world that manages the allocation and registration of Internet number resources for Europe, Russia, the Middle East, and Central Asia.

Note: The RIR registration update feature is not supported in a Multi-Grid configuration.

About the RIPE Database

The RIPE database contains registration details of IP addresses and AS numbers originally allocated by the RIPE NCC (RIPE Network Coordination Center). The database contains information such as organizations that hold IP resources, where the allocations were made, and contact details for the networks. Organizations or individuals that hold the allocated address blocks are responsible for updating information in the database.

The NIOS appliance supports submitting registration and reassignment updates to the RIPE database, which can be accessed through the RIPE API interface or an email template. For more information, see [Configuring RIR Communication Settings](#) on page 440.

Note: Before the NIOS appliance sends registration updates to the RIPE database, it does not validate the data you submit. Therefore, if you enter invalid information that cannot be mapped to the RIPE database, your updates will fail. In addition, the NIOS appliance does not synchronize data from the RIPE database.

Requirements and Permissions

To manage RIR allocated addresses, organizations, and network utilization that contain RIR assignments, you must first enable support for RIR registration updates, and then configure the RIR communication method. Note that once you have enabled support for RIR registration updates, settings and fields that are relevant to this feature are enabled in Grid Manager. You do not need a special license to use this feature.

Only superusers can create, modify, and delete RIR organizations. Limited-access users can manage RIR allocated address blocks if they have the required permissions to the objects.

To view and manage RIR related data, admins must have permissions to the applicable resources. For example, to view RIR networks, admins must have read-only permission to the networks; and to edit them, admins must have read/write permission to them. For more information about admin permissions, see [About Administrative Permissions](#) on page 160.

CONFIGURING RIR REGISTRATION UPDATES

To manage RIR allocated addresses and send registration updates through NIOS, you first add RIR organizations and create RIR allocated networks in NIOS. You can then reassign network addresses within the RIR allocated address block to other organizations based on your requirements, and then configure NIOS to send registration updates directly to the RIPE database. Any data you manage through the Grid is handled by the Grid Master.

When the Grid Master is an HA pair, the active node handles the submission of data. If an HA failover occurs during a submission, the failing node immediately aborts the submission. The new active node resumes the next submission. For information about HA pairs, see [About HA Pairs](#) on page 233.

To manage and submit updates to the RIPE database, you must first enable the Grid to support RIR registration updates. You can then enter RIR information, such as RIR organizations and RIR attributes.

To configure the Grid to manage RIR allocated addresses and submit updates to RIPE, complete the following:

1. Enable support for RIR registration updates, as described in [Enabling Support for RIR Registration Updates](#) on page 439.
2. Define the method to communicate updates to RIPE, as described in [Configuring RIR Communication Settings](#) on page 440.
3. Add and configure RIR organizations and RIR organizational attributes, as described in [Adding RIR Organizations](#) on page 441.
4. Add allocated address blocks and assign specific network addresses to RIR organizations, as described in [Adding and Assigning RIR Networks](#) on page 442.
5. Review and submit registration updates to RIPE, as described in [Previewing Registration Updates](#) on page 451.

You can also perform the following tasks:

- View a list of RIR organizations, as described in [Viewing RIR Organizations](#) on page 451.
- Modify RIR organizations and RIR organizational attributes, as described in [Modifying RIR Organizations](#) on page 442.
- Monitor network utilization for networks that contain RIR assignments. For information, see [Network List](#) on page 470 or [Viewing Networks](#) on page 848.

Enabling Support for RIR Registration Updates

Before you can manage RIR data through Grid Manager, you must first enable support for RIR registration updates.

To enable support for RIR registration updates:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab, and complete the following:
 - **Enable Updates of RIR Registrations:** Select this to enable the support for submitting RIR registration updates to the RIPE database. When you enable this feature, you can configure the appliance to send registration updates to RIPE for network reassignments and reallocations.

Note: Ensure that you configure DNS resolvers for the Grid when you enable this feature. For information about how to configure DNS resolvers, see [Enabling DNS Resolution](#) on page 376.

3. Save the configuration.

Configuring RIR Communication Settings

You can configure the appliance to send RIR address updates to RIPE through the RIPE REST API or through an email using the maintainer email address specified in the RIR organization. Note that when you use the API method to delete a registered address block, you do not need to submit RIR attributes that match the data in the RIPE database. However, when you use the email method, you must enter RIR attribute values that match the data in the database. Otherwise, your submission will fail. To view examples of registration updates that NIOS sends, see [Previewing Registration Updates](#) on page 451.

To configure the RIR communication settings:

1. From the **Administration** tab, select the **RIR Organizations** tab, and then select **RIR Settings -> RIPE** from the Toolbar.
2. In the *RIR Communication Settings - RIPE* editor, select one of the following to determine how the appliance sends updates to RIPE. The default is **API**.
 - **API:** The appliance sends RIR updates to RIPE through the RIPE API. The default destination is <https://rest.db.ripe.net> for accessing the production database and <https://rest-test.db.ripe.net> for accessing the test database. Click **Override** and enter a different URL to override the default value. When you select this as the communication method, the registration status will be updated automatically after the registration update is completed. Note that RIPE supports only secure connections using HTTPS.
 - **Email:** The appliance sends RIR updates to RIPE through the email address displayed in the field. The default is auto-dbm@ripe.net. Click **Override** and enter a different email address to override the default value. The appliance uses a special email template that includes values of certain RIR attributes. If any of the RIR attribute values do not match the database in the RIPE database, your submission will fail. When you select **Email** as the communication method, ensure that you enable email notifications at the Grid level. For information how to enable email notifications, see [Setting SNMP and Email Notifications](#) on page 1044. Note that when you select this as the communication method, the registration status will not be automatically updated. You can manually change the status. For information, see [Modifying RIR Network Data](#) on page 443.
 - **None:** The appliance does not send RIR updates to RIPE.
3. Save the configuration.

MANAGING RIR DATA

An RIR organization provides information about an entity that has registered a network resource in the RIPE Database. This entity can be a company (such as an ISP), a nonprofit group, or an individual. You can add RIR organizations defined in the RIPE database and start managing their data through NIOS.

After you have enabled support for RIR updates and configure the desired communication method for the updates, you can do the following to manage RIR data:

- Add RIR organizations and their associated data, as described in [Adding RIR Organizations](#) on page 441.
- Add the RIR allocated addresses to NIOS and assign specific address blocks to ISP organizations, as described in [Adding and Assigning RIR Networks](#) on page 442.
- View a list of organization objects, as described in [Viewing RIR Organizations](#) on page 451.
- Review the reassignment information before sending the updates to RIPE, as described in [Previewing Registration Updates](#) on page 451.
- Modify RIR organizational data and attributes, as described in [Modifying RIR Organizations](#) on page 442.
- Modify RIR network data and attributes, as described in [Modifying RIR Network Data](#) on page 443.
- Delete RIR organizations, as described in [Deleting RIR Organizations](#) on page 442.
- Delete delegated addresses from an organization, as described in [Deleting RIR Networks](#) on page 444.

Adding RIR Organizations

Before you can submit any RIR updates to the RIPE database, you must first add the RIR organization and its corresponding data to NIOS. You can also create additional organizations for ISP customers.

To add an organization:

1. From the **Administration** tab, select the **RIR Organizations** tab, and then click **Add -> RIPE Organization**.
2. In the *Add RIPE Organization* wizard, complete the following:
 - **Internet Registry:** The default is **RIPE**. This is the RIR that allocates address blocks to your organization. You cannot change this.
 - **Organization Name:** Enter the name of the organization that holds the resources allocated by RIPE NCC. You can enter up to 256 characters. Enter the name in this format: A list of words separated by white space. A word can be made up of letters, digits, the character underscore "_", and the character hyphen "-". The first character of a word must be a letter or digit and the last character of a word must be a letter, digit or a period. For example, you can enter **SPRINT REGION2**.
 - **Organization ID:** Enter the handle or ID of the organization. You can enter up to 23 characters. Enter the ID in this format: Start with **ORG-** followed by two to four characters, then followed by up to five digits and a source specification. Note that the first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length. For example, you can enter **ORG-CA1-RIPE** or **ORG-CB2-TEST**.
 - **Maintainer:** Enter the name of the maintainer for this organization. This is required. You can enter up to 256 characters; however, note that the RIPE database has an 80 characters limit for this field. A maintainer is any registrant or person to whom the authority to update has been delegated by another registrant either directly or indirectly, and who holds an identifier that allows updates to be authenticated and authorized. Data entered here must match exactly how the maintainer appears in RIPE.
 Enter the maintainer name in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit. You cannot use the following words (they are reserved by RPSL): any, as-any, rs-any, peer, as, and, or, not, atomic, from, to, at, action, accept, announce, except, refine, networks, into, inbound, outbound. Also note the following: Names starting with certain prefixes are reserved for certain object types. For example, names starting with "as-" are reserved for as set names. Names starting with "rs-" are reserved for route set names. Names starting with "rtrs-" are reserved for router set names. Names starting with "fltr-" are reserved for filter set names. Names starting with "prng-" are reserved for peering set names. Names starting with "irt-" are reserved for irt names.
 - **Password:** Enter the maintainer password. This is required. You can enter up to 256 characters.
 - **Retype Password:** Enter the same password.
 - **Maintainer Email:** Enter the originating or source email address of the maintainer. This is required.

RIR Organizational Attributes: This table lists all predefined RIR attributes associated with the RIR organization. Click the **Value** field of an attribute in the table to enter a value. The **Required** field indicates whether a value for the corresponding attribute is required.

You can add custom attributes by clicking the Add icon and select an attribute from the drop-down list. You can also delete an RIR attribute by selecting its check box and clicking the Delete icon.

For information about the attributes and how to enter their values, see [RIR Organizational Attributes](#) on page 445.

Note: You cannot leave an optional RIR attribute value empty. If you do not have a value for an RIR attribute, you must delete it from the table. You can enter up to 256 characters for all RIR attributes.

3. Save the configuration. Note that you cannot schedule the creation, modification, or deletion of an RIR organization.

Modifying RIR Organizations

To modify an RIR organization:

1. From the **Administration** tab, select the **RIR Organizations** tab -> *rir_organization* check box, and click the Edit icon.
2. In the *Organization* editor, modify the organization information, as described in [Adding RIR Organizations](#) on page 441. You can also reorder the list of RIR organizational attributes using the up and down arrows.
3. Save the configuration.

Deleting RIR Organizations

You can delete an RIR organization that does not have any networks assigned to it. When you delete an RIR organization, the appliance moves it to the Recycle Bin, if enabled. You can later restore the network if needed. For information about the Recycle Bin, see [Using the Recycle Bin](#) on page 64.

To remove an RIR organization:

1. From the **Administration** tab, select the **RIR Organizations** tab -> *rir_organization* check box, and then click the Delete icon.
2. In the *Delete Confirmation (RIR Organization)* dialog box, click **Yes**.

Adding and Assigning RIR Networks

Before you can assign network addresses within an RIR allocated address block to an organization, you must first add the allocated address block to NIOS. Infoblox supports IPv4 and IPv6 network containers and networks. You can also create network templates that are specific for RIR networks. For information about creating network templates, see [About IPv4 Network Templates](#) on page 829 and [About IPv6 Network Templates](#) on page 837.

Note that when you add network containers or networks to NIOS, the appliance does not validate whether the corresponding networks actually exist in the RIPE database. Even though you can create the networks in NIOS, the submission of updates for the network may fail. For example, if you create a child network and the parent network is not registered in RIPE, the registration update will fail.

In addition, each network can only be associated with an RIR in one network view. If you have a network address block registered with RIPE in a specific network view, you must not register the same address block in a different network view.

When you enable the support for updates of RIR registrations, Grid Manager displays the appropriate data fields that you can use to add or modify RIR related networks. You can do the following to add IPv4 and IPv6 networks:

- Add RIR allocated IPv4 networks to NIOS, or assign addresses to specific organizations. For information see [Adding IPv4 Networks](#) on page 845.
- Add RIR allocated IPv6 networks to NIOS, or assign addresses to specific organizations. For information, see [Adding IPv6 Networks](#) on page 871.
- Add IPv4 network templates that are specific to RIR address allocation. For information, see [About IPv4 Network Templates](#) on page 829.
- Add IPv6 network templates that are specific to RIR address allocation. For information, see [About IPv6 Network Templates](#) on page 837.

You can also do the following to modify specific data about the RIR networks:

- Modify RIR allocated or assigned IPv4 networks. For information see [Modifying IPv4 Networks](#) on page 851.
- Modify RIR allocated or assigned IPv6 networks. For information see [Modifying IPv6 Networks](#) on page 874.
- Modify IPv4 network templates that are specific to RIR address allocation. For information, see [Modifying IPv4 Network Templates](#) on page 831.
- Modify IPv6 network templates that are specific to RIR address allocation. For information, see [Modifying IPv6 Network Templates](#) on page 838.

You can preview the information before the appliance submits updates to the RIPE database. To preview registration updates, click **Preview RIR Submissions** in the *Add IPv4 Network* or *Add IPv6 Network* wizards. For more information, see [Previewing Registration Updates](#) on page 451.

Note: You can also add RIR networks through **Task Dashboard**. For information, see [The Tasks Dashboard](#) on page 99.

After you create an RIR network container or network, you can perform the following:

- Split a network that has an organization ID. A child network that is created does not contain an organization ID by default. You must assign an organization ID to the child network after splitting it. For information about splitting an RIR network, see [Splitting IPv4 Networks into Subnets](#) on page 472 and [Splitting IPv6 Networks into Subnets](#) on page 485.
- Resize an IPv4 RIR network that contains an organization ID and has been registered with RIPE. For more information, see [Resizing IPv4 Networks](#) on page 472.

Viewing RIR networks

You can view a list of IPv4 and IPv6 RIR networks in the **Data Management** tab -> **IPAM** tab or the **Data Management** tab -> **DHCP** tab -> **Networks** tab -> **Networks** section. For more information, see [Network List](#) on page 470 and [Viewing Networks](#) on page 848.

Modifying RIR Network Data

You can modify certain RIR network information in the **RIR Registration** tab of the *IPv4 and IPv6 Network* editors.

To modify RIR network information, complete the following:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* check box, and then click the Edit icon.
or
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* check box, and then click the Edit icon.
2. In the *IPv4* or *IPv6 Network Container* or *Network* editor, click the **RIR Registration** tab, and then complete the following to modify RIR related data for the IPv4 or IPv6 network container or network:
 - **Internet Registry:** Displays the RIR that allocates RIR address blocks. The default is **RIPE**. You cannot change this.
 - **Organization ID:** Displays the organization ID with which this network is associated. You cannot change this.
 - **Registration Status:** Displays the current registration status. This can be **Registered** or **Not Registered**. **Registered** indicates that the network has a corresponding entry in the RIPE database. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. You can modify this by selecting the appropriate status from the drop-down list.
 - **Status of last update:** Displays the registration status, communication method, timestamp of the last registration update. The status can be Pending, Sent, Succeeded, or Failed. The displayed timestamp reflects the timestamp used on the Grid Master. Each time you send a registration update to create, modify, or delete a network container or network, the updated status and timestamp will be displayed here. If you have selected not to send the registration update, the previous status and timestamp are retained.
 - **Registration Action:** From the drop-down list, select what you want to do with the RIR network updates. If you are creating a top-level network block that has already been assigned to the organization, select **None**. If you are creating a child network within the allocated address block, you can select one of the following:
 - **None:** The appliance does not submit the updates.
 - **Create:** The appliance creates the network container or network for the specified organization.
 - **Modify:** Modifies data for this network container or network.

- **Delete:** Deletes the RIR network from the organization. When you select this, you must enter a reason for deleting this entry in the **Delete Reason** field.
- **Do not update registrations:** By default, the appliance sends updates to RIPE if you specify **Create**, **Modify**, or **Delete** as the registration action. Select this if you do not want the appliance to submit updates to the RIPE database.

RIR Network Attributes: Modify the value of RIR network attributes by clicking the **Value** field of an attribute and entering a new value. You can add a new RIR network attribute by clicking the Add icon and selecting an attribute from the drop-down list. You can also select any optional attributes and click the Delete icon to delete them. For information about RIR network attributes, see [RIR Network Attributes](#) on page 447.

You can enter up to 256 characters for all RIR network attributes, unless otherwise noted.

Preview RIR Submissions: Click this to view the updates before the appliance submits them to the RIPE database. This button is enabled only when the registration action is **Create**, **Modify**, or **Delete**, and the **Do not update registrations** check box is not selected. For more information, see [Previewing Registration Updates](#) on page 451.

To schedule this task, click the Schedule icon at the top of the wizard. In the Schedule Change panel, click **Later**, and then specify a date, time, and time zone.

3. Save the configuration.

Deleting RIR Networks

When you delete an RIR network or network container, the appliance moves it to the Recycle Bin, if enabled. You must enter the reason for deleting the RIR network or network container and indicate whether you want to send the deletion update to RIPE. You can delete multiple networks at the same time.

To delete an RIR network or network container:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* check box, and then click the Delete icon. You can choose to delete the network immediately or schedule its deletion.
2. In the *Delete Confirmation (IPv4 or IPv6 Network)* or *Schedule Deletion* dialog box, complete the following:
 - **Justification:** Enter the reason for deleting this network.
 - **Do not update registrations:** Select this check box if you do not want the appliance to submit updates to RIPE.
3. Optionally, you can click **Preview RIR Submissions** to view the RIR network information before deleting the network. Grid Manager displays the preview data in a separate browser window. For information, see [Previewing Registration Updates](#) on page 451
4. Click **Yes**. If you are scheduling a deletion, enter the data and time for execution, and then click **Schedule Deletion**.

MANAGING RIR ATTRIBUTES

Before you can successfully submit RIR data updates, you must ensure that all RIR required attributes contain valid values that can be mapped to data in the RIPE database. The appliance does not validate data with the RIPE database before you submit your updates. The appliance also does not synchronize data from the database.

Note: RIPE does not support UTF-8 data in the **Description** and **Remarks** fields. After an upgrade, the NIOS appliance keeps the UTF-8 data in these fields. However, if you want to modify these fields after the upgrade, you must remove the UTF-8 data before you can save the changes.

When you enter a value for the following RIR attributes that cannot be mapped to a valid reference in the RIPE database, updates to the RIR database will fail. However, these values will still be displayed in the IPv4 or IPv6 network or network container panels of Grid Manager.

- RIPE Routes Maintainer
- RIPE Lower Level Maintainer
- RIPE Reverse Domain Maintainer
- RIPE Admin Contact
- RIPE Technical Contact
- RIPE Computer Security Incident Response Team

You can add multiple values for certain RIR attributes. When you add multiple values of the same attribute, the appliance groups the values in the order they are listed in the attribute table. You can also reorder the RIR attributes using the up and down arrows in the attribute tables.

RIR Organizational Attributes

The following table lists RIR organizational attributes, the format you must use to enter values, and whether they are required or optional.

Organizational Attribute	Corresponding RIPE Attribute	Description and Format	Required/Optional
RIPE Description	descr	Enter a short description about the organization.	Optional
RIPE Country	country	From the drop-down menu, select the country name, followed by the two-letter ISO 3166 country code, of the country or area within the RIPE NCC service region or through Local Internet Registries.	Required
RIPE Admin Contact	admin-c	Enter the name of the on-site admin contact for the organization. Enter the name in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required
RIPE Technical Contact	tech-c	Enter the name of the technical contact for the organization. Enter the name in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required
RIPE Remarks	remarks	Enter remarks about the organization.	Optional
RIPE Notify	notify	Enter the email address to which notifications of changes to the organization will be sent.	Optional
RIPE Changed	changed	Enter the email address of the person who submitted changes to the RIPE database when organization data was updated, followed by the timestamp of the updates in YYYYMMDD format. For example, you can enter <code>jdoe@corp100.com 20120908</code> .	Optional

Organizational Attribute	Corresponding RIPE Attribute	Description and Format	Required/Optional
RIPE Registry Source	source	From the drop-down list, select the registry at which the organization is registered. The default is RIPE . Select RIPE for the RIPE database, which is the authoritative database. Select TEST for the RIPE TEST database that operates in the same way as the RIPE database but contains only test data. Note that test data is cleaned out at the start of each month and a predetermined set of basic objects is re-inserted. You can use the RIPE TEST database to learn how to update the database and try out special scenarios. The RIPE TEST database has fewer restrictions which allows you to create encompassing or parent objects you may need for testing.	Optional
RIPE Organization Type	org-type	<p>From the drop-down list, select one of the following organization type:</p> <ul style="list-style-type: none"> • IANA for Internet Assigned Numbers Authority • RIR for Regional Internet Registry • NIR for National Internet Registry • LIR for Local Internet Registry • WHITEPAGES for special industry people • DIRECT_ASSIGNMENT for direct contract with RIPE NCC • OTHER for all other organizations <p>Note: Only the RIPE database admin can set the organization type, and there are no NIRs in the RIPE NCC service region.</p>	Optional
RIPE Address	address	Enter the organization address.	Optional
RIPE Phone Number	phone	<p>Enter the organization phone number in numeric format starting with the + character, followed by the country code, area code, and the phone number. For example, you can enter +18089991000. You can also use one of the following formats:</p> <ul style="list-style-type: none"> • '+' <integer-list> • '+' <integer-list> "(" <integer-list> ")" <integer-list> • '+' <integer-list> ext. <integer list> • '+' <integer-list> "(" integer list ")" <integer-list> ext. <integer-list> 	Optional

Organizational Attribute	Corresponding RIPE Attribute	Description and Format	Required/Optional
RIPE Fax Number	fax-no	Enter the organization fax number in numeric format starting with the + character, followed by the country code, area code, and the fax number. For example, you can enter +16052529000 . You can also use one of the following formats: <ul style="list-style-type: none"> '+' <integer-list> '+' <integer-list> "(" <integer-list> ")" <integer-list> '+' <integer-list> ext. <integer list> '+' <integer-list> "(" integer list ")" <integer-list> ext. <integer-list> 	Optional
RIPE Email	email	Enter the organization email address.	Required
RIPE Abuse Mailbox	abuse-mailbox	Enter the email address to which abuse complaints are sent.	Optional
RIPE Reference Notify	ref-nfy	Enter the email address to which notifications are sent when a reference to the organization object is added or removed.	Optional

RIR Network Attributes

When you create or edit an RIR associated network, ensure that you enter valid values for the RIR network attributes. The following table lists RIR network attributes, the format you must use to enter values, and whether they are required or optional:

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/Optional
RIPE Admin Contact	admin-c	The name of the on-site admin contact for the network address. This attribute is populated from the organizational attribute. You can modify the value in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required
RIPE Changed	changed	Enter the email address of the person who submitted changes to the RIPE database when the network was updated, followed by the timestamp of the updates in YYYYMMDD format. For example, you can enter jdoe@corp100.com 20120908 .	Required

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/Optional
RIPE Computer Security Incident Response Team	mnt-irt	<p>The name of the Computer Security Incident Response Team (CSIRT) that handles security incidents for the network address.</p> <p>You can enter the value in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The value must start with "irt-", and the last character of a name must be a letter or a digit. You must enter a minimum of five characters.</p>	Optional
RIPE Country	country	The two-letter ISO 3166 country code of the country within the RIPE NCC service region or through Local Internet Registries. This attribute is populated from the organizational attribute. You can select a different country code from the drop-down list.	Required
RIPE Description	descr	Enter a short description about the network.	Required
RIPE IPv4 Status	status	<p>The status of the IPv4 network address. From the drop-down list, select one of the following status:</p> <p>ALLOCATED PA ALLOCATED PI ALLOCATED UNSPECIFIED LIR-PARTITIONED PA LIR-PARTITIONED PI SUB-ALLOCATED PA ASSIGNED PA ASSIGNED PI ASSIGNED ANYCAST EARLY-REGISTRATION NOT-SET</p>	Required
RIPE IPv6 Status	status	<p>The status of the IPv6 network address. From the drop-down list, select one of the following:</p> <p>ALLOCATED-BY-RIR ALLOCATED-BY-LIR ASSIGNED ASSIGNED ANYCAST ASSIGNED PI</p>	Required

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/Optional
RIPE Lower Level Maintainer	mnt-lower	<p>Enter the name of the registered maintainer for hierarchical authorization purposes. This can protect the creation of networks directly (one level) below in the hierarchy of a network container or another network. The authentication method of the maintainer will be used upon creation of any network directly below the network that contains the "mnt-lower:" attribute.</p> <p>Enter the maintainer name in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit. You cannot use the following words (they are reserved by RPSL): any, as-any, rs-any, peer, as, and, or, not, atomic, from, to, at, action, accept, announce, except, refine, networks, into, inbound, outbound. Also note the following: Names starting with certain prefixes are reserved for certain object types. Names starting with "as-" are reserved for as set names. Names starting with "rs-" are reserved for route set names. Names starting with "rtrs-" are reserved for router set names. Names starting with "fltr-" are reserved for filter set names. Names starting with "prng-" are reserved for peering set names. Names starting with "irt-" are reserved for irt names.</p>	Optional
RIPE Network Name	netname	The name of the IP address range. You can enter up to 80 characters. Enter the network name in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit.	Required
RIPE Notify	notify	Enter the email address to which notifications of changes to the object must be sent.	Optional
RIPE Registry Source	source	From the drop-down list, select the registry at which the organization is registered. The default is RIPE . Select RIPE for the RIPE database, which is the authoritative database. Select TEST for the RIPE TEST database that operates in the same way as the RIPE database but contains only test data. Note that test data is cleaned out at the start of each month and a predetermined set of basic objects is re-inserted. You can use the RIPE TEST database to learn how to update the database and try out special scenarios. The RIPE TEST database has fewer restrictions which allows you to create encompassing or parent objects you may need for testing.	Required
RIPE Remarks	remarks	Enter remarks about the network.	Optional

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/Optional
RIPE Reverse Domain Maintainer	mnt-domains	<p>Enter the name of a registered maintainer used for reverse domain authorization. This can protect domain objects. The authentication method of this maintainer will be used for any encompassing reverse domain object.</p> <p>Enter the maintainer name in this format: You can use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit. You cannot use the following words (they are reserved by RPSL): any, as-any, rs-any, peer, as, and, or, not, atomic, from, to, at, action, accept, announce, except, refine, networks, into, inbound, outbound. Also note the following: Names starting with certain prefixes are reserved for certain object types. Names starting with "as-" are reserved for as set names. Names starting with "rs-" are reserved for route set names. Names starting with "rtrs-" are reserved for router set names. Names starting with "fltr-" are reserved for filter set names. Names starting with "prng-" are reserved for peering set names. Names starting with "irt-" are reserved for irt names.</p>	Optional
RIPE Routes Maintainer	mnt-routes	<p>This attribute references a maintainer that is used in determining authorization for the creation of route objects. Enter the name in this format: Start with the reference to the maintainer, followed by an optional list of prefix ranges inside of curly brackets or the keyword "ANY". The default, when no additional set items are specified, is "ANY". For more information, refer to RFC-2622. Example: <mnt-name> [{ list of <address-prefix-range> } ANY].</p>	Optional
RIPE Technical Contact	tech-c	<p>The name of the technical contact for the network. Enter the name in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.</p>	Required

MONITORING RIR DATA

You can view RIR organizations and networks you added to NIOS through Grid Manager. The appliances sends SNMP traps and email notifications about registration updates. It also logs RIR events in the Infoblox syslog. Note that sometimes due to network timeout from RIPE, your registration updates may fail.

You can do the following to monitor RIR data:

- View RIR update events in the syslog, as described in [Viewing the Syslog](#) on page 1016.
- View RIR organizations, as described in [Viewing RIR Organizations](#) on page 451.
- View RIR IPv4 and IPv6 network containers and networks, as described in [Network List](#) on page 470 and [Viewing Networks](#) on page 848.
- Preview RIR updates before submitting them to RIPE, as described in [Previewing Registration Updates](#) on page 451.

Viewing RIR Organizations

You can view the list of RIR organizations that have received address allocation and the ones you have added associated networks.

To view RIR organizations:

1. From the **Administration** tab, select the **RIR Organizations** tab.
2. Grid Manager displays the following information for each RIR organization:
 - **Organization ID:** The RIR organization ID.
 - **RIR:** The RIR that allocates the address block to the organization.
 - **Maintainer:** The name of the maintainer for the organization.

You can also select **Organization Name** and RIR organizational attributes for display.

You can do the following in this tab:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click Save to save the changes. Note that some fields are read only.
- Sort the data in ascending or descending order by column.
- Select an organization and click the Edit icon to modify data, or click the Delete icon to delete it.
- Click the Permissions icon to configure permissions for the admin account.
- Use filters and the Go to function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the Go to field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information about quick filters, see [Using Quick Filters](#) on page 68.
- Print and export the data in this tab.

Previewing Registration Updates

Before the appliance submits RIR updates to RIPE, you can preview the data in Grid Manager. The appliance uses the email template when displaying preview data in a separate browser window.

Preview data includes the subject line for the email, followed by the inetnum or inet6num template for the network and other associated data, such as network name, organization name, and others. When there are multiple operations involved, such as deleting multiple networks, the preview data includes a separate subject line for each operation.

You can preview registration updates when you create a new RIR network. In the Add Networks wizard, click **Preview RIR Submissions** in the wizard. For information about how to create or assign RIR networks, see [Adding and Assigning RIR Networks](#) on page 442.

Following is a sample preview for a network creation request:

```
Subject: CREATE inetnum 100.200.0.0 - 100.200.255.255 KEYWORDS: NEW
inetnum: 100.200.0.0 - 100.200.255.255
netname: corp100_network
descr: RIR network for Corp100.
status: ASSIGNED PA
org: ORG-MC1-TEST
country: US
source: TEST
changed: jdoe@corp100.com 20120809
notify: jdoe@corp100.com
admin-c: NP1-TEST
tech-c: NP1-TEST
mnt-by: JohnDoe
password: ***
password: ***
```

Following is a sample preview for a network modification request:

```
Subect: MODIFY inetnum 100.200.0.0 - 100.200.255.255 KEYWORDS:
inetnum: 100.200.0.0 - 100.200.255.255
netname: corp100_network
descr: RIR network for Corp100.
status: ASSIGNED PA
org: ORG-MC1-TEST
country: US
source: TEST
changed: jdoe@corp100.com 20120809
notify: jdoe@corp100.com
admin-c: NP1-TEST
tech-c: NP1-TEST
mnt-by: JohnDoe
password: ***
```

Following is a sample preview for deleting multiple networks:

```
Subect: DELETE inetnum 100.200.0.0 - 100.200.255.255 KEYWORDS:
inetnum: 100.200.0.0 - 100.200.255.255
netname: corp100_network
descr: RIR network for Corp100.
status: ASSIGNED PA
org: ORG-MC1-TEST
country: US
source: TEST
changed: jdoe@corp100.com 20120809
notify: jdoe@corp100.com
```



```
admin-c: NP1-TEST
tech-c: NP1-TEST
mnt-by: JohnDoe
password: ***
delete: Removed network.
```

```
Subect: DELETE inetnum 100.300.0.0 - 100.300.255.255 KEYWORDS:
inetnum: 100.300.0.0 - 100.300.255.255
netname: corp200_network
descr: RIR network for Corp200.
status: ASSIGNED PA
org: ORG-MC1-TEST
country: US
source: TEST
changed: jsmith@corp200.com 20120809
notify: jsmith@corp200.com
admin-c: NP1-TEST
tech-c: NP1-TEST
mnt-by: JohnSmith
password: ***
delete: Removed network.
```




PART 3 IP ADDRESS MANAGEMENT

IPAM (IP Address Management) is the allocation, administration, reporting, and tracking of IP addresses, network devices, and their associated data. This section provides information about IPAM and how to use the Infoblox tools to perform IPAM tasks and manage your entire IP network. It includes the following chapters:

- [Chapter 12, *IP Address Management*](#), on page 457
- [Chapter 13, *Network Discovery*](#), on page 493
- [Chapter 14, *Network Insight*](#), on page 517



Chapter 12 IP Address Management

This chapter describes how to manage your networks and IP addresses through the Infoblox IPAM (IP Address Management) implementation. It contains the following sections:

- [About IP Address Management](#) on page 458
- [About Host Records](#) on page 459
 - [Assigning Multiple IP Addresses to a Host](#) on page 461
 - [Adding Host Records](#) on page 462
 - [Modifying Host Records](#) on page 463
- [About Network Containers](#) on page 464
 - [Adding IPv4 and IPv6 Network Containers and Networks](#) on page 465
 - [Modifying IPv4 and IPv6 Network Containers and Networks](#) on page 465
 - [Deleting Network Containers](#) on page 465
- [Managing IPv4 Networks](#) on page 466
 - [IPv4 Network Map](#) on page 467
 - [Network List](#) on page 470
 - [Resizing IPv4 Networks](#) on page 472
 - [Splitting IPv4 Networks into Subnets](#) on page 472
 - [Joining IPv4 Networks](#) on page 473
 - [Discovering Networks \(Under Network Insight only\)](#) on page 473
- [Viewing and Managing IPv4 Addresses](#) on page 474
 - [IP Map](#) on page 474
 - [IP Address List](#) on page 476
 - [Managing IPv4 Addresses](#) on page 479
- [Managing IPv6 Networks](#) on page 480
 - [IPv6 Network Map](#) on page 480
 - [IPv6 Network List](#) on page 484
 - [Splitting IPv6 Networks into Subnets](#) on page 485
 - [Joining IPv6 Networks](#) on page 485
- [Viewing IPv6 Data](#) on page 486
- [Managing IPv4 and IPv6 Addresses](#) on page 487
 - [Converting Objects Associated with IP Addresses](#) on page 487
 - [Reclaiming Objects Associated with IPv4 and IPv6 Addresses](#) on page 491
 - [Pinging IP Addresses](#) on page 491
 - [Clearing Active DHCP Leases](#) on page 491

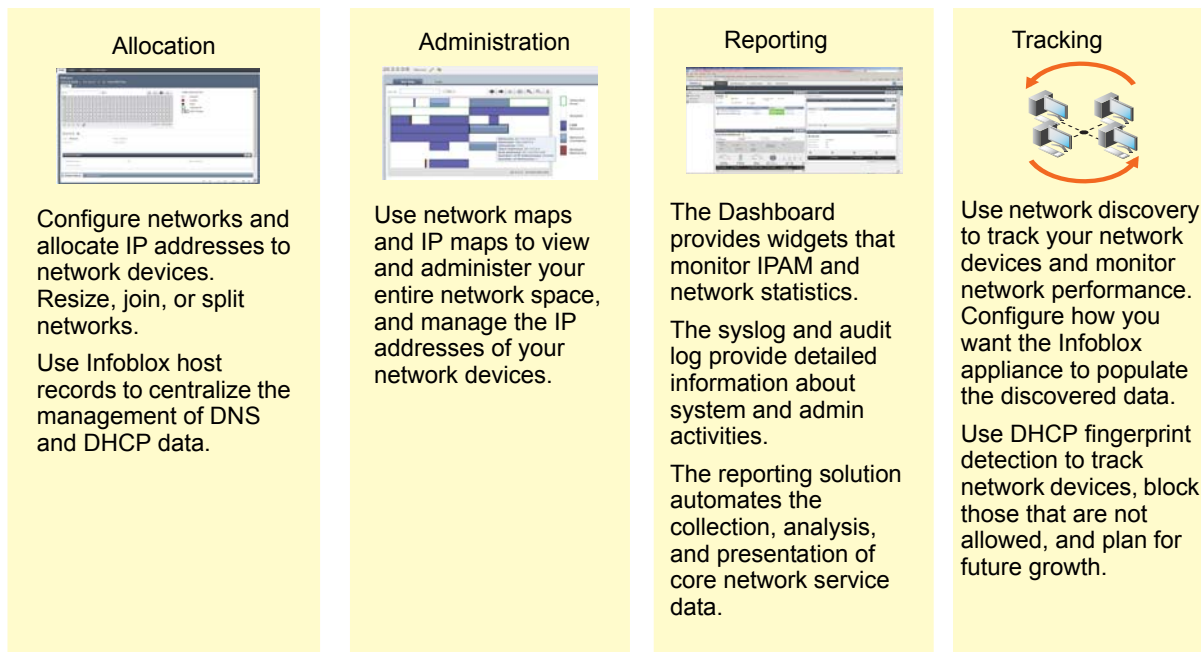
ABOUT IP ADDRESS MANAGEMENT

IPAM is the allocation, administration, reporting, and tracking of public and private IP spaces, network devices, and their associated data. It comprises the deployment of DNS and DHCP services and the monitoring of network devices and performance to ensure data integrity and security of your networks.

The Infoblox IPAM implementation is a feature-rich and easy-to-use solution that encompasses support for IPv4, IPv6, network discovery, and automated monitoring.

Infoblox IPAM provides tools that integrate the allocation, administration, reporting, and tracking of your entire network space. [Figure 12.1](#) highlights the Infoblox IPAM implementation.

Figure 12.1 Infoblox IPAM Features



You can perform the following IPAM tasks to effectively manage and control your network:

- Create host records. Host records integrate the DNS records and DHCP data of a network device. You can use a host record to manage a network device from one central point. For more information, see [About Host Records](#) on page 459.
- Create IPv4 and IPv6 networks. For information, see [About Network Containers](#) on page 464.
- Discover devices in IPv4 and IPv6 networks and other Objects created in IPAM. See [Adding IPv4 and IPv6 Network Containers and Networks](#) on page 465, [Adding Host Records](#) on page 462 and other sections throughout this chapter. View your IPv4 network and address utilization in a graphical mode. For information, see [IPv4 Network Map](#) on page 467 and [IP Map](#) on page 474.
- When necessary, resize, join, or split networks. For information, see [Resizing IPv4 Networks](#) on page 472, [Splitting IPv4 Networks into Subnets](#) on page 472, [Joining IPv4 Networks](#) on page 473, [Splitting IPv6 Networks into Subnets](#) on page 485, and [Joining IPv6 Networks](#) on page 485.
- Manage IPv4 and IPv6 address data. For information, see [Viewing and Managing IPv4 Addresses](#) on page 474, [Viewing IPv6 Data](#) on page 486, and [Managing IPv4 and IPv6 Addresses](#) on page 487.
- Add and manage DNS resource records associated with IP addresses. For information, see [DNS Resource Records](#) on page 655.
- Monitor your core service network data using the Dashboard, audit log, syslog, and reports. For information, see [Part 7 Monitoring and Reporting](#) on page 1001.
- Discover and track network devices. For information, see [Network Discovery](#) on page 493 and [DHCP Fingerprint Detection](#) on page 1031.

ABOUT NETWORK INSIGHT

In some deployments, NIOS IPAM offers powerful network infrastructure device discovery capabilities. NIOS discovery features include both Network Discovery, and Network Insight, which is a superset of Network Discovery that supports SNMP and that enables detailed viewing and assessments of network devices in managed and unmanaged networks, including networks in routed paths, security infrastructure devices, and networks within switched Ethernet segments.

Deployment of Network Insight discovery requires a separate Discovery license and one or more NIOS appliances dedicated to discovery tasks. Types of discovered and catalogued devices include Ethernet switches (including enterprise L3 switches), routers, wireless routers and access points, firewalls, hosts and other devices in end host networks, and much more. Network Insight applies discovery in the following ways:

- Through specification of seed routers, which inform discovery of the various networks that should be examined and catalogued;
- Through discovery of the various Object types you can create under IPAM and DHCP, including IPv4 and IPv6 Networks, Host Records, IPv4 reservations, DHCP ranges, and IPv4/IPv6 fixed addresses.
- An enhanced VM Discovery allowing both scheduled VM discovery and immediate discovery of VMs.

For more information about network discovery, see the chapter [Network Discovery](#) and its various sections.

For a complete discussion of Network Insight, see the chapter [Network Insight](#) and its various sections.

Viewing the Complete List of Discovered Devices

The **Data Management → Devices** list provides a complete view of all discovered devices catalogued by discovery under Network Insight. The list includes routers, switches, firewalls and other security devices, wireless APs, end hosts and servers in end-host networks. Use NIOS standard filtering to narrow down the status table to the devices or values you want to examine.

Click the Action icon for any table row to view the following:

- Click **Interfaces** to display the complete table of interfaces associated with the network device.
- Choose **Networks** and choose any network prefix in the list. The **IP Map** page appears for the chosen network for which the device is a part.
- Choose **Device Details** to view a short list of key information about the selected device.

Values listed in the **Discovered Devices** table include the following:

- **IP Address:** Detected IP address (IPv4 or IPv6).
- **Name:** Detected name of the interface, if any.
- **Type:** The network device type: **Router**, **Switch-Router**, **Firewall**, **NIOS** (Infoblox appliance), **vNIOS**, and others.
- **Model:** The model name as detected from the device during Discovery.
- **Vendor:** The equipment manufacturer.
- **Device Version:** The Operating System version for the network device.
- **Location:** The physical location of the network device as detected from the device during Discovery.
- **Description:** Verbose description of the network device as collected from the device by Discovery.

You can click **Discovery Status** in the Toolbar to view the same list of network devices showing the Discovery data set. This table can also be filtered.

ABOUT HOST RECORDS

Host records provide a unique approach to the management of DNS, DHCP, and IPAM data. By using host records, you can manage multiple DNS records and DHCP and IPAM data collectively, as one object on the appliance.

When you create a host record, you are specifying the name-to-address and address-to-name mappings for the IP address that you assign to the host. The Infoblox DNS server then uses this data to respond to DNS queries for the host. When the server receives a name-to-address query, it responds with an A record for an IPv4 host or an AAAA record for an IPv6 host that contains the data from the host record. Likewise, when it receives an address-to-name query for the host, the appliance responds with a PTR record that contains data from the host record. Additionally, if you specify an alias in the host record, the appliance uses this data as a CNAME record to respond to queries with the alias. It maps the alias to the canonical name and sends back a response with the canonical name and IP address of the host. Thus, a single host record is equivalent to creating A, PTR, and CNAME resource records for an IPv4 address and AAAA and PTR records for an IPv6 address. The appliance supports IDNs for a host record. You can specify alias and domain names in the native character set. For information about IDN support, see [Support for Internationalized Domain Names](#) on page 93.

Hosts also support prefix delegation for IPv6. For example, you can specify an IPv6 prefix in the host record of a router. The router then advertises this prefix on one of its interfaces, so hosts that connect to the interface can generate their IP addresses, using the stateless autoconfiguration mechanism defined in *RFC 2462, IPv6 Stateless Autoconfiguration*.

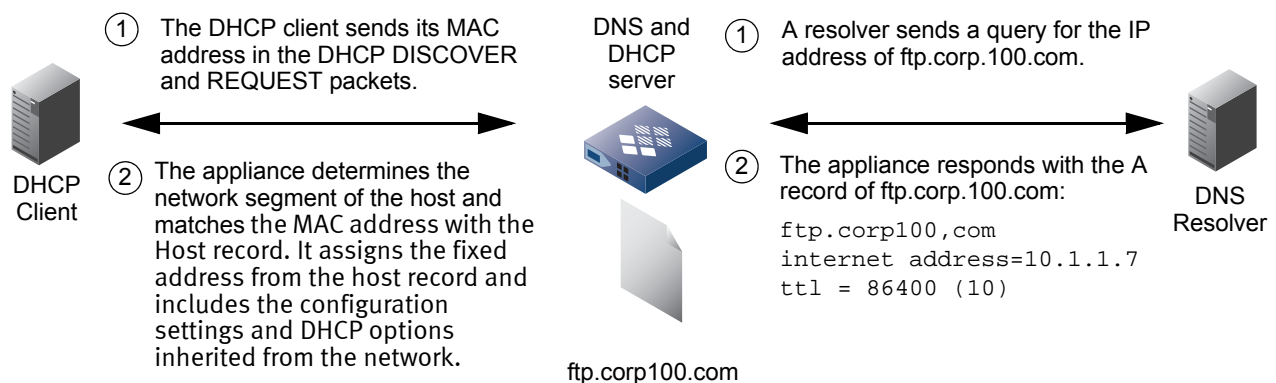
In addition, if the Infoblox DHCP server manages the IP address assigned to the host, the server uses it as a fixed address record as well. The DHCP server assigns the IP address to the host when it receives a DHCP request with the matching MAC address or DUID. Its response includes configuration information, and any DHCP options defined for the host or inherited from the network to which the fixed address belongs. You can also assign multiple IPv4 and IPv6 addresses to a host, as described in [Assigning Multiple IP Addresses to a Host](#) on page 461. Note that you can modify the host record of an IPv4 and make it a reservation, as well. For information, see [Configuring IPv4 Reservations](#) on page 860.

You can execute immediate Discovery on a host record. This simple setting enables you to determine the precise type of device that is associated with the host, along with its IP addresses, its name and other information.

You can define extensible attributes for a host record to further describe the device. You can include information such as its location and owner for IP address management purposes. For information about extensible attributes, see [About Extensible Attributes](#) on page 322.

[Figure 12.2](#) illustrates how the appliance uses the host record for both DHCP and DNS.

Figure 12.2 Using the Host Record for DHCP and DNS



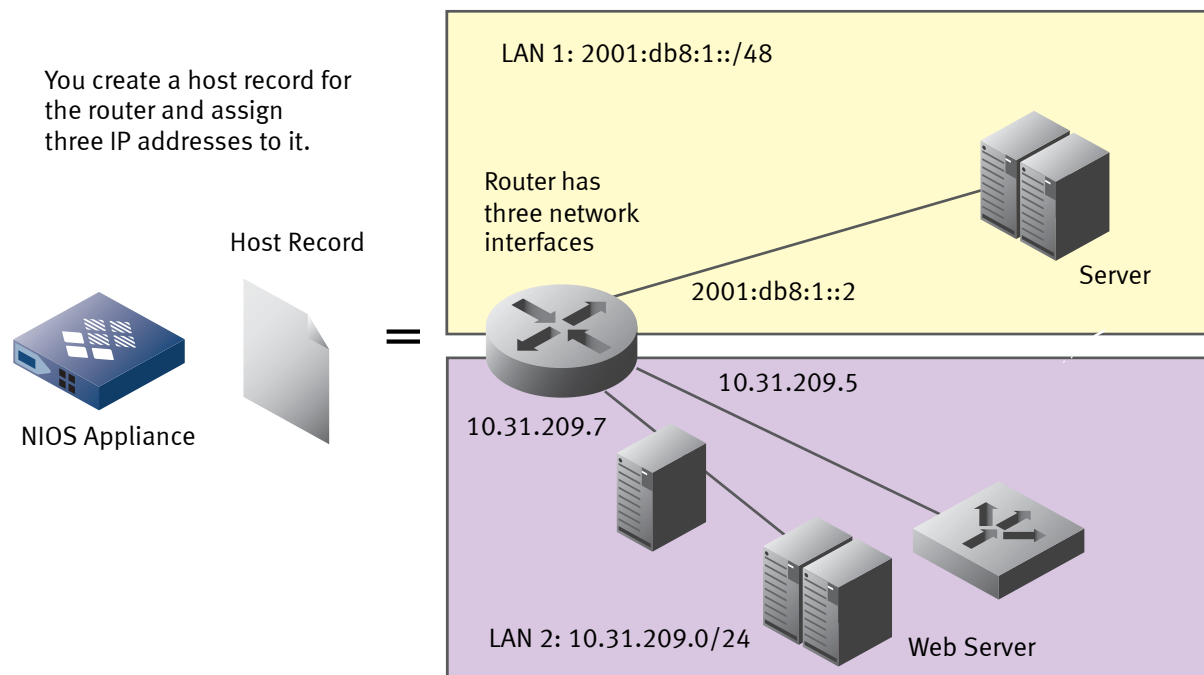
Note that If the zone of the host record is associated with networks, the IP addresses must belong to the associated networks. For example, if the host record is in the `corp100.com` zone, which is associated with `10.1.0.0/16` network, then the IP addresses of the host record must belong to the `10.1.0.0/16` network. For information about associating zones and networks, see [Associating Networks with Zones](#) on page 813.

Assigning Multiple IP Addresses to a Host

You can assign multiple IPv4 and IPv6 addresses to a host depending on the function of the device. For example, you can create a host record for a router that supports three network interfaces in two different networks, and assign IP addresses to each interface, as illustrated in [Figure 12.3](#). When the DNS server responds to DNS queries for the host, it includes an A or AAAA record for each IP address.

In addition, if the IP addresses belong to different networks, they can have different DHCP configurations and options. As shown in [Figure 12.3](#), the configuration information and DHCP options of the interface with the IPv6 address 2001:db8:1::2 may be different from the other two interfaces, 10.31.209.5 and 10.31.209.7, because it is in a different network.

Figure 12.3 Assigning Multiple IP Addresses to one Host Record



Adding Host Records

You can add host records from the Toolbar of the **IPAM**, **DHCP** and **DNS** tabs of the **Data Management** tab and from the Tasks Dashboard. For information about the Tasks Dashboard, see [The Tasks Dashboard](#) on page 99.

When you create a host record, you must specify its zone and at least one IP address. If the zone of the host record is associated with one or more networks, the IP addresses must belong to one of the associated networks. If a zone of a host record contains IDNs, the appliance displays the zone name in the native character set.

To add a host record from the **Data Management** tab:

1. From the **IPAM**, **DHCP** or **DNS** tab of the **Data Management** tab, expand the Toolbar.
2. Click **Add** and select the option to add a host record from the drop-down menu.
3. In the *Add Host Record* wizard, do the following:
 - **Name:** If Grid Manager displays a zone name, enter the host name here. The displayed zone name can either be the last selected zone or the zone from which you are adding the host record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter a unique name for the host. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is corp100.com and you enter admin, then the FQDN is admin.corp100.com.
 - **Enable in DNS:** This is selected by default. It enables DNS service for the host. If you clear this check box, DNS does not serve this host and you cannot assign it to a zone.
 - **DNS View:** Displays the DNS view for the host record. This appears only when you enable the host record in DNS.
 - **Host Name Policy:** Displays the host name policy of the selected zone. This appears only when you enable the host record in DNS.
 - **RRset Order:** Select one of the following RRset orders that the appliance uses to return A and AAAA records of the host. This check box appears only when you have enabled the configuration of RRset order for the Grid and there are multiple IP addresses in this host record. For information about how to enable this feature, see [Enabling the Configuration of RRset Orders](#) on page 561.
 - **Cyclic:** The records are returned in a round robin pattern. This is the default.
 - **Fixed:** The records are returned in the order you specify in this host record. When you select this check box, the appliance displays up and down arrows next to the IPv4 and IPv6 address tables. You can use these arrows to reorder the address list. The appliance returns the A and AAAA records of this host based on the order you define in the address tables.
 - **Random:** The records are returned in a random order.

Note that when you specify **Fixed** as the RRset order, the appliance places the resource records as follows:

- A and AAAA records of the host in the fixed order you specify in the address tables. Note that the order of the returned A and AAAA records are independent of each other.
- Other A and AAAA records in an undefined order.
- Other record types in the default cyclic order.

For more information about RRset order, see [Enabling the Configuration of RRset Orders](#) on page 561.

- In the **IPv4 Addresses** and **IPv6 Addresses** sections, specify the IP addresses of the host record. Click the Add icon do one of the following:
 - Select **Next Available IP Address** to retrieve the next available IP address in a network. Infoblox recommends this option to ensure that you assign an IP address from the appropriate network. If the host record is in a zone that has one associated network, Grid Manager retrieves the next available IP address in that network. When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

If the host record is in zone that has multiple associated networks, the *Network Select* dialog box lists the associated networks. If the zone has no network associations, the *Network Select* dialog box lists the available networks. When you select a network, Grid Manager retrieves the next available IP address in that network.

If you want to enter a link-local IPv6 address, you must enter an IPv4 address and the host MAC address first, and then click the Add (+) icon again to enter the link-local IPv6 address. When you select the link-local IPv6 address, the MAC address is automatically filled in. For information, see [Understanding DNS for IPv6](#) on page 552.

Optionally, you can delete an IP address from the host by selecting an IP address in the table and clicking the Delete icon.

or

- Select **Add Address** to enter an IPv4 or IPv6 address. You can also enter an IPv6 prefix. Note that when you use this option, you could specify an IP address from a network that has not yet been defined. To avoid this, use the **Next Available IP Address** option instead.
- **MAC Address:** For an IPv4 address, enter the MAC address of the network device associated with this host IP address. Note that you must enter a MAC address if DHCP is enabled for the host IP address.

or

- **DUID:** For an IPv6 address, enter the DHCP Unique Identifier (DUID) of the network device associated with this host IP address. Note that you must enter a DUID if DHCP is enabled for an IPv6 host address.
- **DHCP:** Select this to enable the DHCP services to manage the host IP address. If you do not select this option, the host IP address is not managed by the DHCP server.

- **Comment:** Optionally, enter additional information about the host record.

- **Disable:** Select this option to temporarily disable the host record. For example, you might want to disable a host when you need to update the network device.

4. (*Applies only with Network Insight*) Click **Next** to initiate or disable discovery of the new host.

- Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the host record. If you choose **Enable Immediate Discovery**, discovery will execute on the host after you save your settings. You may also choose to leave both options disabled.
- By default, the new host inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** check box and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials. See the section [Configuring Grid SNMPv1/v2 Properties](#) on page 530 and [Configuring Grid SNMPv3 Properties](#) on page 530 for a complete description of SNMP credentials for discovery.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

Modifying Host Records

To modify a host record:


1. From the **Data Management** tab, select the **IPAM**, **DHCP**, or **DNS** tab.
2. In the selected application, search for or navigate to the host record that you want to modify
3. Select the record and click the Edit icon. Grid Manager displays the *Host Record* editor.
4. The *Host Record* editor provides the following tabs from which you can modify data:
 - **General:** Modify the information you entered through the wizard as described in [Adding Host Records](#) on page 462.

- **TTL:** This tab displays the default TTL settings the record inherited from the Grid or the DNS zone, if you enabled override TTL settings at the zone level. You can keep the default settings or override them. To override the inherited value, click **Override** to enable the configuration. Specify how long the record is cached. Select the time period in seconds, minutes, hours, days, or weeks from the drop-down list. To enable the record to inherit the Grid or zone TTL settings, click **Inherit**.
 - **Aliases:** Click the Add icon. Grid Manager displays a new row in the table. Enter a fully qualified domain name (a CNAME record for the host) in the **Aliases** column. You can delete an alias by selecting the alias check box and clicking the Delete icon.
 - **IPv4 Discovered Data:** Displays the discovered data of the IPv4 addresses, if any, of the host record. For information, see [Viewing Discovered Data](#) on page 510.
 - **IPv6 Discovered Data:** Displays the discovered data of the IPv6 addresses, if any, of the host record. For information, see [Viewing Discovered Data](#) on page 510.
 - **Extensible Attributes:** You can add and delete extensible attributes that are associated with a host record. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#) on page 160.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

ABOUT NETWORK CONTAINERS

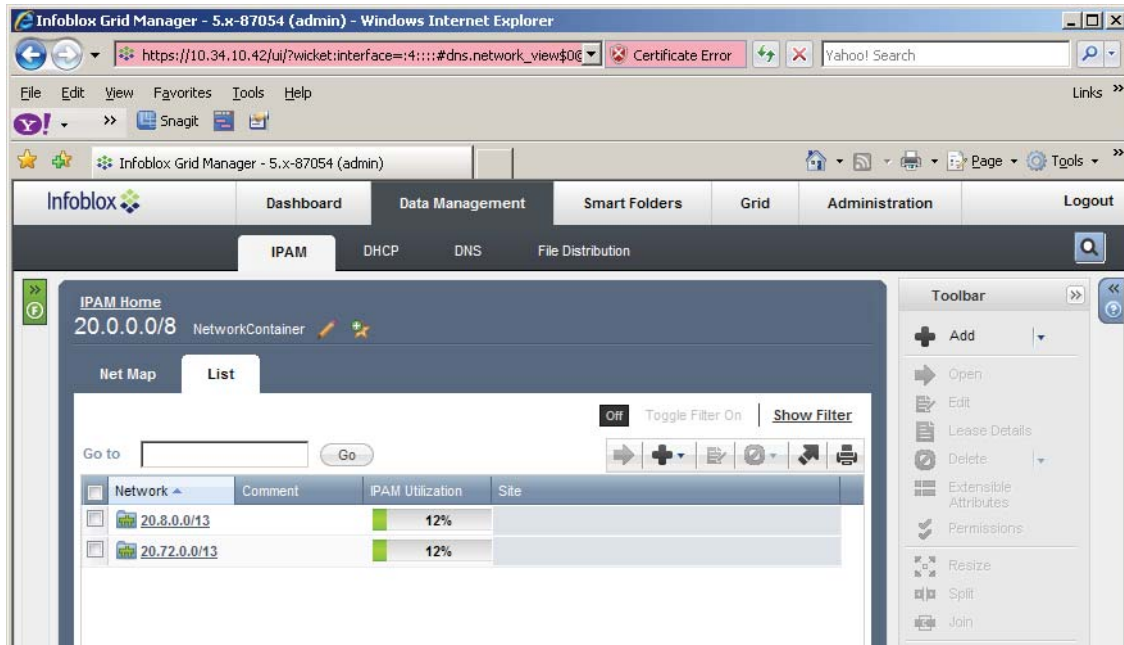
Grid Manager uses network containers to group IPv4 and IPv6 networks. A network container is a parent network that contain other network containers and leaf networks. A leaf network is a network that does not contain other networks. For example, [Figure 12.4](#) illustrates the IPv4 20.0.0.0/8 network, which is a network container with two network containers, 20.8.0.0/13 and 20.72.0.0/13. The 20.8.0.0/13 network has two leaf networks, 20.8.0.0/16 and 20.9.0.0/16. The 20.72.0.0/13 network has one leaf network, 20.72.0.0/16.

Figure 12.4 IPv4 Network Container



From Grid Manager, you can click the link of the network container 20.0.0.0/8 in the IP List panel and drill down to the two network containers, 20.8.0.0/13 and 20.7.0.0/13, as shown in [Figure 12.5](#). You can click the network container links to drill down further to the leaf networks.

Figure 12.5 IP List View of Network Containers



In the **IPAM** tab, when you create an IPv4 or IPv6 network that belongs to a larger network, the appliance automatically creates a network container and puts the leaf network in the container. The appliance also creates network containers when you split IPv4 or IPv6 networks into smaller networks. For information, see [Splitting IPv4 Networks into Subnets](#) on page 472 and [Splitting IPv6 Networks into Subnets](#) on page 485.

Adding IPv4 and IPv6 Network Containers and Networks

To add an IPv4 or IPv6 network container or network:

1. From the **Data Management** tab, select the **IPAM** tab.
2. Click the Add icon and select either **IPv4 Network** or **IPv6 Network**.
3. In the *Add Network* wizard, create a network as described in [Adding IPv4 Networks](#) on page 845 or [Adding IPv6 Networks](#) on page 871.

Modifying IPv4 and IPv6 Network Containers and Networks

You can modify existing network settings, with the exception of the network address and subnet mask.

To modify an IPv4 or IPv6 network container or network:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* check box, and then click the Edit icon.
2. In the *DHCP Network* editor, modify the network settings as described in [Modifying IPv4 Networks](#) on page 851 or [Modifying IPv6 Networks](#) on page 874.

Deleting Network Containers

Depending on the configuration, you may or may not be able to delete or schedule the deletion of a network container and all its contents. Contents in a network container can include other network containers, leaf networks, and associated objects. For recursive deletions, only network containers and networks are considered. Objects such as hosts are not considered for recursive deletions.

Superusers can determine which group of users are allowed to delete or schedule the deletion of a network container and all its contents. For information about how to configure the recursive deletion of network containers, see [Configuring Recursive Deletions of Networks and Zones](#) on page 269.

Note that you must have Read/Write permission to all the contents in order to delete a network container. When you delete a network container only, the appliance reparents the other network containers and leaf networks.

The appliance puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary. When you restore a parent object from the Recycle Bin, all its contents, if any, are re-parented to the restored parent object. For information about the Recycle Bin, see [Using the Recycle Bin](#) on page 64.

To delete a network container:

1. From the **Data Management** tab, select the **IPAM** tab -> *network_container* check box. You can select multiple network containers for deletion.
2. Click the Delete icon.
3. Do one of the following in the *Delete Confirmation* dialog box:
 - Select one of the following. Note that these options appear only if you are allowed to delete the network container and all its contents. For information about how to configure this, see [Configuring Recursive Deletions of Networks and Zones](#) on page 269.
 - **Delete only the network container and re-parent the subnets:** Select this to delete only the network container and re-parent its subnets.
 - **Delete the network container and all its subnetworks:** Select this to delete both the network and its contents.
 - Click **Yes**.

The appliance puts the deleted network container in the Recycle Bin, if enabled. You can also schedule the deletion for a later time. Click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#) on page 76. For information about scheduling recursive deletions of network containers, see [Scheduling Recursive Deletions of Network Containers and Zones](#) on page 76.

MANAGING IPV4 NETWORKS

In Grid Manager, you use the Net Map (network map) and List panels to manage your IPv4 network infrastructure. After you select a network container from the IPAM tab, Grid Manager displays it in the Net Map panel, by default. The Net Map panel provides a graphical view of your networks and has a number of features that simplify network management. The List panel displays the networks in table format.

You can always switch your view of a network container between the Net Map and List panels. Grid Manager keeps track of which panel you last used. When you select a network container, Grid Manager displays it in the Net Map or List panel, depending on which one you last used. For information about each panel, see [IPv4 Network Map](#) on page 467 and [Network List](#) on page 470.

Use the IP Map and List panels to manage the IP addresses in leaf networks. For information, see [Viewing and Managing IPv4 Addresses](#) on page 474.

After you create an IPv4 network, you can modify its properties, resize it, use the split network feature to create subnets, enable discovery to discover routers, switches, firewalls, wireless access points and other device types within it, or join it to another network to create a larger network that encompasses adjacent subnets. You can do the following from both the Net Map and List panels:

- Resize a network. For information, see [Resizing IPv4 Networks](#) on page 472.
- Split a network into subnets. For information, see [Splitting IPv4 Networks into Subnets](#) on page 472.
- Join a network. For information, see [Joining IPv4 Networks](#) on page 473.
- Discover devices in the network. For information, see [Discovering Networks \(Under Network Insight only\)](#) on page 473.

IPv4 Network Map

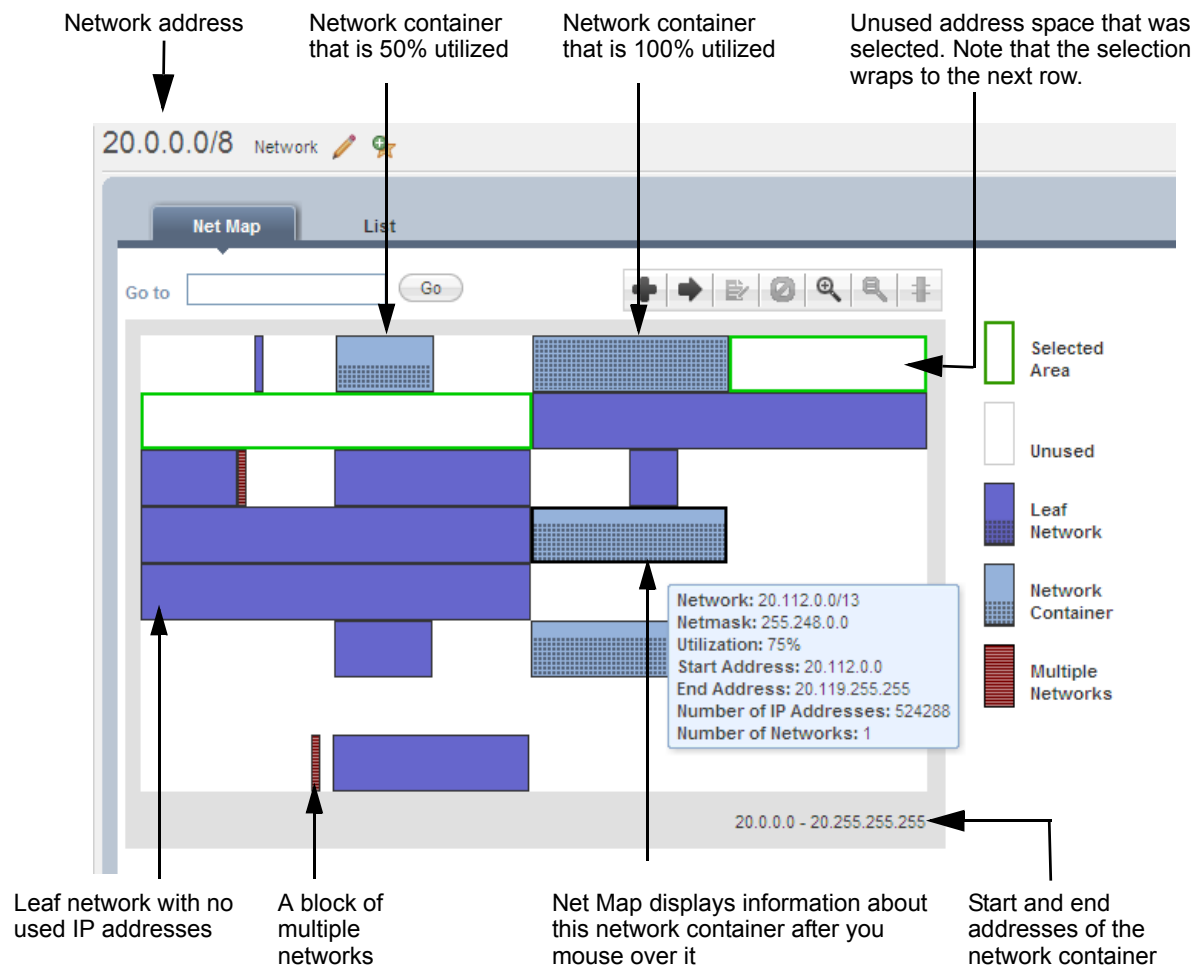
After you select an IPv4 network container from the IPAM tab, Grid Manager displays it in the Net Map (network map) panel, by default. Net Map provides a high-level view of your IPv4 network address space. You can use Net Map to design and plan your network infrastructure, configure and manage individual networks, and evaluate their utilization. Its unique display of the IPv4 network address space across multiple rows is similar to a road map that starts with the first IP address in the network and ends with the last address. Net Map displays the network address space across a maximum of eight rows, depending on the size of the network. It automatically scales the map so that it displays the entire address space of a network container.

The Net Map panel presents a complete view of the network space, including the different types of networks that are in it and its unused address space. IP addresses that belong to a network are blocked off. Each color-coded block represents a network container, a leaf network, or a block of networks that are too small to be displayed individually in the map. For example, in a /8 or /16 network, networks smaller than /20 or /28 respectively and that are beside each other are represented as a multiple network block. In addition, the fill pattern of the blocks indicates their utilization. Therefore, you can quickly evaluate how many and what type of networks are in a network container, their relative sizes, utilization, and how much space you have left.

As you mouse over areas of the map, it displays IP information about the area. Net Map also has a zoom feature that allows you to enlarge or reduce your view of a particular area.

[Figure 12.6](#) displays the network map of a 20.0.0.0/8 network, which is a network container that has network containers and leaf networks.

Figure 12.6 20.0.0.0/8 Network Map



Displaying IP Information

As shown in [Figure 12.6](#), as you mouse over the map, Net Map displays IP information about the area. When you mouse over an unused area, Net Map displays the following information:

- The start and end IP address
- The number of IP addresses that can fit in that space
- The largest possible network
- The number of /16 and /24 networks that can fit in that space

When you mouse over a network, Net Map displays the following information:

- Network address and netmask
- Utilization of the network. For a leaf network, Net Map reports the percentage of used IP addresses, except the broadcast and network addresses. For a network container, Net Map reports the percentage of the IP address space that has been allocated to either network containers or leaf networks.
- The first and last IP address of the network
- The total number of IP addresses in the network

When you mouse over a block of multiple networks, Net Map displays the following information:

- The start and end IP address of that block of networks
- The total number of IP addresses in that block of networks
- The number of networks in that block

Zooming In and Out

Use the zoom function to enlarge and reduce your view of a selected area. You can zoom in on any area in your network. You can zoom in on an area until it displays 128 addresses per row, for a total of 1024 addresses for the map. When you reach the last possible zoom level, the Zoom In icon in the Net Map task bar and the menu item are disabled.

After you zoom in on an area, you can click the Zoom Controller icon to track where you zoomed in. The Zoom Controller lists all the areas that you zoomed in and updates its list dynamically. You can click an item on the list to view that area again. Click the Zoom Controller again to close it.

To select an area and zoom in:

1. Right-click and select **Zoom In**, or click the Zoom In icon in the Net Map task bar.
The pointer changes to the zoom in selector.
2. Select a starting point and drag to the end point. The starting point can be anywhere in the map. It does not have to be at the beginning of a network.
Net Map displays a magnified view of the selected area after you release the mouse button. As you mouse over the zoomed in area, Net Map displays IP information about it.
3. You can do the following:
 - Select an area and zoom in again.
 - Add a network. If you zoom in on an area and click Add without selecting an open area first, Net Map selects the area where it can create the biggest possible network in that magnified area.
 - Select a network and perform any of the following operations:
 - Split the network.
 - Join it to another network.
 - Resize the network.
 - Edit its properties.
 - Open it to display its network or IP map.
 - Right-click and select **Zoom Out**, or click the Zoom Out icon in the Net Map task bar. Each time you click **Zoom Out**, Net Map zooms out one level and the Zoom Controller is updated accordingly.

Net Map Tasks

From Net Map, you can create IPv4 networks, and evaluate and manage your network resources according to the needs of your organization. You can do the following:

- Zoom in on specific areas, as described in [Zooming In and Out](#) on page 468.
- Add a network, as described in [Adding a Network from Net Map](#).
- Select a network and view either its network or Net Map, as described in [Viewing Network Details](#) on page 470.
- Select a network and edit its properties, as described in [Modifying IPv4 and IPv6 Network Containers and Networks](#) on page 465.
- Split a network, as described in [Splitting IPv4 Networks into Subnets](#) on page 472.
- Join networks, as described in [Joining IPv4 Networks](#) on page 473.
- Resize a network, as described in [Resizing IPv4 Networks](#) on page 472.
- (Applies only with Network Insight) Execute **VM Discovery** on the selected network, as described in [Performing VM \(Virtual Machine\) Discovery](#) on page 543.
- (Applies only with Network Insight) View **Discovery Status** for the selected network, as described in [Viewing Discovery Status](#) on page 538.
- (Applies only with Network Insight) Execute **Discovery Diagnostics** on the selected network, as described in [Using Discovery Diagnostics](#) on page 539.
- (Applies only with Network Insight) Direct NIOS to discover devices on the selected network (**Discover Now**). The network must have discovery enabled before this button will be active. See the topics under [About Network Insight](#) on page 519 for more information about requirements and discovery features for Network Insight.

Note: If the **Discover Now** button and other associated discovery elements are disabled on the Toolbar, it indicates that discovery is not enabled for the parent network of the selected network or IP, or the network is not associated with a Discovery appliance.

- Delete one or multiple networks, as described in [Discovering Networks \(Under Network Insight only\)](#) on page 473.
- **Clear All Unmanaged Data** or **Clear All Discovered Data**, as described in the section [Clearing Discovered Data](#) on page 515.
- Switch to the List view of the network. For information, see [Network List](#) on page 470.
 - When you select one or more networks in Net Map and then switch to the List view, the list displays the page with the first selected network.
 - If you select one or more networks in the List view and then switch to the Net Map view, the first network is also selected in Net Map. Although, if you select a network in the List view that is part of a Multiple Networks block in Net Map, it is not selected when you switch to the Net Map view.

Adding a Network from Net Map

When you create networks from Net Map, you can evaluate your network infrastructure and add networks accordingly. You can view the address space to which you are adding a network, so you can determine how much space is available and which IP addresses are not in use. When you mouse over an open area, Net Map displays useful information, such as the largest possible network that fits in that area and the total number of IP addresses. In addition, you can create networks without having to calculate anything. When you add a network, Net Map displays a netmask slider so you can determine the appropriate netmask for the size of the network that you need. As you move the slider, it displays network information, including the total number of addresses. After you select the netmask, you can even move the new network around the open area to select another valid start address.

To add a network from the Net Map panel:

1. Do one of the following:
 - Click the Add icon.

Net Map displays the netmask slider and outlines the open area that can accommodate the largest network.

- Select an open area, and then click the Add icon.
Net Map displays the netmask slider and outlines the largest network that you can create in the open area that you selected.
- 2. Move the slider to the desired netmask. You can move the slider to the netmask of the largest network that can be created in the open area.
As you move the slider, Net Map displays the netmask and its corresponding number of IP addresses. The outline in the network map also adjusts as you move the slider. When you mouse over the outline, it displays the start and end address of the network.
- 3. After you set the slider to the desired netmask, you can drag the new network block around the open area to select a new valid starting address. You cannot move the block to a starting address that is invalid.
- 4. Click **Launch Wizard** to create the network.
The *Add Network* wizard displays the selected network address and netmask.
- 5. You can add comments, automatically create reverse mapping zones, and edit the extensible attributes. (For information, see [Adding IPv4 Networks](#) on page 845.) You cannot change the network address and netmask.
- 6. Save the configuration and click **Restart** if it appears at the top of the screen. Grid Manager updates Net Map with the newly created network.

Viewing Network Details

From the Net Map panel, you can focus on a specific network or area and view additional information about it. If you have a network hierarchy of networks within network containers, you can drill down to individual leaf networks and view their IP address usage.

1. Select a network or area.
2. Click the Open icon.
 - If you selected a network container, Grid Manager displays it in the Net Map panel. You can drill down further by selecting a network or open area and clicking the Open icon again.
 - If you selected a block of multiple networks, Grid Manager displays the individual networks in the Net Map panel. You can then select a network or open area for viewing.
 - If you selected a leaf network, Grid Manager displays it in the IP Map panel.
 - If you selected an open area, Grid Manager displays an enlarged view of that area in the Net Map panel. This is useful when you are creating small networks in an open area.

Network List

The Network list panel is an alternative view of an IPv4 network hierarchy. For a given network, the panel shows all the networks of a selected network view in table format. A network list displays only the first-level subnets. It does not show further descendant or child subnets. You can open a subnet to view its child subnets. Subnets that contain child subnets are displayed as network containers. If the number of subnets in a network exceeds the maximum page size of the table, the network list displays the subnets on multiple pages. You can use the page navigation buttons at the bottom of the table to navigate through the pages of subnets.

The IPAM home panel displays the following:

- **Network:** The network address.
- **Comment:** The information you entered about the network.
- **RIR Organization:** This appears only if support for RIR updates is enabled. This displays the name of the RIR organization to which the network is assigned.
- **RIR Organization ID:** This appears only if support for RIR updates is enabled. This displays the ID of the RIR organization to which the network is assigned.
- **RIR Registration Status:** This appears only if support for RIR update is enabled. This field displays the RIR registration status. This can be **Registered** or **Not Registered**. **Registered** indicates that the network has a corresponding entry in the RIPE database.

- **Last Registration Updated:** Displays the timestamp when the last registration was updated. The displayed timestamp reflects the timestamp used on the Grid Master.
- **Status of Last Registration Update:** Displays the registration status and communication method of the last registration update. The status can be Pending, Sent, Succeeded, or Failed. Each time you send a registration update to create, modify, or delete a network container or network, the updated status will be displayed here. If you have selected not to send registration updates, the previous status is retained.
- **IPAM Utilization:** For a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network.

For a network container that contains subnets, this is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. For example, when you define a /16 network and then 64 /24 networks underneath it, the /16 network container is considered 25% utilized even when none of the IP addresses in the /24 networks is in use.

You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.

- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.

You can select the following columns for display:

- **Disabled:** Indicates whether the network is disabled.
- **Leaf Network:** Indicates whether the network is a leaf network or not. A leaf network is a network that does not contain other networks.
- **Discovery Enabled:** (*Applies only with Network Insight*) Indicates whether discovery is allowed on the network container or the network.
- **Managed:** (*Applies only with Network Insight*) Indicates whether the network is set to Managed status under NIOS.
- **First Discovered:** (*Applies only with Network Insight*) The date and timestamp of the first occasion that NIOS discovered the network.
- **Last Discovered:** (*Applies only with Network Insight*) The date and timestamp of the last occasion that NIOS performed discovery on the network.
- **Extensible attributes and RIR attributes:** You can select the extensible attributes such as Building, Country, Region, State, and VLAN for display. When you enable support for RIR registration updates, you can also select associated RIR attributes for display. For information about RIR attributes, see [Managing RIR Attributes](#) on page 444.

You can sort the list of networks in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#) on page 60.

You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.

Tip: If you select a network from the list and switch to the Net Map panel, the network is also selected in the network map.

Filtering the Network List

You can filter the network list, so it displays only the networks you need. You can filter the list based on certain parameters, such as network addresses, comments and extensible attributes. When you expand the list of available fields you can use for the filter, note that the extensible attributes are those with a gray background.

Resizing IPv4 Networks

You can resize a network to increase or decrease the network size and the number of IP addresses in the network. When you resize a network to a smaller netmask, you increase the number of IP addresses within that network. You can change the size of an IPv4 network when the operation does not affect existing objects in the network. You can resize an existing network only if the resized network does not exceed the upper network limit or create orphan objects, such as hosts and DHCP ranges. When a network has a parent network or subnets, the upper limit of the network size is marked in red in the resize network slider, and you cannot resize beyond this limit. For example, if a network has a /16 parent network, you cannot resize the network to a network that is larger than /16.

Before you resize an RIR allocated network block, ensure that the network block has already been registered at the corresponding RIR. Otherwise, when you reassign addresses within this block, the registration updates may fail. For information about RIR registration updates, see [RIR Registration Updates](#) on page 437.

To resize a network:

1. From the Net Map or List panel, select a network, and then click **Resize** from the Toolbar.
2. In the *Resize Network* editor, do the following:
 - **Address:** Displays the network address. You cannot modify this field.
 - **Netmask:** Displays the netmask of the network as you resize the network. You cannot modify this field.
 - **Resize slider:** Use the resize network slider to specify the appropriate subnet masks for the subnets. When you move the slider, Grid Manager displays the number of subnets and IP addresses within that subnet.
 - **Automatically create reverse-mapping zone:** This is enabled only when you resize a /8, /16, or /24 network. Select this check box to have the appliance automatically create reverse-mapping zones for the subnet. The appliance automatically creates reverse-mapping zones only for /8, /16, and /24 netmasks.
3. Click **OK**.

Splitting IPv4 Networks into Subnets

You can create smaller subnets simultaneously within a network by splitting it. You do not have to configure each subnet individually. You can create smaller subnets with larger netmasks. A larger netmask defines more networks with a smaller number of IP addresses.

These subnets inherit the address properties of the parent network, such as member assignments. The exceptions are the default router and broadcast address configuration. The default router and broadcast address configuration for address ranges and fixed address are disabled by default after splitting a network. You can enable these properties for each subnet after splitting the parent network.

Note that you cannot split a network that is part of a shared network.

To split a network:

1. From the Net Map or List panel, select the check box of a network, and then click **Split** from the Toolbar.
2. In the *Split Network* editor, do the following:
 - **Address:** Displays the network address. You cannot modify this field.
 - **Netmask:** Displays the netmask of the network. You cannot modify this field.
 - **Subnetworks:** Displays the number of subnets and IP addresses for each subnet.
 - **Split network slider:** Use the split network slider to specify the appropriate subnet masks for each subnet. When you move the slider, Grid Manager displays the number of subnets and the IP address range within that subnet.
 - **Immediately Add:** Select one of the following options.
 - **Only networks with ranges and fixed addresses and unmanaged:** Adds only the networks that have DHCP ranges, fixed addresses, and unmanaged addresses.
 - **All possible networks:** Adds all networks that are within the selected netmasks. This is enabled only when you split the /8 networks to /9 or /16 networks.

Note that when you add a large number of networks, it could take a little longer for Grid Manager to display the networks.

- **Automatically create reverse-mapping zone:** Select this check box to have the appliance automatically create reverse-mapping zones for the subnets.

3. Click **OK**.

Joining IPv4 Networks

Joining multiple networks into a larger network is the opposite of splitting a network. You can select a network and expand it into a larger network with a smaller netmask. A smaller netmask defines fewer networks while accommodating a larger number of IP addresses. Joining or expanding a network allows you to consolidate all of the adjacent networks into the expanded network. Adjacent networks are all networks falling under the netmask of the newly-expanded network. You can expand the selected network to a new size and add all other subnets into the new network. When you join networks, you need not define all small networks that cover the address spaces for a larger network.

Each of the adjacent networks join the expanded network and inherit the DHCP member configuration options of the selected network. The expanded network does not inherit the default router and broadcast address configurations of the adjacent networks. Those configurations are disabled by default.

Note: The member assignment for the expanded network combines all member assignments of the joining networks.

Note that the join and resize features work identically only when you have a single network. If the resize feature is disabled and if you have a single network object with additional new networks, then you must use the join feature to combine all networks.

To join or expand a network:

1. From the Net Map or List panel, select a network, and then click **Join** from the Toolbar.
2. In the *Join Network* editor, do the following:
 - **Address:** Displays the network address. You cannot modify this field.
 - **Netmask:** Displays the netmask of the network as you expand the network.
 - **Join Network slider:** Use the join network slider to specify the available subnet masks for the newly expanded network. Select a smaller netmask value, based on your requirements of the newly-expanded network. When you move the slider, a dialog box displays the total number of IP addresses and the IP address range of a selected subnet mask.
 - **Automatically create reverse-mapping zone:** Select this check box to configure the expanded network to support reverse-mapping zones .
3. Click **OK**.

Discovering Networks (Under Network Insight only)

Note: If the **Discover Now** button and other associated discovery elements are disabled on the Toolbar, it indicates that discovery is not enabled for the parent network of the selected network or IP, or that a Discovery appliance (known as a Probe) is not associated with the network that you wish to discover.

To discover IPv4 or IPv6 networks:

1. From the Net Map or List panel, select a network, and then click **Discover Now** from the Toolbar.
 NIOS asks you to confirm that you wish to launch discovery on the selected network.

In the Net Map panel, you can click on IP addresses in the network being discovered. As new data becomes available, NIOS updates the Discovered Data section of the panel with any information found on the device associated with the selected IP.

For more information about requirements and discovery features, see the topics under [About Network Insight](#) on page 519.

Deleting Networks

From the IPAM tab, you can delete multiple IPv4 and IPv6 networks. When you delete a network, all of its data, including all of its DHCP records, subnets, and records in its subnets, is deleted from the database and goes to the Recycle Bin, if enabled. Because of the potentially large loss of data that can occur when you delete a network, Grid Manager requires a confirmation to move the data to the Recycle Bin.

To delete IPv4 or IPv6 networks:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* check box. You can select multiple check boxes for multiple networks.
2. Select **Delete** or **Schedule Delete** from the Delete drop-down menu.
3. To delete the network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule a deletion, see [About Extensible Attributes](#) on page 322.

The appliance puts the deleted network in the Recycle Bin, if enabled.

VIEWING AND MANAGING IPv4 ADDRESSES

You can view and manage IPv4 address data in the IP Map and IP List panels. Grid Manager displays the IP Map and List panels for a specific network after you navigate through the network hierarchy, or when the selected network does not have subnets under it.

IP Map

The IPv4 Map panel provides a graphical representation of all IPv4 addresses in a given subnet. IP Map displays cells that represent IPv4 addresses. Each cell in the map represents an IPv4 address, and its color indicates its status as described in the legend section. You can run a network discovery on the selected network, and the status of each IP address is updated accordingly. For information, see [Chapter 13, Network Discovery](#), on page 493.

Each IP Map panel can accommodate up to 256 cells with each cell representing an IP address. If a given network has more than 256 addresses, additional IP addresses are displayed by paging to the next page. You can use the page navigation buttons to page through the IP addresses. To go to a specific IP address, you can enter the IP address in the **Go to** field or click a specific cell in IP Map.

IP Map has a basic and an advanced view. You can toggle between these views by clicking **Toggle Basic View** or **Toggle Advanced View**.

In the basic view, the IP Map panel displays the following IP address status:

- **Unused:** An IP address that has not been detected and is not associated with any network device or active host on the network.
- **Conflict:** An IP address that has either a MAC address conflict or a DHCP lease conflict detected through a network discovery.
- **Used:** An IP address that is associated with an active host on the network. It can be a resource record, fixed address, reservation, DHCP lease, or host record.
- **Pending:** An IP address that is associated with a scheduled task or approval workflow, and the associated operation has not been executed yet. This IP address is not considered when using the next available IP address function.
- **Selected IP Address:** The IP address that you selected.
- **DHCP Range:** The IP addresses within a DHCP range in the network.
- **Reserved Range:** A range of IP addresses that are reserved for statically configured hosts. They are not served as dynamic addresses. You can allocate the next available IP from the reserved range when you create a static host.

In the advanced view, the IP Map panel displays additional status as follows:

- **Unmanaged:** An IP address that has a discovered host, is not previously known to the appliance, and does not have an A record, PTR record, fixed address, host address, lease, or is not within a DHCP range. You can change an unmanaged address to a host, DHCP fixed address, A record, or PTR record. You can also clear an unmanaged address. All existing administrator permissions apply to the unmanaged addresses.
- **Fixed Address/Reservation:** A host that is either a fixed address or reservation.
- **DNS Object:** An object that is configured for DNS usage.
- **Host Not in DNS/DHCP:** An IP address that is associated with a host record, but is not configured for DHCP or DNS services.
- **Active Lease:** An IP Address that has an active DHCP lease.
- **DHCP Exclusion Range:** A range of IP addresses within a DHCP range. The appliance cannot assign addresses in the exclusion range to a client. You can use these addresses as static IP addresses. This prevents address conflicts between statically configured devices and dynamically configured devices.

Under the IP map, Grid Manager displays the following information for the IP address that you have selected in the map:

- **Type:** The object type that is associated with the IP address. For example, this can be **Lease, IPv4 DHCP Range or Fixed Address**.
- **Comment:** Additional information about the IP address.
- **Lease State:** The lease state of the IP address. This can be one of the following: **Free, Backup, Active, Expired, Released, Abandoned, Reset, BootP, Static, Offered, or Declined**.
- **Name:** The name of the object type associated with the IP address. This field displays the name of the object type in the native character set if a host record contains IDNs. If a host record contains IDNs in punycode, this field displays the name in the punycode representation. For example, if the IP address belongs to a host record, this field displays the hostname. For IDNs, this field displays the name in the native character set. If punycode is used, then the appliance displays name in punycode.
- **MAC Address:** The discovered MAC address of the host. This is the unique identifier of a network device. The discovery acquires the MAC address for hosts that are located on the same network as the Grid member that is running the discovery. This can also be the MAC address of a virtual entity on a specified vSphere server.
- **DHCP Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.

You can do the following in the IP Map panel:

- Click **Go to DHCP View** to view DHCP properties of a selected network.
- Select an address range by clicking once on a start address and then use SHIFT+click on the end address. Click **Add -> Range** from the Toolbar to add the selected range as an IPv4 or IPv6 DHCP range or reserved range.
- Click the Resolve Conflict icon to resolve IP address conflicts. For information, see [Resolving Conflicting Addresses](#) on page 513.
- Click the Ping icon to ping a selected IP address. For information, see [Pinging IP Addresses](#) on page 491.
- Click the Reclaim icon to reclaim an IP address. For information, see [Reclaiming Objects Associated with IPv4 and IPv6 Addresses](#) on page 491.
- Click the Clear icon to clear an active lease. For information, see [Clearing Active DHCP Leases](#) on page 491.

You can also select an IP address from the IP Map panel and view the following information:

- General information, as described in [IP Address Header Panel](#) on page 477.
- Data retrieved through a network discovery or integrated from a PortIQ appliance and Trinetic Network Automation. For information, see [Viewing Discovered Data](#) on page 510.
- The records associated with the IP address, as described in [Related Objects](#) on page 486.
- The audit history, as described in [Audit History](#) on page 486.
- Detailed lease information, as described in [Viewing Detailed Lease Information](#) on page 948.
- Click **DHCP View** to view DHCP properties of the selected network. For information, see [Modifying IPv4 Networks](#) on page 851.

IP Address List

The IP address List panel displays all IPv4 addresses of a selected subnet in table format. The list provides information about the IP addresses in a hierarchy view. You can use this list to view detailed information about each IP address and its related objects in a selected network. This list provides information such as address status, object type, and usage.

You can configure filter criteria to display only IP addresses that you want to see in the table. For example, you can enter “MAC Address begins with 00” as the filter criteria to view only IP addresses that have associated MAC addresses that begin with 00. You can also enter a specific IP address in the **Go to** field to view information about the address.

By default, Grid Manager displays the following information for the IP addresses, except for **Disabled**.

- **IP Address:** The IP address of the corresponding record. The appliance highlights disabled DHCP objects in gray. A DHCP object can be an DHCP address range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **Name:** The name of the object type associated with the IP address. For example, if the IP address belongs to a host record, this field displays the hostname.
- **MAC Address:** The discovered MAC address of the host. This is the unique identifier of a network device. The discovery acquires the MAC address for hosts that are located on the same network as the Grid member that is running the discovery. This can also be the MAC address of a virtual entity on a specified vSphere server.
- **Status:** The current status of the corresponding record, such as **Used**, **Unmanaged**, **Conflict**, or **Unused**.

Note: The appliance displays **Used** for records that are part of a scheduled task.

- **Type:** The object type that is associated with the IP address. For example, this can be **Lease**, **IPv4 DHCP Range**, **Fixed Address**, **Host**, or **Pending** (if the IP address is associated with a scheduled task or approval workflow). Note that all pending IP addresses are considered used and highlighted in grey.
- **Usage:** Indicates whether the IP address is configured for DNS or DHCP.
- **Lease State:** The lease state of the IP address. This can be one of the following: **Free**, **Backup**, **Active**, **Expired**, **Released**, **Abandoned**, **Reset**, **BootP**, **Static**, **Offered**, or **Declined**.
- **User Name:** The name of the user who received the lease for the IP address.
- **Comment:** Additional information about the IP address.
- **First Discovered:** The timestamp when the IP address was initially discovered. This data is read-only.
- **Last Discovered:** The timestamp when the IP address was last discovered. This data is read-only.
- **OS:** The operating system of the discovered host. The OS value can be one of the following:
 - **Microsoft** for all discovered hosts that have a non-null value in the MAC addresses using the NetBIOS discovery method.
 - A value that a TCP discovery returns.
 - The OS of a virtual entity on a vSphere server.

Note that this field sometimes displays the percentage of certainty about the discovered OS.

- **NetBIOS Name:** The returned NetBIOS name from the last discovery.
- **Discovered Name:** The name of the discovered IP address, if any was previously assigned by an administrator.
- **Discoverer:** The identity of the appliance that discovered the IP address.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.
- **Disabled (hidden):** Indicates whether the DHCP or DNS record is disabled.

Note: For an IP address that falls within a DHCP range, Grid Manager displays extensible attribute values for the DHCP range and fixed address or host record. When you view the same IP address in the **DHCP** tab however, Grid Manager displays only the extensible attribute values associated with the fixed address or host record, but not the DHCP range. For example, when you define extensible attribute **State** with the value **California** for DHCP range 1.0.0.1 – 1.0.0.5, and then define extensible attribute **State** with the value of **Alaska** for fixed address 1.0.0.3, Grid Manager displays both **California** and **Alaska** in the **State** field for IP address 1.0.0.3 in the IP Address List view. However, when you view 1.0.0.3 from the **DHCP** tab, the **State** field displays **Alaska** only.

You can display all available extensible attributes. You can also sort the list of IP addresses in ascending or descending order by **IP Address** only. If you enabled the IP Discovery feature, you can configure the IP List panel to display discovered data and fields imported from PortIQ or Trinzie Network Automation appliances. For information about the PortIQ data, see [Integrating Data from PortIQ Appliances](#) on page 508. For information about integrating discovered data from Trinzie Network Automation, see [Integrating Discovered Data From Trinzie Network Automation](#) on page 509.

You can select an IP address from the List panel and view the following information about it:

- General information, as described in [IP Address Header Panel](#) on page 477.
- Data retrieved through a network discovery or integrated from a PortIQ appliance, as described in [Viewing Discovered Data](#) on page 510.
- The records associated with the IP address, as described in [Related Objects](#) on page 486.
- Audit history, as described in [Audit History](#) on page 486.
- Detailed lease information, as described in [Viewing Detailed Lease Information](#) on page 948.

You can also do the following from the IP List panel:

- Click **Go to DHCP View** to view DHCP properties of a selected network. For information, see [Modifying IPv4 Networks](#) on page 851.
- Click the Ping icon to ping a selected IP address. For information, see [Pinging IP Addresses](#) on page 491.

Filtering the IP Address List

You can filter the IP address list, so it displays only the IP addressees you need. You can filter the list based on any combination of extensible attributes and the parameters displayed in the IP address list, such as usage and type. When you expand the list of available fields you can add to the filter, note that the extensible attributes are those with the gray background.

IP Address Header Panel

When you select an IP address from the IP Map or List panel, Grid Manager displays information about the highest priority object associated with the IP address. Depending on the object type, Grid Manager displays all or some of the following information. For example, if the highest priority object is a fixed address, Grid Manager displays only the object type, MAC address, lease state, and comment of the object.

- **Type:** The object or record type, such as A record, PTR record, or host record.
- **Name:** The name of the object. For example, if the IP address belongs to a host record, this field displays the hostname. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP address range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **MAC:** The MAC address of the network device associated with the IP address.
- **Lease State:** The current status of the DHCP lease.
- **Comment:** Comments about the IP address.

Discovered Data

The **Discovered Data** tab displays discovered data through a network discovery or integrated from PortIQ and Trinziq Network Automation appliances. For information about viewing discovered data, see [Viewing Discovered Data](#) on page 510.

Related Objects

The Related Objects tab displays the following information about the records associated with the IP address:

- **Name:** The name of the object. For example, if the IP address belongs to a host record, this field displays the hostname. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **Type:** The object type, such as DHCP lease, host, A record, and bulk host.
- **Comment:** Information about the object.

You can also select the following for display:

- **DNS view:** The DNS view to which the object belongs.

You can do the following in this tab:

- Add a resource record. You can select the following from the drop-down list:
 - Host Record—For information, see [Adding Host Records](#) on page 462.
 - Range—For information, see [Adding IPv4 Address Ranges](#) on page 854.
 - Fixed Address—For information, see [Adding IPv4 Fixed Addresses](#) on page 858.
 - Reservation—For information, see [Adding IPv4 Reservations](#) on page 861.
 - A Record—For information, see [Adding A Records](#) on page 661.
 - PTR Record—For information, see [Adding PTR Records](#) on page 664.
- Edit the properties of the selected object. Depending on the type of object, Grid Manager displays the corresponding editor for the object. For example, if the selected object is a fixed address, Grid Manager displays the fixed address editor. When you select a lease object, Grid Manager displays the lease viewer.
- You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62. Delete a selected object or multiple objects.
- When you select a lease object and click the Show Details icon, you can view the lease start and end dates.
- Depending on the object type, you can convert a selected object to one of the following:
 - **Reservation**
 - **Host**
 - **Fixed Address**
- View detailed lease information about the IP address, as described in [Viewing Detailed Lease Information](#) on page 948.
- Print and export the information in the Related Objects table.

Audit History

By default, the Audit History tab displays the following information about the last five actions performed on the selected IP:

- **Timestamp:** The day, date, and time of the operation.
- **Action:** The type of operation that was performed by the administrator.
- **Object Type:** The object type of the entry.
- **Object Name:** The name of the object.
- **Admin Name:** The name of the administrator who performed the operation.
- **Message:** The description of the administrative activity.

Note: If you change the IP address of an existing record to a new one in the **IP MAP** tab when the **Grid Audit Logging** is set to *Brief*, then NIOS will not display modification or transition details about the new IP address in this tab. You can only view subsequent modifications and deletions to the new IP address. However, you can view the audit log history and transition details of the old IP address, but you cannot view the initial transition from an old IP address to the new IP address.

Managing IPv4 Addresses

You can do the following from the IP Map and List panels:

- Add IP addresses to existing hosts. For information, see [Adding IP Addresses to Existing Host Records](#) on page 479.
- Clear unmanaged IP addresses. For information, see [Clearing Unmanaged Data](#) on page 479.
- Convert objects to other object types. For information, see [Converting Objects Associated with IP Addresses](#) on page 487.
- Reclaim IP addresses. For information, see [Reclaiming Objects Associated with IPv4 and IPv6 Addresses](#) on page 491.
- Ping IP addresses. For information, see [Pinging IP Addresses](#) on page 491.
- Configure and run a network discovery. For information, see [Network Discovery](#) on page 493.
- Resolve discovery conflicts. For information, see [Resolving Conflicting Addresses](#) on page 513.
- Clear discovered data. For information, see [Clearing Discovered Data](#) on page 515.

Adding IP Addresses to Existing Host Records

You can add unused and unmanaged addresses, including all their information, to existing host records. When you add an unmanaged address to a host record, the appliance adds the discovered data to the host record. You can select the desired host to which you want to add the unmanaged address.

To add an unmanaged IP address to an existing host record:

1. From the IP Map or List panel, select an IP address, and then click **Add -> Add to Existing Host** from the Toolbar.
2. In the *Select Host* dialog box, do the following:
 - In the table, select the host to which you want to add the selected IP address. You can also use the filters or the **Go To** field to narrow down the host list. For information, see [Using Filters](#) on page 67 and [Using the Go To Function](#) on page 71.
 - Click the Select icon.
 Grid Manager displays the *Host Record* editor.
3. In the *Host Record* editor, update the host properties as described in [Modifying Host Records](#) on page 463.
4. Save the configuration and click **Restart** if it appears at the top of the screen. To close the editor without saving the changes, click the **Close** icon.

Clearing Unmanaged Data

You can clear the status of unmanaged data at the network and IP address levels. When you clear an unmanaged address, the status of the IP address changes to **Unused**. An unmanaged address is an IP address with a discovered host, is not previously known to the appliance, and does not have an A record, PTR record, fixed address, host address, lease, or is not within a DHCP range. You can change an unmanaged address to a host, a DHCP fixed address, an A record, or a PTR record. You can also clear the unmanaged data associated with the address.

To clear unmanaged data:

1. From the IP Map or List panel, select the IP address for which you want to clear unmanaged data, and then click **Clear -> Clear Unmanaged Data** from the Toolbar. You can select multiple IP addresses.
2. In the *Clear Unmanaged data* dialog box, click **Yes**.

MANAGING IPV6 NETWORKS

In Grid Manager, you can use the IPv6 Net Map (network map) and List panels to manage your IPv6 network infrastructure. After you select a network container from the IPAM tab, Grid Manager displays it in the Net Map panel, by default. The Net Map panel provides a graphical view of your networks and has a number of features that simplify network management. The List panel displays the networks in table format.

You can always switch your view of a network container between the Net Map and List panels. Grid Manager keeps track of which panel you last used. When you select a network container, Grid Manager displays it in the Net Map or List panel, depending on which one you last used. For information about each panel, see [IPv4 Network Map](#) on page 467 and [Network List](#) on page 470.

You can use Grid Manager to manage IPv6 networks and their AAAA, PTR and host resource records. You can configure IPv6 networks and track IP address usage in those networks. You can also split and join IPv6 networks, when necessary.

IPv6 Network Map

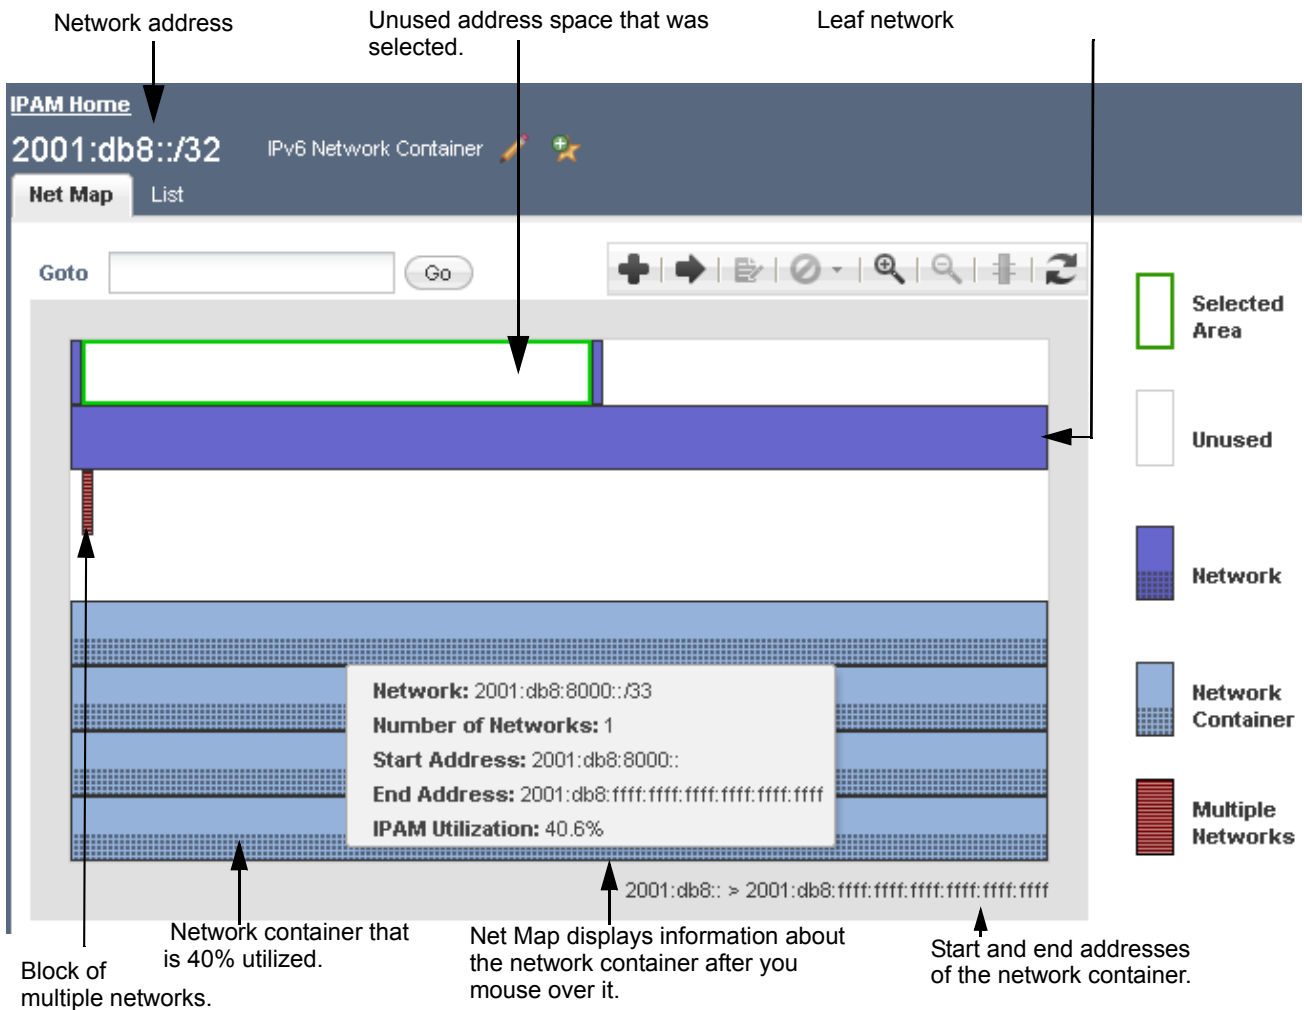
After you select an IPv6 network container from the IPAM tab, Grid Manager displays it in the IPv6 Net Map (network map) panel, by default. Just like the IPv4 Net Map, the IPv6 Net Map provides a high-level view of the network address space. You can use Net Map to design and plan your network infrastructure, and to configure and manage individual networks.

The Net Map panel presents a complete view of the network space, including the different types of networks that are in it and its unused address space. IP addresses that belong to a network are blocked off. Each color-coded block represents a network container, a leaf network, or a block of networks that are too small to be displayed individually in the map. For example, in a /64 or /96 network, networks smaller than /76 or /108 respectively and that are beside each other are represented as a multiple network block. In addition, the fill pattern of the blocks indicates their utilization. Therefore, you can quickly evaluate how many networks are in a network container, their relative sizes, utilization, and how much space you have left.

As you mouse over areas of the map, it displays IP information about the area. Net Map also has a zoom feature that allows you to enlarge or reduce your view of a particular area.

[Figure 12.7](#) displays the network map of a 1111::/16 network, which is a network container that has network containers and leaf networks.

Figure 12.7 IPv6 Network Map



Displaying Network Information

As shown in [Figure 12.7](#), as you mouse over the map, Net Map displays IP information about the area. When you mouse over an unused area, Net Map displays the following information:

- The start and end IP address
- The largest possible network
- The number of /64 networks that can fit in that space

When you mouse over a network container, Net Map displays the following information:

- Network address and netmask
- The first and last IP address of the network
- The number of networks in that block
- IPAM utilization

When you mouse over a network, Net Map displays the following information:

- Network address and netmask
- The first and last IP address of the network

When you mouse over a block of multiple networks, Net Map displays the following information:

- The start and end IP address of that block of networks
- The number of networks in that block

Zooming In and Out

Use the zoom function to enlarge and reduce your view of a selected area. You can zoom in on any area in your network. You can zoom in on an area until it displays 128 addresses per row, for a total of 1024 addresses for the map. When you reach the last possible zoom level, the Zoom In icon in the Net Map task bar and the menu item are disabled.

After you zoom in on an area, you can click the Zoom Controller icon to track where you zoomed in. The Zoom Controller lists all the areas that you zoomed in and updates its list dynamically. You can click an item on the list to view that area again. Click the Zoom Controller again to close it.

To select an area and zoom in:

1. Right-click and select **Zoom In**, or click the Zoom In icon in the Net Map task bar.
The pointer changes to the zoom in selector.
2. Select a starting point and drag to the end point. The starting point can be anywhere in the map. It does not have to be at the beginning of a network.
Net Map displays a magnified view of the selected area after you release the mouse button. As you mouse over the zoomed in area, Net Map displays IP information about it.
3. You can do the following:
 - Select an area and zoom in again.
 - Add a network. If you zoom in on an area and click Add without selecting an open area first, Net Map selects the area where it can create the biggest possible network in that magnified area.
 - Select a network and perform any of the following operations:
 - Edit its properties.
 - Open it to display its IP List.
 - Delete it immediately, or schedule its deletion.
 - Right-click and select **Zoom Out**, or click the Zoom Out icon in the Net Map task bar. Each time you click **Zoom Out**, Net Map zooms out one level and the Zoom Controller is updated accordingly.

Net Map Tasks

From Net Map, you can create IPv6 networks, and evaluate and manage your network resources according to the needs of your organization. You can do the following:

- Zoom in on specific areas, as described in [Zooming In and Out](#) on page 468.
- Use the **Go to** function to find a network in the current zoom level of Net Map.
- Add a network, as described in [Adding a Network from Net Map](#).
- Select a network and view IP address list, as described in [Viewing IPv6 Data](#) on page 486.
- Select a network and edit its properties, as described in [Modifying IPv4 and IPv6 Network Containers and Networks](#) on page 465.
- Split a network, as described in [Splitting IPv6 Networks into Subnets](#) on page 485.
- Join networks, as described in [Joining IPv6 Networks](#) on page 485.
- Delete one or multiple networks, as described in [Discovering Networks \(Under Network Insight only\)](#) on page 473.
- Switch to the List view of the network. For information, see [IPv6 Network List](#) on page 484.
 - When you select one or more networks in Net Map and then switch to the List view, the list displays the page with the first selected network.
 - If you select one or more networks in the List view and then switch to the Net Map view, the first network is also selected in Net Map. Although, if you select a network in the List view that is part of a Multiple Networks block in Net Map, it is not selected when you switch to the Net Map view.

Adding a Network from Net Map

When you create networks from Net Map, you can view the address space to which you are adding a network, so you can determine how much space is available and which IP addresses are not in use. When you mouse over an open area, Net Map displays useful information, such as the largest possible network that fits in that area. In addition, you can create networks without having to calculate anything. When you add a network, Net Map displays a netmask slider so you can determine the appropriate netmask for the size of the network that you need. As you move the slider, it displays network information, including the total number of addresses. After you select the netmask, you can even move the new network around the open area to select another valid start address.

To add a network from the Net Map panel:

1. Do one of the following:
 - Click the Add icon.
Net Map displays the netmask slider and outlines the open area that can accommodate the largest network.
 - Select an open area, and then click the Add icon.
Net Map displays the netmask slider and outlines the largest network that you can create in the open area that you selected.
2. Move the slider to the desired netmask. You can move the slider to the netmask of the largest network that can be created in the open area. You can also move the slider to the smallest network that can be placed in the current zoom level of Net Map.
As you move the slider, Net Map displays the netmask. The outline in the network map also adjusts as you move the slider. When you mouse over the outline, it displays the start and end address of the network.
3. After you set the slider to the desired netmask, you can drag the new network block around the open area to select a new valid starting address. You cannot move the block to a starting address that is invalid.
4. Click **Launch Wizard** to create the network.
The *Add Network* wizard displays the selected network address and netmask.
5. You can add comments, automatically create reverse mapping zones, and edit the extensible attributes. (For information, see [Adding IPv6 Networks](#) on page 871.)
6. Save the configuration and click **Restart** if it appears at the top of the screen.
Grid Manager updates Net Map with the newly created network.

Viewing Network Details

From Net Map, you can focus on a specific network or area and view additional information about it. If you have a network hierarchy of networks within network containers, you can drill down to individual leaf networks and view their IP address usage.

1. Select a network or area.
2. Click the Open icon.
 - If you selected a network container, Grid Manager displays it in the Net Map panel. You can drill down further by selecting a network or open area and clicking the Open icon again.
 - If you selected a block of multiple networks, Grid Manager displays the individual networks in the Net Map panel. You can then select a network or open area for viewing.
 - If you selected a leaf network, Grid Manager displays it in the Network List panel.
 - If you selected an open area, Grid Manager displays an enlarged view of that area in the Net Map panel. This is useful when you are creating small networks in an open area.

IPv6 Network List

The Network list panel is an alternative view of an IPv6 network hierarchy. For a given network, the panel shows all the networks of a selected network view in table format. A network list displays only the first-level subnets. It does not show further descendant or child subnets. You can open a subnet to view its child subnets. Subnets that contain child subnets are displayed as network containers. If the number of subnets in a network exceeds the maximum page size of the table, the network list displays the subnets on multiple pages. You can use the page navigation buttons at the bottom of the table to navigate through the pages of subnets.

The IPAM home panel displays the following:

- **Network:** The network address.
- **Comment:** Information you entered about the network.
- **IPAM Utilization:** For a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.

You can select the following columns for display:

- **Disabled:** Indicates whether the network is disabled.
- **Leaf Network:** Indicates whether or not the network is a leaf network.
- Other available extensible attributes

You can sort the list of subnets in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#) on page 60.

You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.

Tip: If you select a network from the list and switch to the Net Map panel, the network is also selected in the network map.

Filtering the Network List

You can filter the network list, so it displays only the networks you need. You can filter the list based on certain parameters, such as network addresses, comments and extensible attributes. When you expand the list of available fields you can use for the filter, note that the extensible attributes are those with a gray background.

Splitting IPv6 Networks into Subnets

You can create smaller subnets simultaneously within a network by splitting it. You do not have to configure each subnet individually. You can create smaller subnets with larger netmasks. A larger netmask defines a larger number of network addresses and a smaller number of IP addresses.

Note that you cannot split a network that is part of a shared network.

To split an IPv6 network:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* check box, and then click **Split** from the Toolbar.
2. In the *Split Network* editor, do the following:
 - **Address:** Displays the network address. You cannot modify this field.
 - **Netmask:** Specify the appropriate netmask for each subnet.
 - **IPv6 Prefix Collector Network:** If you split a network with prefix delegations that are not tied to specific addresses, specify the network in which all prefix delegations are assigned. If you leave this field blank, the server assigns all prefix delegations that are not tied to specific addresses to the first network.
 - **Immediately create:** Select one of the following
 - **Only networks with ranges and fixed addresses:** Adds only the networks that have DHCP ranges and fixed addresses.
 - **All possible networks:** Adds all networks that are within the selected netmasks. You can select this option only when you increase the CIDR by 8 bits.
 - **Automatically create reverse-mapping zone:** Select this check box to have the appliance automatically create reverse-mapping zones for the subnets. This function is enabled if the netmask of the network is a multiple of four, such as 4, 12 or 16.
3. Click **OK**.

Joining IPv6 Networks

Joining multiple networks into a larger network is the opposite of splitting a network. You can select a network and expand it into a larger network with a smaller netmask. A smaller netmask defines fewer networks while accommodating a larger number of IP addresses. Joining or expanding a network allows you to consolidate all of the adjacent networks into the expanded network. Adjacent networks are all networks that fall under the netmask of the newly-expanded network.

To join or expand a network:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* check box, and then click **Join** from the Toolbar.
2. In the *Join Network* editor, do the following:
 - **Address:** Displays the network address. You cannot modify this field.
 - **Netmask:** Enter the netmask of the expanded network.
 - **Automatically create reverse-mapping zone:** Select this check box to configure the expanded network to support reverse-mapping zones. The appliance automatically creates reverse-mapping zones only if the netmask is between /4 through /128, in increments of 4 (that is, /4, /8, /12, and so on until /128).
3. Click **OK**.

VIEWING IPV6 DATA

To the configured IP addresses in an IPv6 network:

by selecting an IPv6 leaf network from the Network List panel

- For a leaf network that is not in a network container, from the **Data Management** tab, select the **IPAM** tab, and then click the IPv6 network you want to view.
- For a leaf network that is in a network container, from the **Data Management** tab, select the **IPAM** tab -> *network_container*-> *network*.

Grid Manager lists the configured IPv6 addresses. You can export and print the list. It displays the following information about each IP address:

- **IP Address:** The name of the IPv6 DHCP object, which can be a DHCP range, fixed address, host configured for DHCP, or a roaming host with an allocated IP address.
- **Name:** The name of the record associated with the IP address.
- **DUID:** The DHCP Unique Identifier (DUID) of the device that was assigned the IP address.
- **Status:** The status of the IPv6 object, such as Used or Unused.
- **Type:** The object type associated with the IP address, such as **AAAA record**, **IPv6 Fixed Address**, or **Unmanaged**.
- **Usage:** Indicates whether the IPv6 address is configured for DNS or DHCP.
- **Exclude:** (*Applies only with Network Insight*) Denotes whether the IP is excluded from discovery.
- **Lease State:** (*Applies only with Network Insight*) The lease state of the record, such as Active.
- **First Discovered:** (*Applies only with Network Insight*) The date and timestamp of the first occasion that NIOS discovered the IP address.
- **Last Discovered:** (*Applies only with Network Insight*) The date and timestamp of the last occasion that NIOS discovered the IP address.
- **OS:** The operating system of the IP.
- **NetBIOS Name:** The returned NetBIOS name from the last discovery.
- **Device Type(s):** Shows the device type for the device associated with the IP address.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.
- **Comment:** Displays comments about the record.
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.

You can display all available extensible attributes. You can also sort the list of IP addresses in ascending or descending order by **IP Address** only.

You can drill down further and view the records associated with an IP address. To view the associated records of an IP address, select it and Grid Manager displays information about the IP address in the **Related Objects** and **Audit History** tabs.

Related Objects

Grid Manager displays the following information about the records associated with the IP address:

- **Name:** The record name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Type:** The object type. For example, AAAA Record, PTR Record, Host Record, IPv6 Fixed Address.
- **Comment:** Additional information that was entered in the record about the IP address.

Audit History

Grid Manager displays the following information about the last five actions performed on the selected IP:

- **Timestamp:** The day, date, and time of the operation.

- **Action:** The type of operation that was performed by the administrator.
- **Object Type:** The object type of the entry.
- **Admin Name:** The name of the administrator that performed the operation.
- **Message:** Description of the administrative activity.

Filtering the IP Address List

You can filter the IP address list, so it displays only the IP addressees you need. You can filter the list based on any combination of extensible attributes and the parameters displayed in the IP address list, such as usage and type. When you expand the list of available fields you can add to the filter, note that the extensible attributes are those with the gray background.

MANAGING IPV4 AND IPV6 ADDRESSES

Grid Manager uses IP addresses as the entry point to the data set containing Infoblox host, DNS, DHCP, and other information related to that address. You can view the data, modify it, assign extensible attributes to the objects associated with the address, and convert DHCP lease types, such as changing a currently active dynamic lease to a fixed address or host record.

You can view and manage IPv4 address data in the IP Map panel, and view and manage IPv4 and IPv6 data in the IP List panel. You can do the following for IPv4 and IPv6 data from the IP List panel:

- Convert objects to other object types. For information, see [Converting Objects Associated with IP Addresses](#) on page 487.
- Reclaim IP addresses. For information, see [Reclaiming Objects Associated with IPv4 and IPv6 Addresses](#) on page 491.
- Ping IP addresses. For information, see [Pinging IP Addresses](#) on page 491.
- Clear DHCP leases. For information, see [Clearing Active DHCP Leases](#) on page 491.

You can also print and export in CSV format the information displayed in any panel that supports these functions.

Converting Objects Associated with IP Addresses

The NIOS appliance provides a simple mechanism for converting unmanaged IP addresses to resource records, such as host records and A or AAAA records. You can also convert the active lease of a dynamically assigned IPv4 or IPv6 address to a fixed address or host, and convert an IPv4 lease to an IPv4 reservation. Using the conversion mechanism, you can keep the existing information of a network device during the conversion.

The appliance supports the following conversions for IPv4 objects:

- DHCP leases to fixed addresses, reservations, or host records
- Fixed addresses to reservations or host records
- Unmanaged addresses to host records, A records, PTR records, or fixed addresses
- A records to host records
- PTR records to host records

The appliance supports the following conversions for IPv6 objects:

- DHCP leases to fixed addresses or host records
- Fixed addresses to host records
- AAAA records to host records
- IPv6 PTR records to host records

Note: You cannot convert unmanaged IP addresses or leases served by Microsoft DHCP servers to host records.

Converting DHCP Leases

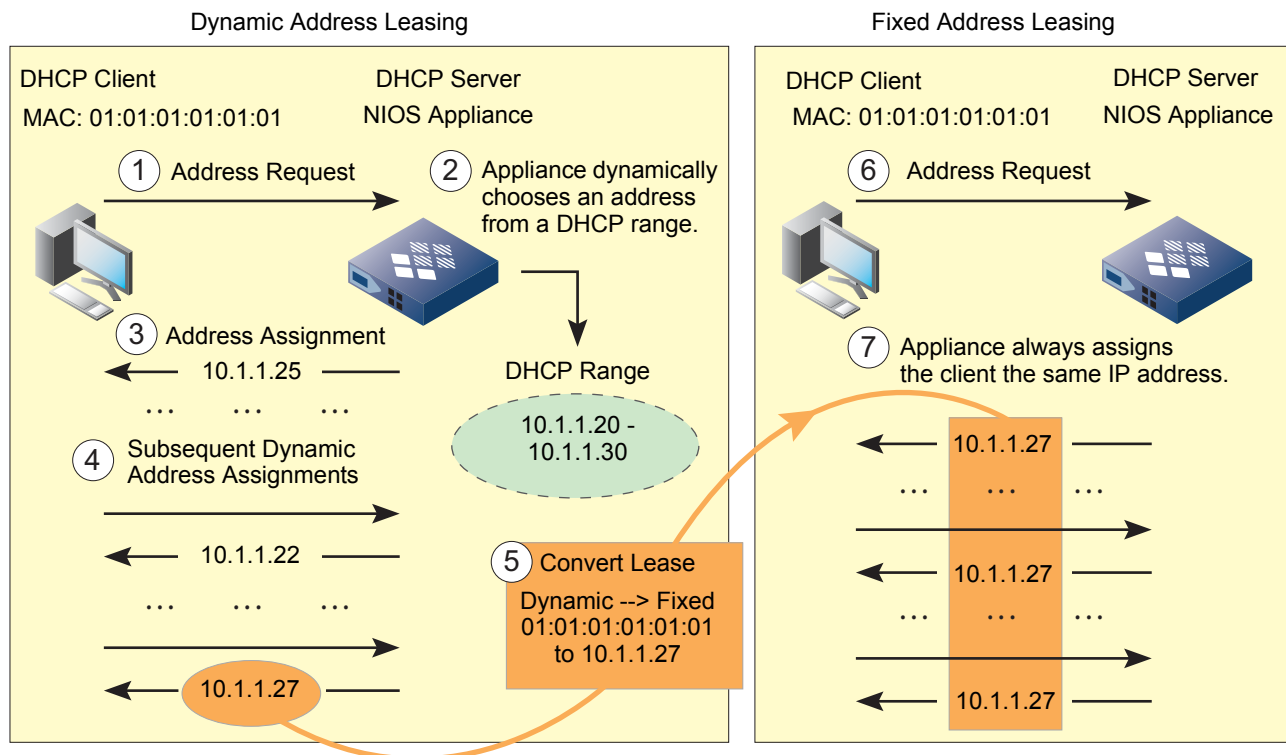
To create a fixed address, you bind an IPv4 address to a MAC address or an IPv6 address to a DUID. You can make that binding by converting an active dynamically leased address to a fixed address. The lease conversion transforms the temporary binding between the IPv4 address and MAC address or the IPv6 address and DUID in the dynamic lease to a persistent one. The lease must be active so that the NIOS appliance has an IPv4-to-MAC address or IPv6-to-DUID binding to convert into a fixed address.

The appliance uses the following rules when converting a DHCP lease:

- If an IPv4 DHCP lease is converted to a fixed address, the appliance copies the client identifier to the fixed address, based on information in the lease. If the appliance finds the client identifier in the lease information, the appliance includes it when it creates the host. If it finds the MAC address, the appliance includes it when it creates the host. If it finds both, the appliance includes only the MAC address (default) when it creates the host.
- If an IPv6 DHCP lease is converted to a fixed address, the appliance copies the DUID to the fixed address.
- If you try to convert an IPv4 DHCP lease or a fixed address with a client identifier, not a MAC address, to a host, the appliance displays an error message in the host editor. This ensures that you do not attempt this operation and lose the data.
- You cannot create two IPv4 fixed addresses with the same client identifier or MAC address in the same network. You cannot create two IPv6 fixed addresses with the same DUID in the same network.
- If the appliance receives a second IPv4 DHCP request with the same client identifier, it provides the same fixed IP address if the lease is still binding.

Figure 12.8 illustrates converting a dynamic IPv4 lease to a fixed lease.

Figure 12.8 Converting a Dynamic IPv4 Lease to a Fixed Lease



An advantage of converting an active dynamic lease is that you do not need to learn the MAC address or DUID of the device to which you want to assign an IP address and manually enter it in the fixed address configuration.

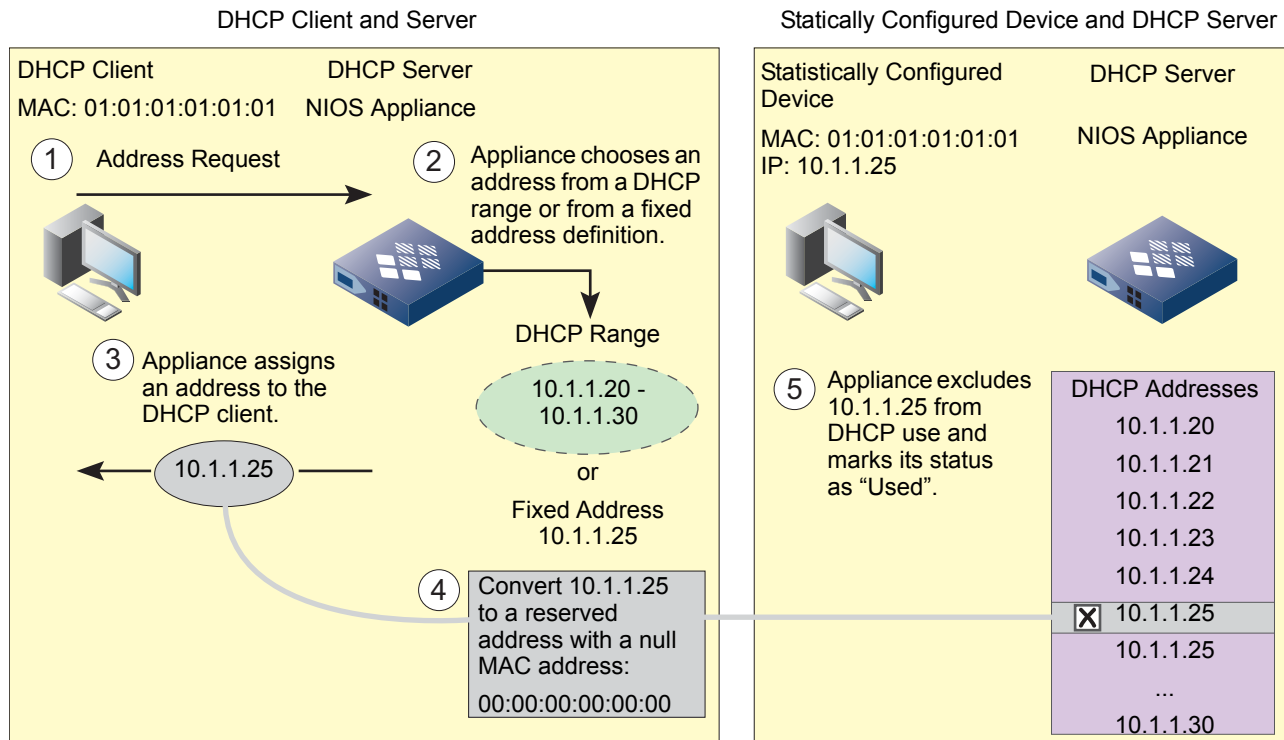
An IPv4 reservation is an address that you exclude from DHCP use because you intend to configure that address manually on a device, such as a firewall, router, or printer. You can also convert an IPv4 fixed address or a dynamic address with an active lease to a reservation.

When you convert an address in a DHCP range to a reservation, you reduce the total number of dynamically assignable addresses in that range by one. Correspondingly, this reduces the number of allocated addresses needed to exceed a high or low watermark threshold for that range.

Note: To return an IP address to its place in a DHCP range after converting it from an active dynamic lease to a fixed address, reservation, or Infoblox host, delete the fixed address, reservation, or host to which you previously converted the IP address. The IP address then becomes part of the DHCP range to which it first belonged.

You can convert IPv4 fixed addresses to reservations, as shown in [Figure 12.9](#).

Figure 12.9 Converting an IPv4 Dynamic Lease or Fixed Address to a Reservation



To convert an object:

1. From the IP Map, select an IPv4 address or from the IP List panel, select an IPv4 or IPv6 address.
2. In the Related Objects tab, select the check box of the object, and then click **Convert** from the Toolbar or navigation bar.
3. Select the object type to which you want to convert the object. Grid Manager displays the corresponding editor for the object type.
4. For all IPv4 conversions, Grid Manager populates the discovered information in the corresponding editor. Depending on the type of conversion, do one of the following:
 - For host record conversions, see [Modifying Host Records](#) on page 463.
 - For IPv4 reservation conversions, see [Modifying Reservations](#) on page 862.
 - For fixed address conversions, see [Modifying IPv4 Fixed Addresses](#) on page 859.
 - For A record conversions, see [Modifying A Records](#) on page 661.
 - For PTR record conversions, see [Modifying PTR Records](#) on page 665.

Note: When you select an object for conversion, Grid Manager displays only the available conversion types for the object. You must save the changes in the editor for the conversion to take place.

Reclaiming Objects Associated with IPv4 and IPv6 Addresses

You can use the reclaim IP function to delete all objects, except the active DHCP lease, that are associated with a selected IP address. To delete a DHCP lease, use the clear lease function as described in [Clearing Active DHCP Leases](#). When you reclaim an IP address, Grid Manager deletes the associated objects and puts them in the Recycle Bin, if enabled. You can reclaim any used and unmanaged IP addresses. You can also select multiple IP addresses for this function. After you reclaim an IP address, the address status changes to Unused. You can then reassign the IP address to other objects. For example, when you reclaim a fixed address, Grid Manager deletes the fixed address object and puts it in the Recycle Bin. When you reclaim an IP address that is associated with a host record and the address is the only address in the host, Grid Manager deletes the host record.

Grid Manager deletes all the objects that are associated with the selected IP addresses and puts them in the recycle bin, with the following exceptions:

- When you reclaim IP addresses that are in a DHCP range, all the objects that are associated with the IP addresses are deleted and the IP addresses remain in the DHCP range.
- When you select an IP address that is part of a host record, only the selected IP address is deleted from the host. However, if the selected address is the only address in the host, Grid Manager deletes the host record.

Grid Manager does not reclaim the following:

- Unused IP addresses
- Bulk hosts

To reclaim an IP address:

1. From the IP Map or List panel, select the IP address you want to reclaim, and then click **Reclaim** from the Toolbar. You can select multiple IP addresses.
2. In the *Delete Confirmation* dialog box, click **Yes**.
Grid Manager puts the deleted objects in the Recycle Bin, if enabled.

Pinging IP Addresses

You can find out whether an IP address is accessible and active by pinging the address. Grid Manager sends a packet to the selected IP address and waits for a reply when you ping the address. You can ping individual IP addresses from the IP Map and IP List panels. You can ping all IP addresses from the IP Map panel only.

To ping an IPv4 or IPv6 address:

- From the IP Map or IP List panel, select the IP address that you want to ping, and then click **Ping** from the Toolbar.

To ping all IPv4 addresses:

- From the IP Map panel, click **Multi-ping** from the Toolbar. Grid Manager pings all IP addresses displayed in the IP Map panel and displays the ping status in the panel.

When the ping or multi-ping is complete, the status bar displays the number of active IP addresses detected through the ping. To close the ping status bar, click the Close icon.

Clearing Active DHCP Leases

A DHCP lease specifies the amount of time that the DHCP server grants to a network device the permission to use a particular IP address. You may sometimes need to terminate an active lease. The following are some of the reasons for clearing active DHCP leases:

- When a network device is moved to another network.
- Reset a DHCP lease to fix other problems.

In Grid Manager, you can select multiple IP addresses and clear their active DHCP leases.

To clear an active lease:

1. From the IP Map or List panel, select the IP address for which you want to clear a DHCP lease, and then click **Clear** -> **Clear Lease** from the Toolbar. You can select multiple IP addresses.
2. In the *Clear DHCP Lease Confirmation* dialog box, click **Yes**.



Chapter 13 Network Discovery

This chapter provides information about the Infoblox discovery process, and how you can use the discovery feature to gather and manage information about predefined networks as well as virtual entities on VMware vSphere servers. It also explains how to integrate and view discovered data from Infoblox PortIQ and Trinzic Network Automation appliances. This chapter includes the following sections:

- [*About Network Discovery*](#) on page 494
 - [*Administrative Permissions*](#) on page 495
 - [*IP Discovery Process*](#) on page 496
 - [*Supported IP Discovery Methods*](#) on page 497
 - [*VM Discovery Process*](#) on page 499
- [*About Configuring a Discovery*](#) on page 500
 - [*Before Starting a Discovery*](#) on page 501
 - [*Selecting a Grid Member*](#) on page 502
 - [*Enabling or Disabling the Merging of Discovered Data*](#) on page 502
 - [*Updating Discovered Data for Managed Objects*](#) on page 503
 - [*Configuring IP Discovery*](#) on page 503
 - [*Configuring VM Discovery*](#) on page 504
 - [*Guidelines for Starting and Scheduling a Discovery*](#) on page 505
 - [*Starting a Discovery Immediately*](#) on page 505
 - [*Scheduling a Discovery*](#) on page 506
 - [*Configuring a Recurring Discovery*](#) on page 506
 - [*Managing a Discovery*](#) on page 507
 - [*Monitoring Discovery Status*](#) on page 507
- [*Integrating Data from PortIQ Appliances*](#) on page 508
- [*Integrating Discovered Data From Trinzic Network Automation*](#) on page 509
- [*Viewing Discovered Data*](#) on page 510
- [*Managing Discovered Data*](#) on page 511
 - [*Managing Unmanaged Data*](#) on page 511
 - [*Resolving Conflicting Addresses*](#) on page 513
 - [*Clearing Discovered Data*](#) on page 515

ABOUT NETWORK DISCOVERY

Note: The discovery features described in this chapter apply to NIOS Grid deployments that do not use the special Discovery license and its accompanying features under Network Insight. The Network Insight discovery features provide the ability to discover, query and catalog routed and switched networks and the devices within them, including infrastructure routers, enterprise switches, security devices such as firewalls, wireless access points, end host computer systems, and more. For more information about Network Insight discovery, see the chapter [Network Insight](#) and its associated topics.

The appliance provides discovery tools for detecting active hosts on predefined networks and on specified VMware vSphere servers. You can use the discovery feature to obtain and manage information about your network hosts. You can start a discovery immediately after you configure it, schedule it for a later date and time, or configure a recurring discovery based on the recurrence pattern. A recurring discovery occurs repeatedly based on the regular schedule you have configured. For information about how to start or schedule a discovery, see [Guidelines for Starting and Scheduling a Discovery](#) on page 505.

Depending on which discovery method you use, the appliance returns information, such as IP addresses, MAC addresses, and operating systems, about the detected hosts and virtual entities. You can include one or both of the following in a discovery task:

- IP discovery: The appliance detects active hosts on specified networks in a network view. You can perform an IP discovery using the following protocols: ICMP (Internet Control Message Protocol), NetBIOS (Network Basic Input/Output System), and TCP (Transmission Control Protocol). For information, see [Supported IP Discovery Methods](#) on page 497.
- VM discovery: The appliance detects active hosts on specified VMware vSphere servers. It also collects vSphere-specific data about the virtual entities on the specified vSphere servers. For information, see [VM Discovery Process](#) on page 499.

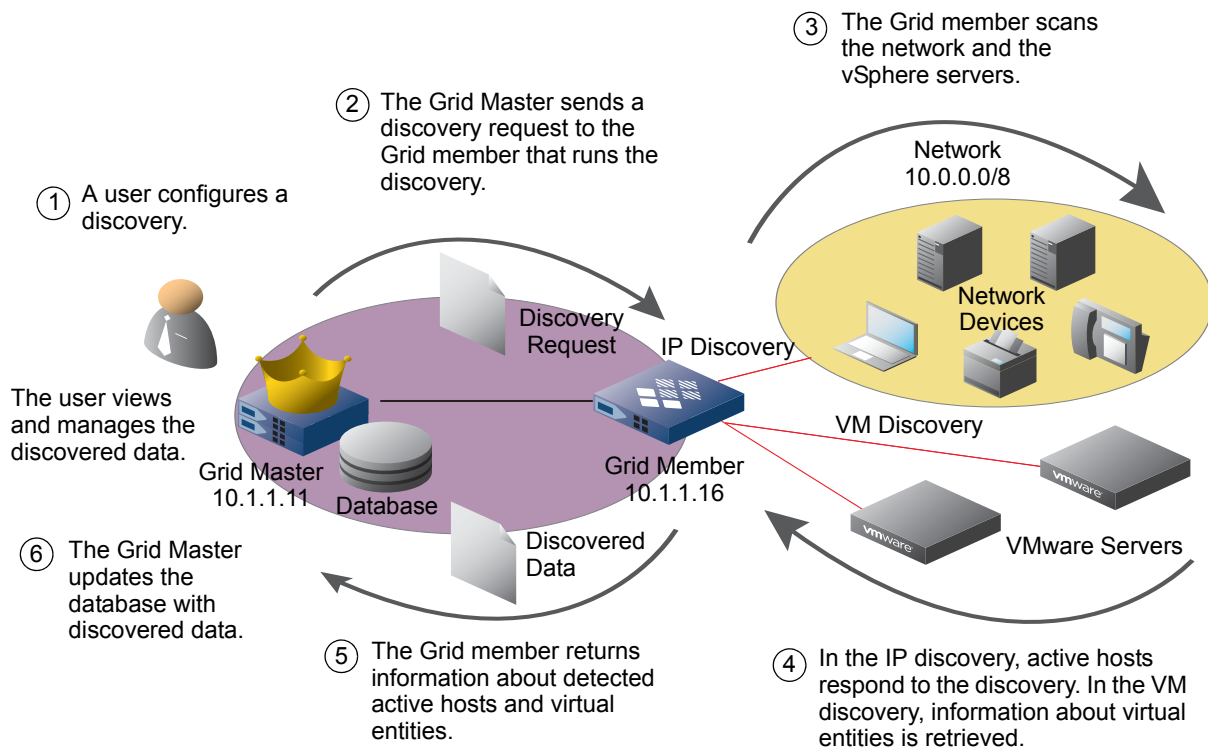
After a discovery, the appliance updates the database with the discovered data. It can either merge the newly discovered data or update only the unmanaged data. Unmanaged data is information that is not configured for DNS or DHCP before the discovery. For information about the guidelines the appliance uses to update the database, see [Before Starting a Discovery](#) on page 501.

You can configure and initiate a discovery from the *Discovery Manager* wizard. You must first select a Grid member to run the discovery. You can run an IP discovery on a set of networks and a VM discovery on a set of VMware vSphere servers. After you configure a discovery task, the Grid Master sends a discovery request to the selected Grid member. The discovery request contains information, such as the target network view, networks, and discovery method. Depending on your configuration, the selected Grid member runs an IP discovery on the predefined networks. When VM discovery is configured, it also collects information about virtual entities from the specified vSphere servers. The Grid member then reports the discovered results to the Grid Master. For information, see [About Configuring a Discovery](#) on page 500.

After a discovery is complete, you can view and manage the discovered data. For information, see [Viewing Discovered Data](#) on page 510 and [Managing Discovered Data](#) on page 511. You can also use the discovered data, such as unmanaged data, last discovered timestamps, and virtual machine data, as filters for Smart Folders. For information, see [Creating Smart Folders](#) on page 142. The appliance records all discovery operations in the audit log.

Figure 13.1 shows a high-level perspective of the discovery process.

Figure 13.1 High-Level Discovery Process



Administrative Permissions

You can initiate a discovery and manage discovered data based on your administrative permissions. For information, see [Managing Administrators](#) on page 149.

You must have read/write permission to discovery to initiate and control a discovery. The following are permission guidelines for initiating and controlling a discovery:

- Superusers can initiate and control a discovery on all networks.
- Administrators with read/write permission to discovery can initiate and control an IP discovery on networks to which they have read/write or read-only permission. They can also initiate a VM discovery. Only the objects with IP addresses to which the administrators have read/write permission are updated to include the vSphere discovered data.

After a discovery is complete, the following permission guidelines apply to viewing and managing discovered data:

- Superusers can view and manage all discovered data.
- Administrators with read/write permission to networks can view all discovered data. They can also add unmanaged data to existing hosts, and resolve IP address conflicts.
- Only administrators with read/write permission to a DNS zone or specific record type can convert unmanaged data to a host, fixed address, reservation, A record, or PTR record.
- Administrators with read-only permission to networks can only view discovered data. They cannot change any discovered data.

IP Discovery Process

Once an IP discovery starts, the Grid member reports the discovery status, such as **Completed**, **Running**, **Paused**, **Stopped**, or **Error**, in the *Discovery Manager* wizard and the *Discovery Status* widget on the Dashboard. In the *Discovery Status* widget, Grid Manager reports the time when the discovery status was last updated and the numbers of each type of discovered data. For information, see [Monitoring Discovery Status](#) on page 507.

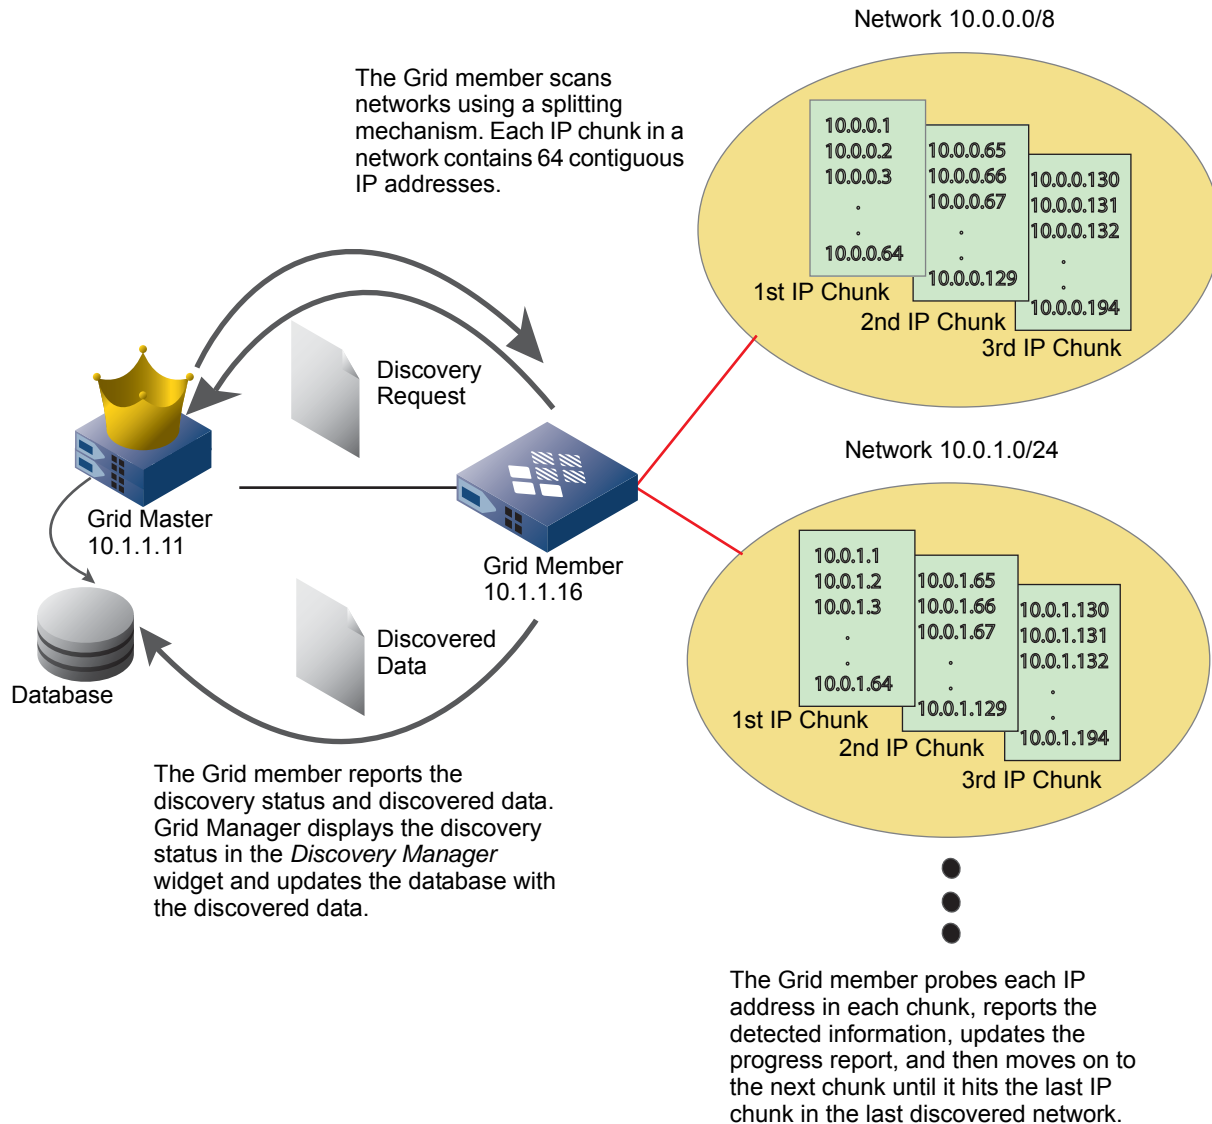
When an IP discovery starts, the appliance divides the IP addresses in a network into chunks, with each chunk containing 64 contiguous IP addresses. The discovery process probes each IP address in parallel and in ascending order, reports the detected information, updates the progress report, and then moves on to the next chunk until it hits the last chunk of IP addresses. The appliance then updates the database with the discovered data.

An IP discovery scans the selected networks in the order the networks appear in the *Discover Manager* wizard.

You can configure discovery processes on the same network, but the same configuration cannot be shared between two discovery processes.

[Figure 13.2](#) illustrates how an IP discovery works.

Figure 13.2 IP Discovery Process



Supported IP Discovery Methods

When you perform an IP discovery, you can choose one of the following discovery methods:

- ICMP as described in [ICMP](#).
- NetBIOS as described in [NetBIOS](#).
- TCP as described in [TCP](#) on page 498.
- Full as described in [Full](#) on page 498.

These methods actively scan predefined networks and probe IP addresses. The Grid member listens for responses from the IP addresses as proof of activities. The IP discovery scans through the specified network ranges and probes IP addresses (except for the network, broadcast, and multicast address types) in each network, including the /31 and /32 subnets. Note that the possible addresses in the /31 and /32 subnets can be used only as source addresses for point-to-point links. In these cases, there are no broadcast or network addresses in the /31 and /32 subnets, and the appliance can discover source addresses in these subnets.

ICMP

This method detects active hosts on a network by sending ICMP echo request packets (also referred to as pings) and listening for ICMP echo responses. The ICMP discovery is a simple and fast discovery that detects whether an IP address exists or not. It returns only the IP address and MAC address (only if the Grid member running the discovery is on the same discovered network) of a detected host. The ICMP discovery might miss some active hosts on the network due to security measures that are put in place to block ICMP attacks.

You configure the timeout value and the number of attempts in the *Discovery Manager* wizard. The ICMP discovery method returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: The discovery returns the MAC address only if the Grid member running the discovery is on the same discovered network.

To use the ICMP discovery method, the ICMP protocol between the Grid member performing the discovery and the target networks must be unfiltered.

NetBIOS

The NetBIOS method queries IP addresses for an existing NetBIOS service. This method detects active hosts by sending NetBIOS queries and listening for NetBIOS replies. It is a fast discovery that focuses on Microsoft hosts or non-Microsoft hosts that run NetBIOS services.

You configure the timeout value and the number of attempts in the *Discovery Manager* wizard. This method returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: Only if the discovered host is running Microsoft.
- OS: This value is set to **Microsoft** for an active host that has a MAC address in the NetBIOS reply.
- NetBIOS name: This value is set to the name returned in the NetBIOS reply.

To use the NetBIOS discovery method, ports 137 (UDP/TCP) and 139 (UDP/TCP) between the Grid member performing the discovery and the target networks must be unfiltered.

TCP

The TCP discovery probes each active host on a list of TCP ports using TCP SYN packets. This method detects all active hosts that generate SYN ACK responses to at least one TCP SYN. The discovery can determine the OS on a host by analyzing how the host reacts to the requests on opened and closed ports. It then uses the TCP fingerprints to guess the OS. To obtain a TCP fingerprint, IP discovery provides two scanning techniques, SYN and CONNECT.

When you use the SYN technique, the discovery sends a TCP SYN packet to establish a connection on a TCP port. If the port is open, the host replies with a SYN ACK response. The discovery does not close the port connection.

The CONNECT technique is a three-way TCP handshake. The discovery starts with the same process as the SYN technique by sending the TCP SYN packet. If the host replies with a SYN ACK response, the discovery then sends a RST packet to close the connection. If the response contains a RST flag, it indicates that the port is closed. If there is no reply, the port is considered as filtered. The TCP discovery is a deliberate and accurate discovery method. It can basically detect all active hosts on a network provided that there are no firewalls implemented on the network.

You can select the TCP ports, the TCP scanning technique, and configure the timeout value and the number of attempts in the *Discovery Manager* wizard. This method returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: The discovery returns the MAC address only if the Grid member running the discovery is on the same discovered network.
- OS: This is set to the highest probable OS reported in the response.

To use the TCP discovery method, the TCP port and a specific set of ports between the Grid member and the discovered networks must be unfiltered. The default set of ports is defined by the factory settings.

Full

The full discovery method is a combination of an ICMP discovery, a NetBIOS discovery, a TCP discovery, and a UDP scan. This method starts by sending an ICMP echo request. If no IP address on the network responds to the ICMP request, the discovery ends. If there is at least one response to the ICMP echo request, a NetBIOS discovery starts. A TCP discovery then follows by skipping through the active hosts that the NetBIOS discovery detects. The TCP discovery also handles the NetBIOS-detected hosts that have no MAC addresses. This method also performs a UDP scan to determine which UDP ports are open.

You configure the timeout value and the number of attempts in the *Discovery Manager* wizard. The full discovery method returns the following information for each detected host:

- IP address
- MAC address
- OS
- NetBIOS name

To use the full discovery, all the filter and firewall requirements in the ICMP, NetBIOS, and TCP discovery methods apply.

The following is a summary of the supported IP discovery methods:

Discovery Type	Returned Data	Guideline	Mechanism
ICMP	<ul style="list-style-type: none"> IP address MAC address 	Use ICMP for a rough and fast discovery	ICMP echo request and reply
NetBIOS	<ul style="list-style-type: none"> IP address MAC address OS NetBIOS name 	Use NetBIOS for discovering Microsoft networks or non-Microsoft networks that run some NetBIOS services	NetBIOS query and reply
TCP	<ul style="list-style-type: none"> IP address MAC address OS 	Use TCP for an accurate but slow discovery	TCP SYN packet and SYN ACK packet
Full	<ul style="list-style-type: none"> IP address MAC address OS NetBIOS name 	Use Full for a general and comprehensive discovery	<ol style="list-style-type: none"> 1. ICMP echo request and reply 2. NetBIOS query and reply 3. TCP SYN packet and SYN ACK packet

The method you select to run an IP discovery determines the kind of information the discovery returns and the time it takes to complete an IP discovery. If time is a concern, the following are factors you may consider when configuring an IP discovery:

- The timeout value
- The number of attempts
- The number of ports the discovery scans
- The size of network you want to discover

VM Discovery Process

When you perform a VM discovery, the appliance communicates with the specified vSphere servers to collect vSphere-specific data. Unlike IP discovery, VM discovery processes all the IP addresses on the specified vSphere servers. Therefore, VM discovery can discover IP addresses in all the networks within the selected network view.

The following is a summary of the VM discovery:

Discovery Type	Returned Data	Guideline	Mechanism
VM discovery	<ul style="list-style-type: none"> IP address MAC address OS Discovered name Virtual entity type Virtual entity name Virtual cluster Virtual datacenter Virtual switch Virtual host Virtual host adapter 	Add the VMware vSphere servers on which you want to perform the VM discovery	The appliance communicates with the vSphere servers to collect discovered data

ABOUT CONFIGURING A DISCOVERY

You can configure and control a discovery from the *Discovery Manager* wizard, which is accessible from the *Discovery Status* widget on the Dashboard or from the **Data Management** tab -> **IPAM** tab. In a Grid, only one member can run a discovery at a time. Multiple members cannot run discoveries simultaneously.

The following are guidelines for configuring a discovery task:

- You must have read/write permission to the discovery process to initiate a discovery.
- After you start a discovery, you cannot change the configuration of the discovery.

You can perform the following tasks from the *Discovery Manager* wizard:

- Select the Grid member from which you want to run the discovery. For an IP discovery, the Grid member does not need to be assigned to the discovered network or within a DHCP range. For information, see [Selecting a Grid Member](#) on page 502.
- Optionally, enable or disable the merging of discovered data with existing data. This function is enabled by default. For information, see [Enabling or Disabling the Merging of Discovered Data](#) on page 502.
- Configure an IP discovery. For information, see [Updating Discovered Data for Managed Objects](#) on page 503. You must first define the networks on which you want to run an IP discovery. For information, see [Configuring DHCP for IPv4](#) on page 843.
- Add the vSphere servers on which you want to run a VM discovery. For information, see [Configuring VM Discovery](#) on page 504.

You can include both the IP discovery and VM discovery in a discovery task. Note that the IP discovery and VM discovery are enabled by default for a new NIOS installation. For an upgrade, the IP discovery is enabled and the VM discovery is disabled.

After you configure a discovery, you can start the discovery process immediately or schedule it for a later date. For information, see [Starting a Discovery Immediately](#) on page 505 and [Scheduling a Discovery](#) on page 506. You can also configure a recurring discovery that repeats on a regular basis. For information, see [Configuring a Recurring Discovery](#) on page 506. After you start a discovery, you can pause or stop it. For information, see [Managing a Discovery](#) on page 507. The appliance saves the configuration of the last discovery.

You can do the following after a discovery is complete:

- View the discovery status. You can view the current discovery status in the *Discovery Status* widget on the Dashboard. For information, see [Dashboards](#) on page 97.
- View the discovered data. For information, see [Viewing Discovered Data](#) on page 510.
- Manage the discovered data. For information, see [Managing Discovered Data](#) on page 511.

Before Starting a Discovery

The following are some guidelines for consideration before you start a discovery.

Database updates

After the Grid Master receives discovery data from the Grid member, it integrates the data based on the following rules:

- For a discovered host with a new IP address, the appliance marks the IP address “unmanaged”.
- For a discovered host associated with one of the following, the appliance updates the data of the associated object:
 - A fixed address reservation or host address reservation
 - A host address not configured for DHCP services
 - A fixed address or host address with the same MAC address as that of the discovered host
 - An A or PTR record
 - A DHCP lease with the same MAC address as that of the discovered host
- For a DHCP lease that does not have any associated object, such as a fixed address or host record, the appliance updates the IP address with the discovered data. When the lease expires and the IP address has no associated objects, the appliance marks the IP address “unmanaged”. When the lease expires and the IP address is associated with the same MAC address, the appliance preserves the discovered data.
- For a discovered host associated with one of the following, the appliance updates all data except the MAC address and marks the IP address as a conflict. For information, see [Resolving Conflicting Addresses](#) on page 513.
 - A fixed address with a different MAC address than that of the discovered host
 - A DHCP lease with associated objects and with a different MAC address than that of the discovered host
 - An Infoblox host address configured for DHCP services and with a different MAC address than that of the discovered host
- For a discovered host that is part of a DHCP range but does not have a fixed or leased address or is not within an exclusion range, the appliance assigns a DHCP range conflict to the IP address.
- For a discovered host through a VM discovery, the appliance adds the discovered data to the database. The data is displayed in the IP Map and IP List panels, the **Discovered Data** tab of an object editor, and the Discovered Data section of the IP Address panel.
- The OS of an IP address obtained by an IP discovery supersedes that obtained by a VM discovery, and the newly discovered name of a host supersedes the last discovered data.
- When a VM discovery cannot obtain the IP address of a virtual entity, it does not return any discovered data for the entity.
- Only the objects with IP addresses to which the administrators have read/write permission are updated to include the VM discovery data.

Database Capacity

When the Grid Master database reaches its maximum capacity (the maximum capacity varies based on the appliance model), the Grid Master stops updating the database and requests that the Grid member stop the discovery. When the discovering Grid member database reaches its capacity, the Grid member pauses the discovery. The appliance displays a dialog to inform you that the discovery pauses. The Grid member resumes the discovery once the database falls below its capacity. When a discovery pauses because of capacity issues, you cannot resume the discovery or start a new discovery. You can check the capacity of your appliance database before starting a discovery.

HA Failover

In an HA pair, if the Grid Master fails over to the passive node, the passive node takes over and continues with the discovery from the last known state. If an independent appliance fails, the appliance stops the discovery process and keeps the discovery in a paused state. The appliance resumes the discovery once it starts up again.

Selecting a Grid Member

You must select a Grid member from which you want to run a discovery.

1. From the **Data Management** tab, select the **IPAM** tab, then select **Immediate** or **Recurring** from the **Discovery** drop-down list in the **Toolbar**.
or
From the *Discovery Status* widget, select **Immediate** or **Recurring** from the drop-down list and click **Discovery Manager**.
2. In the *Discovery Manager* wizard, click the **General** tab and complete the following:
 - **Current Status:** Displays the last discovery status and timestamp. This data is read-only.
 - **Member Name:** Click **Select Member**. In the *Member Selector* dialog box, select the Grid member from which you want to run the discovery. You can also use filters or the Go to function to find a specific member. For information, see [Using Filters](#) on page 67 and [Using the Go To Function](#) on page 71.

After you select a Grid member, you can do the following:

- Enable or disable the merging of discovered data, as described in [Enabling or Disabling the Merging of Discovered Data](#).
- Configure an IP discovery, as described in [Configuring IP Discovery](#) on page 503.
- Configure a VM discovery, as described in [Configuring VM Discovery](#) on page 504.

Enabling or Disabling the Merging of Discovered Data

You can decide whether to merge the newly discovered data with the current data in the database. This function is enabled by default.

To enable or disable the merging of discovered data:

1. From the **Data Management** tab, select the **IPAM** tab, then select **Immediate** or **Recurring** from the **Discovery** drop-down list in the **Toolbar**.
or
From the **Discovery Status** widget, select **Immediate** or **Recurring** from the drop-down list and click **Discovery Manager**.
2. In the *Discovery Manager* wizard, click the **General** tab and complete the following:
 - **Merge the discovered data with existing data:** When you select this check box, the appliance merges the discovered data with the existing data. It appends newly discovered data to existing data and preserves the existing data when there is no newly discovered data. This check box is selected by default.

Note: When you clear this check box, the appliance replaces the existing data with the newly discovered data and if there are no newly discovered values for some fields, the appliance removes the existing values for these fields.

You can do the following:

- Configure an IP discovery, as described in [Configuring IP Discovery](#) on page 503.
- Configure a VM discovery, as described in [Configuring VM Discovery](#) on page 504.

Updating Discovered Data for Managed Objects

You can decide whether you want the appliance to update the data of existing A records, PTR records, host records, and fixed addresses.

To update discovered data for managed objects:

1. From the **Data Management** tab, select the **IPAM** tab, then select **Immediate** or **Recurring** from the **Discovery** drop-down list in the **Toolbar**.
or
From the *Discovery Status* widget, select **Immediate** or **Recurring** from the drop-down list and click **Discovery Manager**.
2. In the *Discovery Manager* wizard, click the **General** tab and complete the following:
 - **Update discovered data for managed objects:** Select this check box if you want the appliance to update the data of existing A records, PTR records, host records, and fixed addresses. If you do not select this check box, the appliance updates only the unmanaged objects.

You can do the following:

- Configure an IP discovery, as described in [Configuring IP Discovery](#).
- Configure a VM discovery, as described in [Configuring VM Discovery](#) on page 504.

Configuring IP Discovery

When you start an IP discovery from the IPAM Home, Net Map or network List panel, you can select the networks on which you want the discovery to run. When you start an IP discovery from the IP Map or IP List panel, the discovered network is the one to which the IP addresses belong. You can include additional networks when you configure the IP discovery from the *Discovery Manager* wizard. You can run an IP discovery on multiple networks in one network view.

1. From the **Data Management** tab, select the **IPAM** tab, then select **Immediate** or **Recurring** from the **Discovery** drop-down list in the **Toolbar**.
or
From the *Discovery Status* widget, select **Immediate** or **Recurring** from the drop-down list and click **Discovery Manager**.
2. In the *Discovery Manager* wizard, click the **IP Discovery** tab, and then complete the following in the **Basic** tab:
 - **Mode:** Select the IP discovery method you want to use. For information, see [Supported IP Discovery Methods](#) on page 497. If you select **TCP** or **FULL**, ensure that you configure the TCP ports in the **Advanced** tab. The default is **Full**.
 - Click the Add icon to add networks. In the *Network Selector* dialog box, select the network view and networks. Use SHIFT+click and CTRL+click to select multiple networks. You can also use filters or the Go to function to find a specific network. For information, see [Using Filters](#) on page 67 and [Using the Go To Function](#) on page 71.

You can do the following in the table:

- Click the Add icon again to add more networks.
 - Select a network or multiple networks in the network table and click **Delete** to delete them.
 - Click the Export icon to export the data in CSV format.
 - Click the Print icon to print the data.
 - **Disable:** Select this to exclude IP discovery from the discovery task and to run a VM discovery only. IP discovery is enabled by default. You can configure a VM discovery in the same discovery task. You cannot start a discovery if you disable both IP discovery and VM discovery.
3. If you select **TCP** or **FULL** in **Mode**, click the **Advanced** tab and complete the following:
 - **TCP Scan Technique:** Select the TCP technique you want to use for the discovery. The default is SYN. For information, see [TCP](#) on page 498.

- In the port table, select the check box of the port you want to configure. You can select all ports by clicking the check box in the header.

Optionally, you can click the Add icon and complete the following to add a new service to the list.

- **Port:** Enter the port number you want to add to the list. You must enter a number between 1 and 65535.
- **Service:** Enter the name of the service.

You can also delete a specific TCP port in the list. You can select multiple ports for deletion.

- **Timeout (ms):** Enter the timeout value in milliseconds for the discovery. The timeout value determines how long the discovery waits for a response from an IP address after probing it. The minimum is 5 and the maximum is 4000. The default is 1000.
 - **Attempts:** Enter the number of times you want the discovery to probe an IP address when scanning a network. The minimum is 1 and the maximum is 5. The default is 2.
4. Start the discovery or save the configuration, as described in [Starting a Discovery Immediately](#) on page 505. You can also schedule the discovery to run at a later date, as described in [Scheduling a Discovery](#) on page 506 or configure a recurring discovery, as described in [Configuring a Recurring Discovery](#) on page 506.

Configuring VM Discovery

Note: If you install NIOS appliances to perform SNMP-based device discovery with a special Discovery license, usage methods change for VM Discovery. For related information, see the section [Performing VM \(Virtual Machine\) Discovery](#) on page 543.

A VM discovery retrieves information about vSphere servers and the virtual entities running on the servers. You can add more than one vSphere server to the discovery. When you disable specific servers, the appliance excludes them from the VM discovery.

1. From the **Data Management** tab, select the **IPAM** tab, then select **Immediate** or **Recurring** from the **Discovery** drop-down list in the **Toolbar**.
or
From the *Discovery Status* widget, select **Immediate** or **Recurring** from the drop-down list and click **Discovery Manager**.
2. In the *Discovery Manager* wizard, click the **VM Discovery** tab, and then complete the following:
Network View: Select the network view in which you want to perform the VM discovery. This appears only when there are multiple network views. Otherwise, the VM discovery is performed in the default network view.
Click the Add icon and complete the following in the Add vSphere Server section:
 - **Server:** Enter the FQDN or IP address of the vSphere server.
 - **Protocol:** Select the protocol that is used to connect to the vSphere server. The default is HTTPS.
 - **Port Number:** Enter the number of the port the appliance uses to communicate with the vSphere server. The default is 443.
 - **Username:** Enter the username the appliance uses to log in to the vSphere server. The user account on the vSphere server should have at least read-only permission.
 - **Password:** Enter the password of the vSphere server account.
3. Click **Test** to test the settings before adding them to the table.
4. Click **Add** to add the vSphere server to the table. You can also do the following in the table:
 - Click the Add icon again to add more vSphere servers.
 - Select the **Disable** check box in the table to exclude a specific vSphere server from the VM discovery. The appliance keeps the server configuration when you disable the server. All servers in the list are included in the VM discovery by default.
 - Select a server and click the Edit icon to modify its configuration.
 - Select a server and click the Delete icon to delete the server.

— Click the Export icon to export the data in CSV format.

5. Optionally, select the **Disable** check box to exclude VM discovery from the discovery task. VM discovery is enabled by default. The appliance keeps the server configurations when you disable VM discovery.
6. Start the discovery or save the configuration, as described in [Starting a Discovery Immediately](#) on page 505. You can also schedule the discovery to run at a later date, as described in [Scheduling a Discovery](#) on page 506 or configure a recurring discovery, as described in [Configuring a Recurring Discovery](#) on page 506.

Guidelines for Starting and Scheduling a Discovery

After you configure a discovery, you can start the discovery process immediately or schedule it for a later date. You can also configure a recurring discovery that repeats on a regular basis. When you start a discovery immediately or schedule for a later date after you configure it, the discovery happens only once and it will not be repeated. For more information about how to start a discovery immediately or schedule it for a later date, see [Starting a Discovery Immediately](#) on page 505 and [Scheduling a Discovery](#) on page 506.

To repeat a discovery regularly, you can configure a recurring discovery. A recurring discovery occurs repeatedly based on the schedule you have configured. When you select **Recurring** from the drop-down list, you can schedule the discovery process to recur on an hourly, daily, weekly, or monthly basis. For information about how to configure a recurring discovery, see [Configuring a Recurring Discovery](#) on page 506.

You can configure these discovery jobs independent of each other and each one contains a specific set of networks and discovery settings.

Note the following guidelines about immediate, regular and recurring discovery tasks:

- You cannot run regular and recurring discovery processes concurrently.
- If a recurring discovery is scheduled to start when a discovery is in progress, the recurring discovery will be postponed to the next schedule time. The current recurring discovery will not be performed.
- You can pause and resume all discovery tasks.
- You cannot start a discovery when another one has been paused.
- You cannot use the `start` command to start a recurring discovery.
- Discovery permissions are applicable to all discovery tasks.

Starting a Discovery Immediately

You can start a discovery immediately after you complete its configuration. To configure an immediate discovery:

1. From the **Data Management** tab, select the **IPAM** tab, and then select **Immediate** from the **Discovery** drop-down list in the **Toolbar**.

or

From the *Discovery Status* widget, select **Immediate** from the drop-down list and click **Discovery Manager**.

2. In the *Discovery Manager* wizard, select one of the following:
 - **Restore to Defaults:** Restores the discovery configuration using the default values.
 - **Start:** Starts the discovery.
 - **Save:** Saves the discovery configuration.
 - **Close:** Cancels the configuration. If you have started a discovery, the discovery runs in the background when you click **Close**. For information, see [Running Tasks in the Background](#) on page 85.

Note: Once you start a discovery, you cannot change the discovery configuration. After you click **Start**, the button changes to **Pause**. You can click **Pause** to pause a discovery. When the discovery is paused, the button changes to **Resume**. You can click **Resume** to continue the paused discovery.

Scheduling a Discovery

After you configure a discovery, you can schedule to run it at a later date and time. Note that you can only schedule the start of a discovery, you cannot schedule it to pause, stop, or resume. After a scheduled discovery starts, you can then pause, stop, or resume it.

To schedule a discovery, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone. Click **Schedule Start** to schedule the discovery. If applicable, you can select **Click here to view/manage the scheduled items to reschedule a discovery or view all scheduled discoveries**.

You can schedule only one discovery at a time. Once you schedule a discovery, you cannot change the configuration until the task is cancelled or executed.

Configuring a Recurring Discovery

After you configure a discovery, you can schedule a recurring discovery by configuring a recurrence pattern. The appliance automatically starts a recurring discovery based on the configured schedule and detects any newly added or removed networks. Note that you can only schedule the start of a discovery, you cannot schedule a discovery to pause, stop, or resume. After a scheduled discovery starts, you can then pause, stop, or resume it. You can schedule only one discovery at a time. Once you schedule a discovery, you cannot change the configuration until the task is cancelled or executed. You can disable the recurring network discovery task. When you disable this task, it will not recur during the scheduled interval.

If the discovery task fails during a scheduled interval, then the task stops and will not continue for the corresponding occurrence. The scheduled task resets and the discovery starts at the next scheduled time. For example, when you configure a recurring discovery to occur every five hours, discovery starts at the following hours on each day: 00:00, 05:00, 10:00, 15:00, and 20:00. If the discovery scheduled for 05:00 fails, the discovery starts at the next recurrence, which is at 10:00. For information about failed discovery, see [About Configuring a Discovery](#) on page 500.

The following examples explain when a recurring discovery starts based on your configuration:

Example 1

When you configure a recurring discovery to occur every five hours, the discovery starts at the following hours on each day: 00:00, 05:00, 10:00, 15:00, and 20:00. The first occurrence on each day starts at 00:00.

Example 2

When you configure a recurring discovery to occur every two days during a week, the discovery starts on the following days every week: Monday, Wednesday, Friday, and Sunday. The first occurrence starts on Monday of each week.

To configure a recurring discovery:

1. From the **Data Management** tab, select the **IPAM** tab, and then select **Recurring** from the **Discovery** drop-down list in the **Toolbar**.
or
From the *Discovery Status* widget, select **Recurring** from the drop-down list and click **Discovery Manager**.
2. In the *Discovery Manager* wizard, select an option under **Recurrence** to schedule the network discovery task to run on a recurring basis. Complete the following:
 - **Disable:** Select this check box to disable recurring network discovery. If you select this check box, network discovery process will not recur during the scheduled interval. Clear the check box to enable recurring network discovery.
 - **Hourly:** Select this option if you want the discovery task to recur on an hourly basis.
 - **Recur every _ hour(s) at:** Specify the number of hours after which the discovery task must recur. The value must be an integer between 1 and 24. By default, this value is displayed as 1.
 - **_ minutes past the hour:** Specify minutes past the hour after which the discovery task must recur. The value must be an integer between 0 and 59. By default, this value is displayed as 0.
 - **Time Zone:** Select a time zone from the drop-down list.

- **Daily:** Select this option if you want the discovery task to recur on a daily basis.

Recur daily

- Every day: Select this option if the discovery task must recur every day of the week.
- Every weekday: Select this option if the discovery task must recur every day of the week, except Saturday and Sunday.
- Time: Enter a time in hh:mm:ss AM/PM format or select a time from the wizard.
- Time Zone: Select a time zone from the drop-down list.

- **Weekly:** Select this option if you want the discovery task to recur on a weekly basis.

Recur every week on

- Sunday/Monday/Tuesday/Wednesday/Thursday/Friday/Saturday: Select a day of the week when you want the discovery task to recur.

Note: When you schedule the recurring task on a weekly basis, Monday is considered to be the first day of the week.

- Time: Enter a time in hh:mm:ss AM/PM format or select a time from the wizard.
- Time Zone: Select the time zone from the drop-down list.

- **Monthly:** Select this option if you want the discovery task to recur on a monthly basis.

Recur the day of the month

- Day _ every _ month(s): Select a day of the month when you want the discovery task to recur. You can schedule the task to recur once in a specified number of months. Specify an integer between 1 and 31 for the day field. By default, the day is set to one. For month(s), specify an integer between 1 and 12. The default value is one.

Note that when you define the task to occur on a specific day number that doesn't exist in a month, the occurrence is skipped for the corresponding month. For example, if you schedule the task to occur on the 30th of every month, this task will not be executed in February. If you schedule the task to occur on the 31st of every month, an error message is displayed for months which have fewer days than 31. For these months, recurrence will happen on the last day of the month.

- Time: Enter a time in hh:mm:ss AM/PM format or select a time from the wizard.
- Time Zone: Select a time zone from the drop-down list.

3. Save the configuration.

Managing a Discovery

You can do the following after you start a discovery:

- **Pause:** The appliance pauses the discovery at the current chunk of IP addresses.
- **Resume:** The appliance continues the discovery from the last **Pause** state. It resumes the discovery at the beginning of the first unprocessed chunk of IP addresses on the network.
- **Stop:** The appliance stops and terminates the discovery. It marks the operation as complete. You cannot resume this discovery. All discovered data remains intact in the database.

Monitoring Discovery Status

You can monitor the discovery status through the *Discovery Status* widget on the Dashboard. You can also start, pause, resume, and stop a discovery from the widget. For information, see [Discovery Status](#) on page 128.

INTEGRATING DATA FROM PORTIQ APPLIANCES

Infoblox PortIQ appliances discover and track where devices connect to your network switches, and provide information about the switch ports to which the devices connect. You can integrate the data discovered by PortIQ appliances into the NIOS appliance database, and then view the data in the IP Map and List panels of Grid Manager. For information about PortIQ appliances, refer to the *Infoblox Administrator Guide for PortIQ Appliances*.

You can import the following data about the IP addresses that PortIQ appliances discover:

- **Discovered Name:** The name of the network device associated with the discovered IP address.
- **Discoverer:** Specifies whether the IP address was discovered by a PortIQ or NIOS discovery process.
- **First Seen:** The date and time the IP address was first seen.
- **Attached Device Description:** A textual description of the switch that is connected to the end device.
- **Attached Device Address:** The IP address of the switch that is connected to the end device.
- **Attached Device Name:** If a reverse lookup was successful for the IP address associated with this switch, the host name is displayed here.
- **Attached Device Port Description:** A textual description of the switch port that is connected to the end device.
- **Attached Device Port:** The number of the switch port connected to the end device.
- **Attached Device:** Identifies the switch that is connected to the end device.
- **Port Duplex:** The negotiated or operational duplex setting of the switch port connected to the end device.
- **Port Link:** The link status of the switch port connected to the end device. Indicates whether it is connected.
- **Port Speed:** The interface speed, in Mbps, of the switch port.
- **Port Status:** The operational status of the switch port. Indicates whether the port is up or down.
- **VLAN Description:** The description of the VLAN of the switch port that is connected to the end device.
- **VLAN Name:** The name of the VLAN of the switch port.
- **VLAN:** The ID of the VLAN of the switch port.

Do the following to integrate data from PortIQ appliances into the NIOS appliance:

1. Configure the PortIQ appliance to synchronize its data with the NIOS appliance. In a Grid, PortIQ appliances must synchronize their data with the Grid Master. For information, refer to the *Infoblox PortIQ Appliance User Guide*.
2. Specify the data to be displayed in the IP Map and IP List panels.
 - When you select an IP address from the IP Map or IP List panel, this information can be displayed in the Discovered Data section. For information, see [Viewing Discovered Data](#) on page 510 and [Managing Discovered Data](#) on page 511.
 - In the IP List panel, you can add data fields. For information, see [Managing Discovered Data](#) on page 511.

INTEGRATING DISCOVERED DATA FROM TRINZIC NETWORK AUTOMATION

Trinzic Network Automation appliances discover and track IPv4 and IPv6 network devices and provide information about the discovered IP addresses. You can integrate IPv4 and IPv6 discovered data into the NIOS appliance database, and then view the data in the IP List panel of Grid Manager as well as in the **Discovered Data** tab of certain DNS and DHCP object editors. For information about Trinzic Network Automation network discovery and how to import discovered data from a Trinzic Network Automation appliance to the NIOS appliance, refer to the *Trinzic Network Automation Administrator Guide*.

Note: The NIOS appliance does not import IPv6 leases that contain prefixes and link-local IPv6 addresses. This data is discarded during an import.

The appliance can import the following IPv4 and IPv6 data that Network Automation discovers:

- **IP Address:** The discovered IPv4 or IPv6 address.
- **Discovered MAC Address:** The MAC address of the discovered host.
- **Last Discovered:** The date and time the IP address was last discovered.
- **NetBIOS Name:** The name returned in the NetBIOS reply or the name you manually register for the discovered host.
- **OS:** The operating system of the detected host.
- **First Discovered:** The date and time the IP address was first discovered.
- **Discoverer:** Specifies whether the IP address was discovered by a Network Automation discovery process.
- **Discovered Name:** The name of the network device associated with the discovered IP address.
- **Attached Device Description:** A textual description of the switch that is connected to the end device.
- **Attached Device Address:** The IP address of the switch that is connected to the end device.
- **Attached Device Name:** If a reverse lookup was successful for the IP address associated with this switch, the host name is displayed here.
- **Attached Device Port Description:** A textual description of the switch port that is connected to the end device.
- **Attached Device Port Name:** The name of the switch port connected to the end device.
- **Attached Device Port:** The number of the switch port connected to the end device.
- **Attached Device:** Identifies the switch that is connected to the end device.
- **Port Duplex:** The negotiated or operational duplex setting of the switch port connected to the end device.
- **Port Link:** The link status of the switch port connected to the end device. Indicates whether it is connected.
- **Port Speed:** The interface speed, in Mbps, of the switch port.
- **Port Status:** The operational status of the switch port. Indicates whether the port is up or down.
- **VLAN Name:** The name of the VLAN of the switch port.
- **VLAN:** The ID of the VLAN of the switch port.

VIEWING DISCOVERED DATA

After a discovery or after integrating data from PortIQ and Trinzie Network Automation appliances, you can view the discovered data in the following:

- IP Map panel for IPv4 addresses, as described in [IP Map](#) on page 474.
- IP List panel, as described in [IP Address List](#) on page 476.
- **Discovered Data** tab in certain DNS and DHCP object editors, as described below.

To specify the data fields to display in the **Discovered Data** tab of the IP Map and IP List panels:

1. Expand the **Discovered Data** tab.
2. Click the Configure icon.
3. Select a field from the **Available** list and click the ➤ arrow to move it to the **Selected** list. You can always move the fields between the two lists. Use SHIFT-click and CTRL-click to select multiple fields.
4. You can also change the order in which the fields are displayed by moving the fields up and down in the **Selected** list. To move a field up in the list, select it and click the Up arrow. To move a field down, select it and click the Down arrow.

The Discovered Data section displays the fields you specified. A discovery creates only **Unmanaged** and **Conflict** data. Depending on the source of the discovered data, when you modify certain DNS and DHCP objects, Grid Manager can display the following IPv4 and IPv6 discovered data (if any) in the **Discovered Data** tab:

- **NetBIOS Name:** The name returned in the NetBIOS reply or the name you manually register for the discovered host.
- **OS:** The operating system of the detected host or virtual entity. The OS can be one of the following:
 - **Microsoft** for all discovered hosts that have a non-null value in the MAC addresses using the NetBIOS discovery method.
 - A value that a TCP discovery returns.
 - The OS of a virtual entity on a vSphere server.
- **Discovered MAC Address:** The discovered MAC address for the host. This is the unique identifier of a network device. The discovery acquires the MAC address for hosts that are located on the same network as the Grid member that is running the discovery. This can also be the MAC address of a virtual entity on a specified vSphere server.
- **Discovered DUID:** For IPv6 address only. The DHCP unique identifier of the discovered host. This is an optional field, and data might not be included.
- **Last Discovered:** The date and time the IP address was last discovered.
- **First Discovered:** The date and time the IP address was first discovered.
- **Discoverer:** Specifies whether the IP address was discovered by a Network Automation or NIOS discovery process.
- **Discovered Name:** The name of the network device associated with the discovered IP address.

If you imported data from Infoblox PortIQ or Trinzie Network Automation appliances, Grid Manager may display the following information, if available. You can also select all or some of this data for display in the Discovered Data tab of the IP Map and IP List panels. For information about the data imported from PortIQ and Trinzie Network Automation appliances, see [Integrating Data from PortIQ Appliances](#) on page 508 and [Integrating Discovered Data From Trinzie Network Automation](#) on page 509.

- **Attached Device Description:** A textual description of the switch that is connected to the end device.
- **Attached Device Address:** The IPv4 or IPv6 address of the switch that is connected to the end device.
- **Attached Device Name:** If a reverse lookup was successful for the IP address associated with this switch, the host name is displayed here.
- **Attached Device Port Description:** A textual description of the switch port that is connected to the end device.
- **Attached Device Port Name:** The name of the switch port connected to the end device.

- **Attached Device Port:** The number of the switch port connected to the end device.
- **Attached Device:** Identifies the switch that is connected to the end device.
- **Port Duplex:** The negotiated or operational duplex setting of the switch port connected to the end device. You can modify this in the IPv6 fixed address and AAAA record editors.
- **Port Link:** The link status of the switch port connected to the end device. Indicates whether it is connected.
- **Port Speed:** The interface speed, in Mbps, of the switch port. You can modify this in the IPv6 fixed address and AAAA record editors.
- **Port Status:** The operational status of the switch port. Indicates whether the port is up or down.
- **VLAN Description:** The description of the VLAN of the switch port that is connected to the end device.
- **VLAN Name:** The name of the VLAN of the switch port.
- **VLAN:** The ID of the VLAN of the switch port.

For IP addresses discovered through a VM discovery, Grid Manager displays the following additional information, if available. You can also select this information for display in the **Discovered Data** tab of the IP Map and IP List panels:

- **Virtual Host Adapter:** The name of the physical network adapter through which the virtual entity is connected to the appliance.
- **Virtual Datacenter:** The name of the vSphere datacenter or container to which the virtual entity belongs.
- **Virtual Cluster:** The name of the VMware cluster to which the virtual entity belongs.
- **Virtual Entity Name:** The name of the virtual entity.
- **Virtual Entity Type:** The virtual entity type. This can be blank or one of the following: Virtual Machine, Virtual Host, or Virtual Center. Virtual Center represents a VMware vCenter server.
- **Virtual Host:** The name of the VMware server on which the virtual entity was discovered.
- **Virtual Switch:** The name of the switch to which the virtual entity is connected.

MANAGING DISCOVERED DATA

In addition to viewing the discovered data, you can do the following to manage the data:

- Manage an unmanaged address by adding it to a host, converting it to managed data, or clearing its unmanaged status. For information, see [Managing Unmanaged Data](#) on page 511.
- Resolve conflicting addresses. For information, see [Resolving Conflicting Addresses](#) on page 513.
- Clear all discovered data for a network or network view. See [Clearing Discovered Data](#) on page 515.
- Clear discovered data. For information, see [Clearing Discovered Data](#) on page 515.

Managing Unmanaged Data

You can manage unused and unmanaged addresses by doing one of the following:

- Add to an existing host, as described in [Adding to an Existing Host](#) on page 512.
- Convert to a fixed address, host, A record, or PTR record, as described in [Converting Unmanaged Data](#) on page 512.
- Clear the unmanaged status, as described in [Clearing Unmanaged Data](#) on page 513.

Note: You cannot convert unmanaged IP addresses served by Microsoft DHCP servers to host records.

Adding to an Existing Host

You can add an unmanaged address, including all its information, to an existing host. You can select the desired host to which you want to add the unmanaged address.

To add an unmanaged address to an existing host:

1. From the *IP Map* or *List* panel, select an unmanaged address you want to add to a host, and then click **Add -> Add to Existing Host** from the Toolbar.
2. In the *Select Host* dialog box, select a host from the table. You can also search for a host using filters or the Go to function. For information, see [Using Filters](#) on page 67 and [Using the Go To Function](#) on page 71. Click the Select icon to select the desired host.

Note: Depending on the page size configuration, the search results are limited to the page size that you set. If the search results exceed the page size limit, the appliance displays an error message to inform you to refine your search criteria or to change the page size limit. In the *Host Record* editor, complete the information as described in [Modifying Host Records](#) on page 463.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Converting Unmanaged Data

You can convert an unmanaged address to a host, an A record, a PTR record, or a fixed address.

To convert an unmanaged address:

1. In the *IP Map* or *List* panel, select an unmanaged address you want to convert, and then select **Convert** from the Toolbar.
2. In the drop-down list, select the type of address to which you want to convert the unmanaged address. For IPv4 addresses, you can select **To Host**, **To A Record**, **To PTR Record**, or **To Fixed Address**. For IPv6 addresses, you can select **To Host**, **To AAAA**, **To PTR Record**, or **To IPv6 Fixed Address**.

Depending on the record type you select, Grid Manager displays the corresponding editor. It also populates the attributes of the unmanaged address in the editor. Enter the appropriate information in the editor.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: After the conversion, the status of the unmanaged address changes to **Used**.

The following are some conditions for a conversion:

- **A and AAAA records:** You must select a DNS zone when converting an unmanaged address to an A or AAAA record.
- **PTR record:** You must select a DNS zone when converting an unmanaged address to a PTR record.
- **IPv4 and IPv6 Fixed Address:** Grid Manager displays a confirmation dialog box to ensure that you want to create a fixed address for the unmanaged address.
- **IPv4 and IPv6 Host record:** You can use the unmanaged address to enable a host record for DNS or DHCP.

Clearing Unmanaged Data

You can clear the status of unmanaged data. When you clear an unmanaged address, the status of the IP address changes to **Unused**. You can clear individual and multiple unmanaged addresses in the *IP Map* or *List* panel. You can also clear all the unmanaged data within a network in the *Net Map* and *List* panels.

You may also clear unmanaged data for any selected network or network view that has Unmanaged status.

To clear an unmanaged address:

1. In the *IP Map* or *List* panel, select an unmanaged address, and then click **Clear -> Clear Unmanaged Data** from the Toolbar.
or
In the *Net Map* or *List* panel, select a network or networks, and then click **Clear -> Clear All Unmanaged Data** from the Toolbar to clear all unmanaged addresses in the networks.
2. In the *Clear Unmanaged Data* confirmation dialog box, click **Yes**.

Note: When you clear unmanaged addresses in a given network view, all unmanaged IPv4 and IPv6 addresses of all networks in the network view are cleared. When you select an entire network or a specific network in the *Net Map* or *List* panel, all the unmanaged addresses in the network are cleared. After you clear the unmanaged data, the status of the IP addresses changes to **Unused**.

Resolving Conflicting Addresses

Conflicts happen when discovered data does not match the existing IP address data. The *IP Map* panel displays conflicting addresses in red. The *List* panel displays **Conflict** as the status for all conflicting addresses. Depending on the conflict, you can do one of the following to resolve it:

- For a DHCP lease conflict, you can clear the existing lease and create either a fixed address or a reservation for the IP address. You can also keep the existing data and clear the discovered data. For information, see [Resolving DHCP Lease Conflicts](#).
- For a fixed address conflict, you can either keep the existing fixed address data or update the existing data with the discovered data. For information, see [Resolving Fixed Address Conflicts](#) on page 514.
- For a DHCP range conflict, you can create a fixed address, create a reservation, or clear the discovered data. For information, see [Resolving DHCP Range Conflicts](#) on page 514.
- For a host conflict, you can either keep the existing host record data or update the existing data with the discovered data. For information, see [Resolving Host Conflicts](#) on page 515.

You must resolve conflicting addresses individually. You cannot resolve multiple conflicts at the same time.

Note: Once the conflict is resolved, the status of the IP address changes depending on how you resolved the conflict.

To resolve a conflict:

1. In the *IP Map* or *List* panel, select a conflicting address, and then click **Resolve Conflict** from the Toolbar.
2. The *Resolve Conflict* dialog box displays the reason of the conflict and lists the existing information and discovered information of the address in the **Description** field. Depending on the type of conflict, the appliance displays the corresponding resolution options. You can compare the existing and discovered data and decide how you want to resolve the conflict.

Resolving DHCP Lease Conflicts

When an IP address has a DHCP lease and the discovered MAC address is in conflict with the existing MAC address, the IP address has a DHCP lease conflict.

To resolve a DHCP lease conflict:

1. In the *Resolve Conflict* dialog box, select one of the following:
 - **Clear lease and create fixed address from discovered data:** Clears the existing DHCP lease and creates a fixed address with the discovered data. The *Fixed Address* editor appears with the discovered data populated.
 - **Clear lease and create a reservation from discovered data:** Clears the existing DHCP lease and creates a new reservation using the discovered data. The *Reservation* editor appears with the discovered data populated. This option does not apply to leases served by Microsoft DHCP servers because they do not support Infoblox reservations. For information about managing DHCP data served by Microsoft servers, see [Chapter 34, Managing Microsoft DHCP Services](#), on page 983.
 - **Keep the existing and ignore this conflict:** Keeps the current DHCP lease for the address and ignores the lease conflict.
2. Click **OK**.

Resolving Fixed Address Conflicts

When the discovered MAC address of an IPv4 address does not match its existing MAC address, or when the DUID of an IPv6 address does not match its existing DUID, the IP address has a fixed address conflict.

To resolve a fixed address conflict:

1. In the *Resolve Conflict* dialog box, select one of the following:
 - **Keep fixed address and clear discovered data:** Keeps the existing fixed address and clears the discovered data.
 - **Update fixed address with discovered data:** Updates the existing fixed address data with the discovered data.
2. Click **OK**.

Resolving DHCP Range Conflicts

When an IP address is in a DHCP range and does not match an existing DHCP lease, fixed address, or exclusion range and it shows an active state during a discovery, the IP address has a DHCP range conflict.

To resolve a DHCP range conflict:

1. In the *Resolve Conflict* dialog box, select one of the following:
 - **Create a fixed address:** Creates a fixed address with the discovered data.
If the fixed address is served by a Microsoft server, but is outside of a scope, you must then navigate to the *Fixed Address* editor and assign the fixed address to the appropriate Microsoft server.
 - **Create a reservation:** Creates a reservation with the discovered data. This creates an Infoblox reservation and therefore cannot be used for IP addresses served by Microsoft servers. Note that you cannot convert an IPv6 address to a reservation.
 - **Clear discovered data:** Clears the discovered data and no object is created for the IP address.
2. Click **OK**.

Resolving Host Conflicts

When the MAC address of an IPv4 address that belongs to a host record does not match its existing MAC address, or when the DUID of an IPv6 address that belongs to a host record does not match its existing DUID, the IP address has a host conflict.

To resolve a host conflict:

1. In the *Resolve Conflict* dialog box, select one of the following:
 - **Keep host record and clear discovered data:** Keeps the existing data and clears the discovered data.
 - **Update host record with discovered data:** Updates the existing host record data with the discovered data.
2. Click **OK**.

Clearing Discovered Data

You can clear discovered data on selected IPv4 or IPv6 networks. This is useful, for example, if the network topology has changed since the last time you ran a discovery on the network, and a new discovery needs to be run. You may perform this action whether or not the network is in a Managed or Unmanaged state.

To clear discovered data:

1. In the *Net Map* or *List* panel, select a network, and then click **Clear** -> **Clear Discovered Data** from the Toolbar.
2. In the *Clear Discovered Data* dialog box, click **Yes**.

You can also clear discovered data on all networks in a network view as follows:

1. In the *Net Map* or *List* panel, select a network, and then click **Clear** -> **Clear All Discovered Data** from the Toolbar.
2. In the *Clear All Discovered Data* dialog box, click **Yes**.

Note: When you clear all discovered data in a given network view, all imported discovered data for managed addresses in all IPv4 and IPv6 networks in the network view are cleared.



Chapter 14 Network Insight

This chapter provides information about the Network Insight device discovery feature, and how you can use discovery to collect, catalogue and manage information about network infrastructure devices and switched Ethernet network segments. Specific NIOS appliances are dedicated to the task of discovery, providing a device discovery extension layer for the NIOS Grid. You use the familiar NIOS interface to view the operating state of all discovered network infrastructure devices and newly discovered IP networks, including routers, firewalls, load balancers, Ethernet L2/L3 switches, end hosts and end host networks, and more. Discovery also provides cataloguing of all interfaces discovered for every device, and provides separate and highly specific information about them.

Discovery uses SNMP (Simple Network Management Protocol) along with other lower-level protocols to enable one or more dedicated NIOS appliances to become network management systems for monitoring, controlling and diagnosing network devices.

The NIOS IPAM feature set also provides control mechanisms for discovery, specifically for including and excluding networks and IP addresses for discovery. You can schedule and define when Discovery services take place on any network.

You can convert unmanaged Discovered networks and network assets to be managed in NIOS.

This chapter includes the following sections:

- [*About Network Insight*](#) on page 519
 - [*Consolidators and Probes*](#) on page 520
 - [*Administrative Permissions*](#) on page 521
- [*Supported Discovery Methods*](#) on page 522
 - [*SNMP*](#) on page 523
 - [*ICMP*](#) on page 523
 - [*TCP*](#) on page 523
 - [*Port Scanning*](#) on page 524
 - [*NetBIOS*](#) on page 524
- [*Starting and Stopping the Discovery Service*](#) on page 525
 - [*Changing the Discovery Member Type*](#) on page 526
 - [*Choosing a Probe Member Interface for Discovery*](#) on page 527
- [*Discovering IPs and Networks*](#) on page 527
 - [*Performing Discovery on an Existing Object*](#) on page 529
 - [*Smart Folders and Discovered Devices*](#) on page 529
- [*Configuring Grid Properties for Discovery*](#) on page 529
 - [*Activating DHCP Routers as Seed Routers*](#) on page 529
 - [*Configuring Grid SNMPv1/v2 Properties*](#) on page 530
 - [*Configuring Grid SNMPv3 Properties*](#) on page 530


- [*Defining Advanced Discovery Polling Techniques for the Grid*](#) on page 531
- [*Scheduling Switch Port Discovery and Data Collection*](#) on page 532
- [*Configuring Member Properties for Discovery*](#) on page 533
 - [*Defining Probe SNMPv1/v2 Properties*](#) on page 533
 - [*Configuring Probe SNMPv3 Properties*](#) on page 534
 - [*Defining Seed Routers*](#) on page 534
 - [*Choosing a Probe Member Interface for Discovery*](#) on page 527
- [*Excluding IP Addresses from Discovery*](#) on page 535
 - [*Disabling Discovery for a Network*](#) on page 537
 - [*Enabling IPs for Immediate Discovery or for Discovery Exclusion*](#) on page 535
 - [*Quick Exclusion of IPs from Discovery*](#) on page 535
 - [*Creating a New Fixed Address Object and Excluding it from Discovery*](#) on page 536
- [*Viewing the List of Discovered Devices*](#) on page 537
- [*Viewing Discovery Status*](#) on page 538
- [*Using Discovery Diagnostics*](#) on page 539
- [*Viewing Discovered Interface Information*](#) on page 540
 - [*Viewing the List of Networks Associated with a Discovered Device*](#) on page 540
 - [*Viewing the Management State of IPs in Discovered Networks*](#) on page 541
 - [*Viewing the List of IP Addresses Associated with a Discovered Device*](#) on page 541
 - [*Viewing the List of Assets Associated with a Discovered Interface*](#) on page 541
- [*Converting Unmanaged Networks to Managed Networks*](#) on page 542
- [*Adding Discovery Device Support*](#) on page 543
- [*Performing VM \(Virtual Machine\) Discovery*](#) on page 543
 - [*Executing a VM Discovery*](#) on page 543
 - [*Scheduling a VM Discovery Session*](#) on page 544

ABOUT NETWORK INSIGHT

Discovery enables you to know what you have in your networks down to individual ports, without making assumptions or performing guesswork. You can use discovery to catalogue the devices in any network even when you do not want devices to be directly managed by NIOS or to have their network identities provisioned by NIOS through DNS and DHCP.

All NIOS appliances that perform discovery require a separate Discovery license. (See [Starting and Stopping the Discovery Service](#) on page 525 for more information.) Appliances with this license operate only for discovery tasks and do not perform core DNS or DHCP functions. Discovery appliances, called Probes, collect network and device data. The Infoblox appliances performing discovery are all members of the Infoblox Grid, separately dedicated to the task of polling and discovery of networks and devices. A separate appliance, called a Consolidator, aggregates and organizes all collected device information from the Probes and synchronizes with the Infoblox Grid Master. (See [Consolidators and Probes](#) on page 520 for related information on Discovery appliances.) Consolidators also require a Discovery license.

Figure 14.1 IPAM-Discovery-licensed NIOS Appliances Added as new Discovery members to the Grid

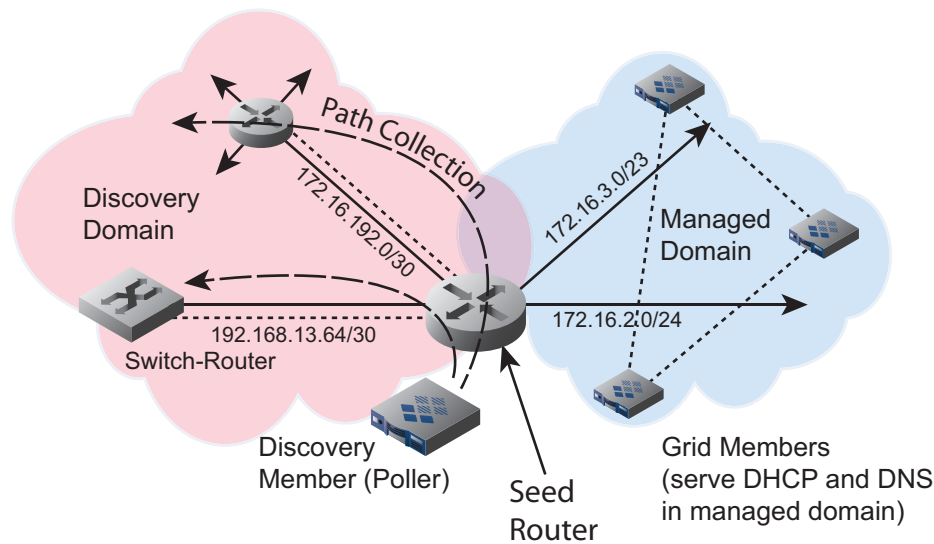


Appliances use SNMP and other protocols to discover and catalogue a diverse assortment of device types, including the following:

- Routers
- Enterprise Switches
- Firewalls and Security Appliances
- Load Balancers
- Enterprise Printers
- Wireless Access Points
- VoIP Concentrators
- Application Servers
- End Hosts

Discovery provides a powerful tool for administrators to gather key information about complex unmanaged networks, including discovery of routed paths and the host clouds behind enterprise switches, even in organizations where a partial NIOS deployment already exists. In [Figure 14.2](#), a NIOS appliance running discovery connects to an enterprise router, and uses its information to determine more about the networks that exist deeper within the unmanaged network, termed the Discovery Domain in this example:

Figure 14.2 Discovery in Action



As indicated in [Figure 14.2](#), discovery can trace through multiple hops and perform device discovery at every step, filling out the maps of unmanaged networks for the administrator. These unmanaged networks may have their own non-Infoblox services for DHCP and DNS. Discovery enables the administrator to begin the process of integrating devices in the Discovery Domain into the Managed Domain.

The collection of unmanaged network information extends to the networks of distribution Ethernet switches. Data collection also includes end hosts and application/file servers connected to edge switches in enterprise offices. Discovery terms these devices **Assets**, and separately catalogues them in detail. See [Viewing the List of Assets Associated with a Discovered Interface](#) on page 541 for related information.

The Probes forward all data to the Consolidator, which subsequently synchronizes device data for cataloguing with the Grid Master. Once information about discovered networks and devices resides on the Grid Master, you can convert unmanaged networks to managed networks, adding them to the NIOS database.

You provide one or more routers as seed routers to act as the initial gateways for discovering other networks and their devices in the discovery domain (an example appears in [Figure 14.2](#)). You can also use DHCP routers (e.g., routers serving DHCP leases) as seed routers to aid in faster discovery.

You can also create new networks and DHCP ranges and activate Discovery upon them. Once you create the network, discovery will locate, poll and catalogue the network devices comprising the networks. This information is then synchronized with the Grid Master. See [Discovering IPs and Networks](#) on page 527 for more information.

Consolidators and Probes

Infoblox appliances fall into two distinct classes for Discovery operations:

Consolidator—The central repository of discovery data for the entire managed network. The Consolidator is a single appliance that contains all data about all devices detected through discovery. The Consolidator communicates with the Grid Master as a normal Grid Member and transfers all its data to the Grid Master, as indicated in [Figure 14.1](#). Consolidators compile their information from one or more associated Probe appliances. The Consolidator appliance requires a Discovery license. If you have one or more Probe appliances or virtual appliances, the Consolidator performs no discovery on its own. If you plan to use a single dedicated appliance for Discovery, that appliance must be licensed for Discovery and be configured as a Consolidator. A Grid Master cannot be licensed as a Consolidator. For related licensing configuration, see [Starting and Stopping the Discovery Service](#) on page 525.

Note: Any Infoblox Grid supports a single Consolidator appliance.

Probes—A Probe is an Infoblox appliance or virtual appliance that performs the direct querying, probing and polling of network devices and the initial data collection. Probe appliances also require the Discovery license. Infoblox recommends using one or more Probe appliances with any Consolidator. Each Probe can override the Grid level discovery credentials with its own discovery credentials.

Data synchronization occurs continuously between the Consolidator and all associated Probe appliances and between the Consolidator and the Grid Master.

Note: Every Probe appliance can be assigned to a single network view. Network view assignments can be changed at any time. Multiple Probe appliances can be assigned to the same network view.

Consolidator-Probe Appliance—You may also choose to operate a Consolidator-Probe NIOS appliance as a single Discovery system. In this deployment, the appliance operates as both a Consolidator and a Probe, performs all discovery operations, aggregates all databases within it, and synchronizes with the Grid Master. As noted, a Standalone appliance is defined as a Consolidator, but also performs the operations of a Probe.

Keys to Discovery Operation

You provide several pieces of information to allow discovery appliances to do their work:

- One or more seed routers for polling to discover networks along routed paths. See [Defining Seed Routers](#) on page 534.
- One or more networks to be discovered.
- SNMP credentials (including SNMPv3 authentication/encryption credentials) for network devices. See [SNMP](#) on page 523 and [Configuring Grid SNMPv1/v2 Properties](#) on page 530 for more information.
- Exclusions. You may need to exclude IPs from discovery for various reasons. See [Exclusion from Discovery](#) on page 521.

Standalone appliances cannot be installed in a network that already has existing Probes and a Consolidator. For more information, see [Changing the Discovery Member Type](#) and [Consolidator and Probe Appliance Deployment Guidelines](#) on page 527.

For related information, see the topic [Starting and Stopping the Discovery Service](#) on page 525.

Exclusion from Discovery

You can optionally exclude IP addresses from discovery. The basic principle is that some devices do not need to be discovered, perhaps because they are already managed as part of a NIOS Grid and hence should not be subjected to discovery; because a device does not support SNMP; or for other organizational reasons. In [Figure 14.2](#) on page 520, networks 172.16.2.0/24 and 172.16.3.0/23 are excluded from discovery because they are already fully managed by a NIOS Grid.

For more information, see the various topics under [Excluding IP Addresses from Discovery](#) on page 535.

Administrative Permissions

You can initiate a discovery and manage discovered data based on your administrative permissions. For information, see [About Administrative Permissions](#) on page 160.

Initiating and controlling a discovery requires specific administrative permissions. The following are permission guidelines for initiating and controlling a discovery:

- The **IPAM Discovery Admin** role provides a pre-configured list of permissions by which assigned admin accounts may perform discovery tasks. Administrators with these permissions can initiate and control discovery on any existing network. The **IPAM Discovery Admin** role supports the following permissions:
 - All permissions associated with the Network Discovery feature set (active if you do not have a Network Insight license)
 - Read-Only on all Network Views, network containers, networks and ranges
 - Read-Only on all Hosts
 - Read-Only on all Members

- Read-Write Network Discovery permissions
- Editing network, network container or range discovery properties: Read-Only for each type. For member assignment, the user also needs additional read-only permission for the assigned member
- Editing fixed address, host or reservation discovery properties: Read-only for each type
- Excluding an IP address or an IP Range (from the Network Editor ‘s “Discovery Exclusions” tab or from the IPAM IP List view): read-only permission for the network
- Discover Now for Network, DHCP Range or IP: Read-Only permission for each.
- If the user does not possess the Network Discovery permission, Network Insight permissions are disabled.
- Superusers can initiate and control discovery on all networks. Some discovery functions require superuser permissions:
 - Grid Discovery properties
 - Uploading, Viewing and deletion of device support bundles
 - Launching Discovery Diagnostics
 - Launching Discovery Status

Similar to Network Discovery, devices and end hosts discovered through discovery can undergo conversion from unmanaged status to managed status. This entails converting an unmanaged IP address to a Host, an A record or AAAA record, a PTR record, or to a fixed address.

After a discovery is complete, the following permission guidelines apply to viewing and managing discovery data:

- Superusers can view and manage all discovered data.
- IPAM Discovery admins can convert unmanaged networks to managed networks and can change discovery settings for networks.
- Administrators with read permission to networks can view all discovery data without editing.
- Administrators with read/write permission to a DNS zone or specific record type can convert unmanaged data to a host, fixed address, reservation, A record, or PTR record.
- If a user has read-only permissions for a device’s management IP address, the device will be visible in the **Data Management → Devices** tab.
- For unmanaged networks: users may Delete, Convert to Managed, Clear Unmanaged and Clear Discovered Data if one of the following is true:
 - User has read-write permission for the network –or–
 - User has Network discovery permission plus Read-Only for the network.

SUPPORTED DISCOVERY METHODS

When you perform a discovery, you can choose any or all of the following discovery methods:

- SNMPv1/v2c device polling as described in [SNMP](#) on page 523.
- SNMPv3 device polling as described in [SNMP](#) on page 523.
- ICMP Ping Sweep and Smart Subnet Ping Sweep as described in [ICMP](#) on page 523.
- TCP as described in [TCP](#) on page 523.
- NetBIOS as described in [NetBIOS](#) on page 524.

These methods actively scan predefined networks and probe IP addresses. The appliance listens for responses from the IP addresses as proof of activity. The IP discovery scans through the specified network ranges and probes IP addresses (except for the network, broadcast, and multicast address types) in each network, including the /31 and /32 subnets. Note that addresses in the /31 and /32 subnets can be used only as source addresses for point-to-point links. In these cases, no broadcast or network addresses exist in the /31 and /32 subnets, and the appliance can discover source addresses in these subnets.

SNMP

Note: Infoblox does not recommend using vendor default SNMP credentials on network devices. Should you need to use vendor defaults for a given device type, you must enter those values in the list of SNMP credentials on the Grid Master.

Discovery supports discovery of devices and networks through SNMPv1/v2c and through SNMPv3 protocols. Discovery acquires information from standard SNMP MIB Object IDs (OIDs) to correctly identify and catalogue devices. You enter or import lists of SNMP credentials with which the appliances query devices on the network to perform discovery.

SNMPv1 and SNMPv2c protocols are combined into a set termed SNMPv1/v2 for Discovery. SNMPv1/v2 discovery requires standard read community strings as entries on the Grid Master.

Accounts using SNMPv3 use a standard suite of authentication and security protocols. If NIOS will use SNMPv3 to collect data from devices supporting the protocol, you can define specific user credentials with combinations of authentication and protocol support, and the unique keys for each protocol. NIOS also supports multiple entries for the same username string, enabling checking of similar SNMPv3 credentials that use different authentication and security protocols.

Some devices found by discovery may not have SNMP credentials that are known or entered into the SNMP credentials sets.

Note: SNMP Credentials from the Grid or from the Member credential list are always tried in the specified order unless a credential is associated with a host, fixed address or reservation being discovered.

ICMP

Discovery uses different variations of Ping traces to perform higher-performance, brute-force device discovery. ICMP is the last resort when devices do not support SNMP management protocols or an SNMP credential is lacking.

The ICMP Smart Ping Sweep option enables brute-force subnet Ping sweeps on IPv4 networks. Subnet ping sweeps are used as a last resort in the discovery process. A subnet ping sweep is performed if NIOS is unable to identify any network devices in a given subnet. Subnet ping sweeps are performed no more than once per day, and will end the ping sweep on a given subnet once NIOS discovers a network device and is able to collect data from it. You can configure the timeout value (Ping Sweep Timeout) and the number of attempts (Ping Sweep Attempts).

Note: Smart subnet ping sweeps will not be performed on subnets larger than /22. Ping sweeps of any kind do not apply on IPv6 networks because of the dramatically greater scale of network addresses in the IPv6 realm.

Complete Ping Sweep differs from the Smart Subnet ping sweep in the following ways:

- The discovery ping sweep will run only against the specified range.
- The sweep will run regardless of the range size.
- The sweep will run regardless of the number of discovered devices within the specified range.

Discovery also performs automatic Ping traceroutes when needed for path collection. Path collections run without user intervention or configuration.

TCP

TCP scanning probes each active host on a list of TCP ports using TCP SYN packets. This method detects all active hosts that generate SYN ACK responses to at least one TCP SYN. The discovery can determine the OS on a host by analyzing how the host reacts to the requests on opened and closed ports. It then uses the TCP fingerprints to guess the OS. To obtain a TCP fingerprint, IP discovery provides two scanning techniques, SYN and CONNECT.

When you use the SYN technique, the discovery sends a TCP SYN packet to establish a connection on a TCP port. If the port is open, the host replies with a SYN ACK response. The discovery does not close the port connection.

The CONNECT technique is a three-way TCP handshake. The discovery starts with the same process as the SYN technique by sending the TCP SYN packet. A response containing a RST flag indicates that the port is closed. If the host replies with a SYN ACK response, discovery sends a RST packet to close the connection. If there is no reply, the port is considered filtered. TCP scanning is a deliberate and accurate discovery method, enabling detection of all active hosts on a network provided that there are no firewalls blocking TCP packet exchanges.

You can choose the TCP ports and the TCP scanning technique in the *Discovery Manager* wizard. This method returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: The discovery returns the MAC address only if the Probe member running the discovery is on the same discovered network.
- OS: This is set to the highest probable OS reported in the response.

To use the TCP discovery method, the TCP port and a specific set of ports between the Probe member and the discovered networks must be unfiltered. The default set of ports is defined by the factory settings.

Port Scanning

By enabling port scanning, NIOS probes the list of TCP ports enabled in the **Advanced** tab, to determine whether they are open. You can control some settings for port scanning behavior, including the choice of a TCP scanning technique.

- **Profile Device:** If enabled, NIOS attempts to identify the network device based on the response characteristics of its TCP stack, and uses this information to determine the device type. In the absence of SNMP access, the Profile Device function is usually the only way to identify devices that do not support SNMP. If you disable Profile Device, devices accessible via SNMP will still correctly identify; all other devices are assigned a device type of Unknown. Profile Device is disabled by default for discovery polling.

The Profile Device option uses the editable list of TCP protocol ports from the **Grid Discovery Properties → Polling → Advanced** tab as its profile, and polls each of the ports enabled in that list, using the configured timeout value and the number of polling attempts for each port.

See [Defining Advanced Discovery Polling Techniques for the Grid](#) on page 531 for more information.

Should you disable Port Scanning, discovery attempts no port probes other than SNMP on any device.

NetBIOS

The NetBIOS method queries IP addresses for an existing NetBIOS service. This method detects active hosts by sending NetBIOS queries and listening for NetBIOS replies. It is a fast discovery that focuses on Microsoft hosts or non-Microsoft hosts that run NetBIOS services.

NetBIOS discovery returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: Listed only if the discovered host is running Microsoft, otherwise blank.
- OS: This value is set to **Microsoft** for an active host that has a MAC address in the NetBIOS reply.
- NetBIOS name: This value is set to the name returned in the NetBIOS reply.

To use the NetBIOS discovery method, ports 137 (UDP/TCP) and 139 (UDP/TCP) between the Grid member performing the discovery and the target networks must be unfiltered.

The following table summarizes the supported discovery methods:

Discovery Type	Returned Data	Guideline	Mechanism
Smart IPv4 Subnet Ping Sweep	<ul style="list-style-type: none"> • IP address • MAC address 	Apply on known subnetworks on which no devices are readily found. Limited to networks of /22 and smaller.	ICMP echo request and reply.

Complete Ping Sweep	<ul style="list-style-type: none"> • IP address • MAC address 	Last resort for discovery. Use ICMP for a rough and fast discovery. Enables path tracing.	ICMP echo request and reply, ICMP traceroute.
NetBIOS	<ul style="list-style-type: none"> • IP address • MAC address • OS • NetBIOS name 	Use NetBIOS for discovering Microsoft networks or non-Microsoft networks that run some NetBIOS services	NetBIOS query and reply.
TCP	<ul style="list-style-type: none"> • IP address • MAC address • OS 	Use TCP for an accurate but slow discovery	TCP SYN packet and SYN ACK packet.
Port Scanning/ Profile Device	<ul style="list-style-type: none"> • Open and Closed TCP ports • IP Address 	Disabled by default, use for non-SNMP devices.	Scans specified list of TCP ports, using TCP SYN packet.
SNMPv1/v2 SNMPv3	<ul style="list-style-type: none"> • Open and Closed TCP ports • IP Address • System Description • System Up Time • Routing Neighbors • ARP tables • SNMP credentials 	Most important protocols for discovery. Ensure you have the SNMP credentials necessary for probing devices using SNMP.	Queries and collects system OIDs such as SysDescr and sysUpTime.
VM discovery	<ul style="list-style-type: none"> • IP address • MAC address • OS • Discovered name • Virtual entity type • Virtual entity name • Virtual cluster • Virtual datacenter • Virtual switch • Virtual host • Virtual host adapter 	Add the VMware vSphere servers on which you want to perform the VM discovery	The appliance communicates with the vSphere servers to collect discovery data on virtual machine instances

STARTING AND STOPPING THE DISCOVERY SERVICE

Note: The Discovery service can only be started on Grid members of Consolidator and Probe appliance types.

Each discovery member requires separate Discovery licensing, and must have a running Discovery service. To start or stop the Discovery service, you select each appliance under the Grid Manager page. Some items to keep in mind include:

- A Grid Master does not run the Discovery service.
- Appliances running a Discovery license and the Discovery service, do not support HA pairs.
- All appliances running Discovery must have the Discovery license installed first before starting the service.

- Appliances running Discovery do not run core network services such as DNS and DHCP. Discovery appliances may also run the NTP service.
- If you expect to run a single appliance in the Grid for discovery, the appliance is designated as a Consolidator, and will also perform Probe discovery operations.
- When you add a new Grid member with a Discovery license, the appliance is set automatically to the following:
 - A Consolidator, if no other member with a Discovery license exists in the Grid.
 - A Probe, when at least one appliance exists in the Grid

Note: When a Member joins the Grid and applies a Discovery license for the first time, the admin user will need to log off and log in again to the NIOS UI to see the Discovery-enabled functionality.

For information about discovery configuration at the service level, see the topic [Configuring Grid Properties for Discovery](#) on page 529.

To start the Discovery service on a correctly licensed Consolidator or Probe appliance:

1. From the **Grid** tab, select the **Grid Manager** tab, and click the **Services** tab.
2. Click the **Discovery** icon to display the list of members running the Discovery service.
3. Select the Discovery member or members for which you wish to start the service.
4. Open the Toolbar and click the Start button.
NIOS asks you to verify that you want to proceed with starting the service for the selected member.
5. Click **Yes**.

To stop the Discovery service on a Consolidator or Probe appliance:

1. From the **Grid** tab, select the **Grid Manager** tab, and click the **Services** tab.
2. Click the **Discovery** icon to display the list of members running the Discovery service.
3. Select the Discovery member or members for which you wish to stop the service.
4. Open the Toolbar and click the Stop button.
NIOS asks you to verify that you want to proceed with stopping the service for the selected member.
5. Click **Yes**.

Changing the Discovery Member Type

On occasion, you may wish to change an appliance with a Discovery license to a Consolidator, or to a Probe, or change a Consolidator to a standalone discovery appliance. To do so, you must first turn off the Discovery service on the appliance (for information, see [Starting and Stopping the Discovery Service](#) on page 525, and the following section [Consolidator and Probe Appliance Deployment Guidelines](#) for information on specific cases).

You can also change the network view for any Probe appliance, and choose a different interface through which the appliance will perform discovery.

To change the Discovery appliance type:

1. From the **Grid** tab, select the **Grid Manager** tab, and click the **Services** tab.
2. Click the **Discovery** icon to display the list of members running the Discovery service.
3. Select the Discovery member for which you wish to change the appliance type.
4. Expand the Toolbar and click the **Edit → Member Discovery Properties** button.
5. In the **General** tab of the *Member Discovery Properties* editor, choose the **Member Type**:
 - **Polling**: turns the appliance into a discovery Probe appliance.
 - **Consolidator**: turns the appliance into a discovery Consolidator appliance.
 - **Unassigned**: disables the Discovery features in the appliance.

6. If you are changing an appliance into a Probe, choose the **Network View**. If only a single network view is present, the Network View section is not displayed.
7. Click **Save & Close**. If you wish to change the interface over which the appliance sends and receives discovery traffic, see [Choosing a Probe Member Interface for Discovery](#) on page 527.

Consolidator and Probe Appliance Deployment Guidelines

When you wish to install and deploy Discovery appliances, use the following installation guidelines:

Installing a Standalone in the Grid—Before you designate an appliance as a Standalone discovery appliance, no previously installed Probes should be present on the network and joined to the NIOS Grid. If you install a new appliance intended as a Standalone, in a network that already has one or more Probe instances (perhaps for testing or evaluation purposes), before Discovery service is stopped on the Probe instances, the new “Standalone” appliance will automatically detect the Probe instances and start as a Consolidator appliance, preventing it from acting to probe and detect devices as a standalone appliance. Consolidators cannot be assigned to network views or to discovery in network objects such as IPv4 or IPv6 network containers.

Converting a Consolidator to a Standalone—Also consider the example of a Consolidator appliance operating with one or more instances running as Probes, each with respective Discovery licenses. If you wish to convert the Consolidator to a standalone Discovery appliance, stop the Discovery service on all associated probes. Then, stop and restart the Discovery service on the Consolidator appliance. The appliance then will be selectable for discovery of network objects, acting as a Standalone Discovery appliance.

Adding new Probe Instances to a Standalone deployment—Finally, consider the use of a Standalone Discovery appliance to which you wish to associate a new Probe instance or instances. This process converts a Standalone to a Consolidator. After the new Probe instances join the NIOS Grid, stop the Discovery service on the standalone Discovery appliance. Then, start the Discovery service on the new Probe or Probes. Next, restart the Discovery service on the previously defined Standalone appliance. It will detect the newly active Probe instances and activate as a Consolidator.

In all cases, you must maintain proper IPAM Discovery licensing.

Choosing a Probe Member Interface for Discovery

Probe members must have a designated interface over which all Discovery traffic exchanges take place. You may designate one interface for discovery on each Probe appliance. The LAN1, LAN2 or MGMT ports may be specified. (These ports must be defined in the member network settings before they can be used for discovery.)

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
2. Select a Probe appliance from the Discovery page.
If you do not select an appliance from the list, the Edit Member Discovery Properties option is disabled.
3. Expand the Toolbar and click **Edit → Member Discovery Properties**.
4. Click the **General** side tab.
5. Choose an interface from the Discovery Interface list.
6. Click **Save & Close**.

DISCOVERING IPs AND NETWORKS

To start discovery on connected networks:

1. Ensure that your appliances are licensed for discovery.
2. Add the needed seed routers to each Probe appliance (see [Defining Seed Routers](#) on page 534).

3. Add the necessary SNMPv1/v2 and SNMPv3 credentials at the Grid level or Member/Probe level (see [Configuring Grid SNMPv1/v2 Properties](#) and [Configuring Grid SNMPv3 Properties](#) on page 530, and [Defining Probe SNMPv1/v2 Properties](#) on page 533 and [Configuring Probe SNMPv3 Properties](#) on page 534).
4. If necessary, enable the use of DHCP routers and servers as seeds to increase device discovery ([Activating DHCP Routers as Seed Routers](#) on page 529).
5. If you have extensive end host networks connected to Ethernet switches, enable switch port discovery (see [Scheduling Switch Port Discovery and Data Collection](#) on page 532).

With these settings, the Probe appliances will automatically begin discovering network infrastructure devices.

You can elect to immediately discover new Objects that you create and enable in NIOS IPAM. Objects that allow immediate discovery include the following:

- **IPv4 Fixed Address** (see [Configuring IPv4 Fixed Addresses](#) on page 857 for the complete procedure).
You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the IPv4 fixed address, and override the SNMP credentials if necessary.
- **IPv6 Fixed Address** (see [Configuring IPv6 Fixed Addresses](#) on page 878 for the complete procedure).
You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the fixed address, and override the SNMP credentials if necessary.
- **IPv4 Reservation** (see [Configuring IPv4 Reservations](#) on page 860 for the complete procedure).
You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the IPv4 reservation, and override the SNMP credentials if necessary.
- **Host** (see [Adding Host Records](#) on page 462 for the complete procedure).
You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the host, and override the SNMP credentials if necessary.
- **IPv4 Network** (see [Configuring IPv4 Networks](#) on page 845 for the complete procedure).
You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new network.
- **IPv6 Network** (see [Configuring IPv6 Networks](#) on page 870 for the complete procedure).
You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new network.
- **IPv4 DHCP Range** (see [Configuring IPv4 Address Ranges](#) on page 854 for the complete procedure).
You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new IPv4 DHCP range.
- **IPv6 DHCP Range** (see [Configuring IPv6 Address Ranges](#) on page 876 for the complete procedure).
You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new IPv6 address range.

During configuration, you can choose to **Exclude from Network Discovery** if you wish to postpone discovery for specific object types.

Note: Individual IP addresses within a network, and specific Object types (IPv4 Reservation, Fixed IP Address, and Host), may be Excluded from discovery. You must explicitly **Enable Discovery** for other object types (IPv4 and IPv6 Ranges, IPv4 and IPv6 Networks) and optionally can **Enable Immediate Discovery**.

If you choose not to perform immediate discovery, but do **Enable Discovery**, the new network or other object will be discovered at a normal time determined by NIOS.

You can manually perform discovery on any object at any time by selecting the object and choosing **Discover Now** from the Toolbar. See the section [Performing Discovery on an Existing Object](#) for more information.

By default, Grid Discovery settings are the prevailing settings for all newly created Objects. You can override basic Discovery polling options for networks and DHCP ranges allowing immediate discovery. In such cases local settings will take priority. Credentials cannot be overridden for networks and DHCP ranges,

Performing Discovery on an Existing Object

Note: If a network or IP does not enable the **Discover Now** button after selecting it, make sure the network or other Object has a Discovery Probe member assigned to it.

After you create any of the eight objects listed in the previous section, you may wish to perform discovery on an object at a later time.

At a later time, you can simply select the object and discover it.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
2. Select the network or other object over which you want to perform discovery.
Depending on the object type, you will need to navigate from the network level to the individual IP table in the **List** page to locate the object for immediate discovery.
3. Expand the Toolbar and click the **Discover Now** button.

You can also click the Action icon for the network and choose **Discover Now** from the menu.

The Probe member associated with the network or other object initiates a Discovery procedure.

Smart Folders and Discovered Devices

NIOS maintains a Smart Folder entitled Discovered Switches/Routers, under which is a list of all routers, switches and switch-routers that thus far have been discovered and catalogued through NIOS' discovery feature. Simply open the Smart Folders category under the Finder menu and click on the **Discovered Switches/Routers** folder. Clicking on a device name opens the device page under **Data Management → Devices** and shows the **Interfaces** page for the chosen device. For related information, go to [Predefined Smart Folders](#) on page 142.

A similar folder, **Unmanaged**, provides a list of unmanaged discovered networks. Clicking on any listed network displays the **Data Management → Devices → IP Map** page with its graphical map of IP addresses.

CONFIGURING GRID PROPERTIES FOR DISCOVERY

Note: Grid Discovery Properties operations require superuser permissions under NIOS.

Each individual appliance can override credential settings for its own SNMP operations during discovery. Basic discovery polling settings may be overridden for DHCP ranges or IPv4/IPv6 Network objects.

Some settings, such as seed router definition, take place only on Probe appliances.

Activating DHCP Routers as Seed Routers

Default gateways for DHCP Ranges defined in NIOS and for networks served by DHCP, can be leveraged as seed routers for faster discovery. Enabling this option allows discovery to employ all currently defined DHCP-configured devices in the discovery domain.

Unlike defining normal seed routers, which is done for each Discovery member appliance, you enable DHCP Seed router support at the Grid level.

With this setting enabled, the default gateways for any DHCP ranges and networks (that also have Discovery enabled) will automatically be leveraged for discovery.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
If you do not select an appliance from the list, the Grid Discovery Properties option is enabled.
2. Expand the Toolbar and click **Edit → Grid Discovery Properties**.
3. Click the **Advanced** tab. The Advanced Polling settings page appears.

4. Check the **Use DHCP Routers as Seed Routers** checkbox. The Probe members can then use the default gateways for associated DHCP ranges and networks as seed routers to more quickly discover and catalogue all devices (such as endpoint hosts, printers and other devices). All such default gateways will automatically be leveraged by discovery, and no further configuration is necessary unless you wish to exclude a device from usage.
5. Click **Save & Close**.

Note: Check for a list of configured DHCP seed routers for any discovery Probe member in the Member Discovery Properties page, **Seed** → **Advanced** tab.

Configuring Grid SNMPv1/v2 Properties

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
If you do not select an appliance from the list, the Grid Discovery Properties option is enabled.
2. Expand the Toolbar and click **Edit** → **Grid Discovery Properties**.
3. Click the **Credentials** side tab. The **SNMPv1/v2** page appears.
Add read community strings for polling SNMPv1/v2 devices on this page. Probe members can inherit all SNMPv1/v2 credential information from the Consolidator, or can elect to Override, to add SNMPv1/v2 credential records specific to Discovery operations on the Probe appliance.
4. Click the Add icon to add a new community string entry to the list. Enter a read community text string that the management system will send together with its queries to the network device during discovery.
A community string is similar to a password in that the discovered device accepts queries only from management systems that send the correct community string. Note that this community string must exactly match the value that is entered in the managed system.
5. If you have a substantial list of community strings in this list and need to find a specific string, enter the value in the **Go To** field and click **Go**.
6. To remove a community string entry: check the check box and click the Delete icon.
7. To export the entire list of community strings in a table file readable by a spreadsheet program, click the Export icon and choose **Export Data in Infoblox CSV Import Format**.
 - a. To export all data in a different format, click the Export icon and choose **Export Visible Data**.

Configuring Grid SNMPv3 Properties

SNMPv3 allows the use of two secret keys for every credential—one for authentication, and another for encryption. NIOS allows flexible application of keys—authentication but no encryption, for example. You define users in one of three ways:

- SNMPv3 user, with no authentication or privacy credentials
- SNMPv3 user, with authentication but no privacy credentials
- SNMPv3 user, with both authentication and privacy credentials

You can import sets of SNMPv3 credentials from an Infoblox CSV Import format data file.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
If you do not select an appliance from the list, the Grid Discovery Properties option remains enabled.
2. Expand the Toolbar and click **Edit** → **Grid Discovery Properties**.
3. Click the **Credentials** side tab. The **SNMPv1/v2** page displays automatically.
4. Click the **SNMPv3** tab. The page automatically displays the set of SNMPv3 credentials defined at the Grid level.

Note: Any SNMPv3 credentials you enter here will automatically be applied to all Probe appliances. They can be overridden on each Probe appliance.

5. Click the Add icon to add a new SNMPv3 authentication entry to the list. Enter the **Name** for the new credential; followed by the **Auth Protocol**, **Auth Password**, **Privacy Protocol**, **Privacy Password**, and the **Order** value, which is the order used for attempting use of the SNMP credentials. You can press Tab to navigate across the fields for the credential entry.
6. If you have a substantial list of SNMPv3 entries and need to find a specific entry, enter the value in the **Go To** field and click **Go**.
7. To remove an SNMPv3 authentication entry: check the check box and click the Delete icon.
8. To export the entire list of community strings in a table file readable by a spreadsheet program, click the Export icon and choose **Export Data in Infoblox CSV Import Format**.
 - a. To export just the subset of data that is visible in the dialog, click the Export icon and choose **Export Visible Data**.
 - b. A **Show Passwords** option allows the secret keys to be visible in the import.

Defining Advanced Discovery Polling Techniques for the Grid

You may need to work with Advanced Polling settings to ensure full discovery of devices on your chosen networks. Advanced SNMP polling settings consist of choosing the TCP Scan Technique, along with a number of specialized settings for Ping Sweeps and other operations.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
If you do not select an appliance from the list, the Grid Discovery Properties option is enabled.
2. Expand the Toolbar and click **Edit → Grid Discovery Properties**.
3. Click the **Advanced** tab. Advanced Polling settings include the following:
 - **TCP Scan Technique:** Two options SYN and CONNECT.
 - **SYN** (the default) quickly performs scans on thousands of TCP ports per system, never completing connections across any well-known port. SYN packets are sent and the Probe waits for a response while continuing to scan other ports. A SYN/ACK response indicates the protocol port is listening while a RST indicates it is not listening. The SYN option presents less impact on the network.
 - Use the **CONNECT** option for scanning IPv6 networks. Unlike the SYN option, complete connections are attempted on the scanned system and each successive TCP protocol port being scanned.
 - Enable one, selected, or all service protocols, listed by port number, service name and transport protocol (TCP). Click the top of the check box column to enable all protocols for Discovery checking.
 - **Purge expired assets after:** Remove records of discovered end hosts that are no longer reachable after a specified period of time. The default is set to one day.
 - **Purge expired device data after:** Remove records of discovered network infrastructure devices that are no longer reachable after a specified period of time. The default is set to seven days, a more forgiving value given that devices sometimes require maintenance, upgrades or repairs, or in cases where hosts leave the network on long trips.
 - **ARP Aggregate Limit:** Determines the largest ARP table collectible by NIOS Discovery. The default is set to 30 ARP table entries.
 - **Route Limit:** Limits the size of the routing table that NIOS Discovery will be required to collect from any given device. Some routers can have tables in the hundreds of thousands of entries, and collecting such a large body of data can impose performance problems in the network and in discovery data collection. This setting defaults to 3000, and automatically excludes BGP routes from collection. Consult Infoblox Technical Support before making changes to this value.
 - **ARP Cache Refresh:** Defines the time period between ARP refreshes by NIOS across all switch ports. Before any other switchport polling and discovery operations take place (including any global Discovery polling operations initiated by the NIOS administrator), another ARP refresh is carried out by the Probe appliance regardless of the time interval. The default is five minutes, because switch forwarding tables are frequently purged from LAN switching devices. (The default on Cisco switches is five minutes/300 seconds.) NIOS primarily uses ARP Cache refreshes to improve the accuracy of end-device discovery. Without this feature, some endpoints may not be discovered and cataloged.

- **Ping Sweep Frequency:** Defaults to 1, because Ping Sweep should not be executed more than once a day when the feature is enabled at the Grid level or for a given Discovery range. This setting affects the **Smart Ping Sweep** and **Complete Ping Sweep** features under Grid Discovery Properties.
- **Disable discovery of networks not in IPAM:** Enabling this setting disallows NIOS Discovery from executing discovery on any infrastructure networks that are not presented in the Infoblox IPAM system; e.g. present and managed in a network view or network container.
- **Authenticate and poll using SNMPv2c or later only:** For credential discovery and device polling exclusively using SNMPv2c and SNMPv3, excluding SNMPv1, enable this check box.
- **Use DHCP Routers as Seed Routers:** Enabling this setting, the Consolidator directs all associated Members to use the default gateway IPs for any DHCP ranges or DHCP networks as seed routers to more quickly Discover and catalogue all devices (such as endpoint hosts, printers and other devices). See [Activating DHCP Routers as Seed Routers](#) for more information.
- **Log IP discovery events in Syslog:** If enabled, directs all discovery members in the NIOS Grid to log IP discovery events (any discovered IPv4 or IPv6 device) to a configured Syslog server.
- **Log network discovery events in Syslog:** If enabled, directs all discovery members in the NIOS Grid to log network discovery events (discovery of devices in an IPv4 or IPv6 network) to a configured Syslog server.

4. After completing your settings, click **Save & Close**.

Scheduling Switch Port Discovery and Data Collection

Discovery schedules are Grid-wide and apply across all Probe members unless overridden at the network or DHCP range level. Discovery treats Ethernet switch port data collection as a separate task that can be separately scheduled. You may also enable periodic polling of switch ports, in which case all enabled polling settings automatically execute over defined time periods.

Substantial data may be collected from devices on switched Ethernet segments; you may schedule discovery for these networks for off-peak hours.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.

If you do not select an appliance from the list, the **Grid Discovery Properties** option remains enabled.

2. Expand the Toolbar and click **Edit → Grid Discovery Properties**.

3. Select the **Polling → Basic** tab.

4. Enable the **Switch Port Data Collection** checkbox.

You schedule Discovery execution on discovered Ethernet switch ports based on the following settings:

- a. Select how often you want to execute Discovery. You can choose **Once**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

When you select **Once**, enter the day in the **Day__of__** field and select a month from the drop-down list.

- Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
- Choose the **Time Zone**.

When you select **Hourly**, complete the following:

- **Schedule every hour(s) at:** Enter the hours between each Discovery instance to be run, from 1 to 24.
- **Minutes past the hour:** Enter the minute after the hour when the Consolidator member (not the Grid Master) executes the Discovery through its associated Probe members. For example, enter 5 if you want discovery to start on the managed networks five minutes after the hour.
- Choose the **Time Zone**.

When you select **Daily**, click either **Every Day** or **Every Weekday**.

- Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
- Choose the **Time Zone**.

When you select **Weekly**, complete the following:

- **Schedule every week on:** Select the check box for any day of the week.
- **Time:** Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list.

- Choose the **Time Zone**.

When you select **Monthly**, complete the following:

- **Schedule the day of the month:** A Discovery can be executed monthly on a specific day, or instances can be executed more than one month apart on a specific day, in the **Day__every__month(s)** field.
- **Time:** Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list.
- Choose the **Time Zone**.

5. Save your settings.

CONFIGURING MEMBER PROPERTIES FOR DISCOVERY

Member-level Discovery settings can be used to override Grid-level settings for that member. By default, all Probe appliances use the Grid-level Discovery settings. Probe member Discovery settings consist of the following:

- SNMPv1/v2 community strings.
- SNMPv3 credentials.
- Choosing a Discovery Interface.
- Choosing a Discovery member type.
- Defining Seed Routers (defined only at the Member level).

After overriding and creating local credentials on the Probe appliance, you may choose to restore the global settings by changing the setting back to **Inherit**. Doing so restores the inherited settings (NIOS also retains the previously entered local credentials, keeping them available for later use by another **Override** action).

Defining Probe SNMPv1/v2 Properties

You can define SNMP credentials on each Probe appliance. By default, each Probe inherits its community strings and credentials lists from the Consolidator.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.

2. Check the check box for one or more Probe appliances.

If you do not select an appliance from the list, the **Member Discovery Properties** option is disabled.

3. Expand the Toolbar and click **Edit → Member Discovery Properties**.

4. Click the **Credentials** side tab.

Add read-only community strings for polling SNMPv1/v2 devices on this page. Probe members inherit all SNMPv1/v2 credential information from the Grid's Discovery Properties. You can **Override**, to add SNMPv1/v2 credential records specific to Discovery operations on the Probe appliance.

5. Click **Override** to enter new SNMPv1/v2 records for the currently selected Probe member.
6. Click the Add icon to add a new community string entry to the list. Enter a Read community text string that the management system will send together with its queries to the network device during discovery. A community string is similar to a password in that the discovered device accepts queries only from management systems that send the correct community string. Note that this community string must exactly match the value that is entered in the managed system.
7. If you have a substantial list of community strings in this list and need to find a specific string, enter the value in the **Go To** field and click **Go**.
8. To remove a community string entry: check the check box and click the Delete icon.
9. To export the entire list of community strings in a table file readable by a spreadsheet program, click the Export icon and choose **Export Data in Infoblox CSV Import Format**.
 - a. To export all data in a different format, click the Export icon and choose **Export Visible Data**.

Configuring Probe SNMPv3 Properties

You can define SNMPv3 credentials on Probe appliances, that apply only to the network view to which the Probe associates.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
2. Check the checkbox for any Probe appliance on the Discovery page.
3. Expand the Toolbar and click **Edit → Member Discovery Properties**.
4. Click the **Credentials** side tab. The **SNMPv1/v2** page displays automatically.
5. Click the **SNMPv3** tab. The page automatically displays the set of SNMPv3 credentials defined at the Grid level.
6. Click **Override** to apply new SNMPv3 credentials on the chosen Probe appliance.

Note: Any SNMPv3 credentials you enter here will apply only to devices in the network view for which the Probe is associated.

7. Click the Add icon to add a new SNMPv3 authentication entry to the list. Enter the **Name** for the new credential; followed by the **Auth Protocol**, **Auth Password**, **Privacy Protocol**, **Privacy Password**, and the **Order** value, which is the order in which the SNMPv3 credentials are tried by the Probe appliance.

Note: You can press Tab to navigate across the fields for the credential entry.

8. If you have a substantial list of SNMPv3 entries and need to find a specific entry for editing, enter the value in the **Go To** field and click **Go**.
9. To remove an SNMPv3 authentication entry: check the check box and click the Delete icon.
10. To export the entire list of community strings in a table file readable by a spreadsheet program, click the Export icon and choose **Export Data in Infoblox CSV Import Format**.
 - a. To export all data in a different format, click the Export icon and choose **Export Visible Data**.
 - b. A **Show Passwords** option allows the secret keys to be visible in the data export.

Defining Seed Routers

Seed Routers can be defined only on Probe appliances. The **Seed** tab enables definition of seed routers which discovery uses to accelerate and perform discovery. Definition of seed routers is highly recommended for IPv4 networks and is required for IPv6 networks. All NIOS Probe members automatically use their default gateway as a seed router.

Note: If you do not define a seed router, it is recommended that you enable discovery for a network or DHCP range.

Click the Add icon and enter the IP address for the desired IPv4 or IPv6 seed router in the text field.

You can check **Discovery Status** to see whether a seed router is successfully being reached and whether the seed is providing information. By reviewing discovery status for each seed router, you can determine whether NIOS should be able to discover the network successfully, or if there are possible configuration errors preventing network discovery, without having to wait to see what NIOS finds. For seed routers, **Reached Status** and **Overall Status** should both read **Passed**.

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
2. Check the checkbox for any Probe appliance on the Discovery page.
3. Expand the Toolbar and click **Edit → Member Discovery Properties**.
4. Click the **Seed** side tab.
5. Click the Add icon to add a new seed router entry to the list.
6. Enter the IP address value for the new seed router in the new field.

IPv6 Seed Router Usage

For discovery of any IPv6 network, you must use seed router values, comprised of at least one well-connected IPv6 router, preferably with routes to all other networks to be managed. In some cases, seed routers may not have the full routing tables or be unable to provide full information for some reason. The general rule of thumb is that more seed routers are better, but the connectivity of seed router(s) also helps determine how many seed routers you need. Avoid having more seed router entries than necessary.

Note: For effective use of seed routers, provide SNMP credentials to the Probe member to allow it to pull the key routing and connectivity information, including the IPv6 routing table and the local Neighbor Discovery Cache, from the device. See the topic [Defining Probe SNMPv1/v2 Properties](#) on page 533 for more information.

EXCLUDING IP ADDRESSES FROM DISCOVERY

Note: You use the IPv4 or IPv6 Network editors to exclude IP addresses or ranges of IP addresses from discovery within the specified network. See the section [Disabling Discovery for a Network](#) on page 537 for more information.

Host Records, Fixed Address and IPv4 Reservations can be excluded from discovery. You may also exclude an IP address or a range of IPs within a network from Discovery.

Enabling IPs for Immediate Discovery or for Discovery Exclusion

Note: You may create a network and choose not to discover it at that time, by disabling both **Enable Immediate Discovery** and **Enable Discovery**. Conversely, you can explicitly exclude specific IPs or IP ranges from discovery. Discovery will never take place on these IPs unless the admin specifically changes their exclusion setting.

Administrators can specify IPv4 and/or IPv6 addresses that must be immediately discovered by the appliance.

Some devices may need exclusion because they do not support SNMP, or for other organizational reasons.

Devices matching IP addresses selected for immediate discovery are given one-time priority over other discovered devices, for data collection and counting toward any device found matching the license limits. The process is similar to a seed router, except that in the latter, NIOS considers the specified device a router, and specifying it as such accelerates discovery and data collection through that device.

Note: A device specified through an IP address can also be excluded from discovery or management. See [Excluding IP Addresses from Discovery](#).

In NIOS, each defined IP can be represented in a host, a fixed address, an IPv4 reservation or an Unmanaged IP.

Quick Exclusion of IPs from Discovery

You can use the IPAM IP Map or IP List page to quickly exclude IP addresses and selected ranges of IP addresses from discovery. For example, you may have NIOS appliances or routers that provide the gateway to networks that are already managed by NIOS, or known IPs associated with devices that do not support SNMP.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
2. Click on any network or network container in the list. The **IP Map** appears for the selected network.

Note: You may also use the **List** page for the selected network to exclude IPs or selected ranges of IPs. However, you will have to page through or search through the pages comprising the list view to locate the IPs you want to exclude. (If you know the IP address value in the **List** view but it does not appear on the page, enter it in the **Go to** field to search for the IP.) The **IP Map** view allows you to view every IP address in a selected network, such as a /24 prefix.

3. Select one or more IPs in the map. SHIFT+click to select a series of contiguous IPs. CTRL+click to select non-contiguous IPs.
 4. Expand the Toolbar and click **Exclusion → Enable Exclusion**. The selected IP addresses are excluded from any Discovery actions.
-

Note: You can click the Actions icon for any List record and choose **Exclusion → Enable Exclusion**.

To locate an IP to exclude within a network container:

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
 2. Select the network container by clicking it. The IPAM Home page changes to display the List page, showing the list of networks within the container.
 3. Click the network that has the IP in its space that you wish to exclude.
 4. Select the network IP address from the **List** table. If you know the IP address value but it does not appear on the page, enter it in the **Go to** field to search for the IP.
 5. Expand the Toolbar and choose **Exclusion → Enable Exclusion**.
-

Note: You can click the Actions icon for any List record and choose **Exclusion → Enable Exclusion**.

A parent network container may exclude IPs, and you may add and remove discovery exclusions within network containers. You can drill down to the child networks in the **Net Map** view to perform exclusions on IPs. For example, consider a /16 network container that has a number of smaller /24 child networks within it. Right-clicking on any child network in the network container and choosing **Edit** from the popup menu, opens the editor with its **Discovery Exclusions** tab, where you can perform exclusions within the child network.

Creating a New Fixed Address Object and Excluding it from Discovery

You can create a new Fixed IP through the **Add IPv4 Fixed Address Wizard** or the **Add IPv6 Fixed Address Wizard**.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
 2. Select a network from the IPAM home page by checking the network's check box.
 3. Expand the Toolbar and click **Add → Fixed Address → IPv4** or **Add → Fixed Address → IPv6**.
 4. Click **Next Available IP** to obtain the next available IP address in the chosen network. For more information about obtaining the next available IP address, see [About the Next Available Network or IP Address](#) on page 844.
-

Note: The appliance displays an error message if the obtained next available IP address is already being used by other users. You can request for another unused IP address or enter a new one.

5. If the network of the IP address is served by a Grid member, Grid Manager displays the **Assign IP Address by** section, with its **MAC Address**, **DHCP Client Identifier** and **DHCP Relay Agent** settings. Select the different options as needed to define a fixed IP Address. For more information, see [Configuring IPv4 Fixed Addresses](#) on page 857.
 6. Click **Next** to continue to the DHCP Options page in the wizard.
-

Note: See [About the Next Available Network or IP Address](#) on page 844 for more information on DHCP Options configuration.

7. If you do not wish to configure DHCP Options for this Fixed IP, click **Next** to go to the following Wizard step, for defining Discovery settings.
8. Choose from the options on the Step 4 Wizard page:
 - Check **Exclude from Network Discovery** to prevent the Fixed IP from being probed by discovery.
 - (Enabled by default) If you want immediate discovery of the current Fixed IP, check the **Enable Immediate Discovery** checkbox.
Both check boxes may be disabled. Doing so does not disable discovery for the current object—discovery is simply performed by NIOS on its own internal timetable.
 - To override SNMP credentials for either SNMPv1/v2 or for SNMPv3 for the current Fixed IP, check the **Override Credentials** checkbox.
You can enter both a SNMPv1/v2 Read community string and an SNMPv3 credential, or enter only the single type you need. Each can be selected and edited in turn.
 - Select **SNMPv1/v2** and enter the **Read** community string;
—or—
 - Select **SNMPv3** and enter the device admin account name, and the **Auth Protocol**, **Auth Password**, **Privacy Protocol** and **Privacy Password** values where necessary.
9. Click **Next** to go to the final Wizard step.
10. If necessary, add or apply any extensible attributes necessary for the new record.
11. Click **Save & Close**.

Excluded IP Addresses in NIOS

You may exclude IP addresses from Discovery from within a number of different contexts in NIOS. Under IPAM, you can exclude in the IP Map and IP List pages. The IP List page provides an Exclude data column that directly shows the exclusion status for all IPs in the selected network. Various objects, such as Host records, IPv4 and IPv6 Fixed Addresses, and IPv4 reservations, may be excluded from Discovery.

You may view excluded IP addresses in the IP List page or in the network editor's Exclusions tab.

DISABLING DISCOVERY FOR A NETWORK

You go to the DHCP feature under **Data Management** to disable discovery for a network.

To disable discovery for an IP network:

1. From the **Data Management** tab, select the **DHCP** tab and click **Networks**.
2. In the **Networks** page, select the IP network that you want to exclude from discovery.
3. Expand the Toolbar and click the Edit icon.
4. Click the **Discovery** side tab.
5. Child networks inherit their discovery default settings from their parent networks. Click **Override** to change the **Enable Discovery** setting. (The **Discovery Member** setting will remain unchanged.)
6. Deselect the **Enable Discovery** check box.
7. Click **Save & Close**.

VIEWING THE LIST OF DISCOVERED DEVICES

IPAM provides a Devices page for a complete listing of every device Discovery finds within its reachable networks. You can explore a considerable body of information about the complete list of discovered devices starting on the home Devices page, and the ability to drill down to specific information about every discovered device.

1. From the **Data Management** tab, select the **Devices** tab. The Devices page provides a list of all discovered network infrastructure devices and end hosts that have been discovered and catalogued.

Values listed in the Devices page include the following:

- **IP Address:** Detected IP address (IPv4 or IPv6), if any.
- **Name:** Detected name of the device. Each device name provides a link to the list of interfaces associated with the device.
- **Type:** The network device type: **Router**, **Switch-Router**, **Firewall**, **NIOS** (Infoblox appliance), **vNIOS**, and others.
- **Model:** The model name as detected from the device during Discovery.
- **Vendor:** The equipment manufacturer (Cisco, Juniper, Fortinet, F5, and many others).
- **Device Version:** The Operating System version for the network device.
- **Location:** The physical location of the network device as detected from the device during Discovery.
- **Description:** Verbose description of the network device as collected from the device by Discovery.

You can click **Discovery Status** to view the same list of network devices showing the Discovery data set. You can sort the table by any displayed value. Use NIOS-standard filtering to isolate device names, IP addresses or other values in which you are interested.

For each listed device, the Actions icon also provides the following options depending on the listed device type:

Interfaces: Displays the Interfaces page for the chosen device. (See [Viewing the List of Networks Associated with a Discovered Device](#) on page 540.)

Show IPAM IP Address: Shows the Management IP address for the device that has a network in IPAM—the IPAM tab appears, showing details for the IP address.

Networks: a drop-down list of all IPv4/IPv6 networks to which the currently selected device connects.

Device Details: a basic list of information about the chosen device, including the IP address by which the device is discovered, operational status, IPAM Type (whether the device is Managed or Unmanaged), the Device Type and the number of Interfaces.

VIEWING DISCOVERY STATUS

Note: Opening Discovery Status for viewing requires Superuser permissions under NIOS.

You can view a list showing the complete Discovery status of all device or of selected devices.

To isolate devices for evaluation, use filtering to reduce the list. Click **Use Filter** at the top of the table and choose **IP Address**, **Name** or **Overall Status** as the filter.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
2. Expand the Toolbar and click the **Discovery Status** button.

The Discovery Status table lists detailed information about network devices and end hosts discovered through all methods, including SNMP, ICMP ping sweeps and other processes.

- **IP Address:** the IPv4 or IPv6 address of the discovered device. You can filter the table by this value.
- **Name:** The name of the discovered device as reported through SNMP. You can filter the table by this value.
- **Type:** The discovered device type. Examples include Router, NIOS, Switch-Router, Firewall, Load Balancer, and numerous others.
- **Overall Status:** Indicates the overall success or failure of the discovery operation on the device. Hover the mouse over the device to see more detailed information about the discovery status of the device, including the timestamp of the last Discovery event, confirmation of detection (“Device Exists”), and the means of detection, which are usually methods such as SNMP, reading the ARP table or location through a Seed router. You can filter the table by this value.
- **Reached Status:** Indicates the reachability of the discovered device. Typically, devices will be reported as Reachable through SNMP, a path trace through ICMP, or UDP-based path tracing for an IPv6 address.

- **SNMP Collection Enabled:** Indicates whether the managed device allows SNMP as a management protocol. This value will show **Yes** or **No**.
- **SNMP Credential Status:** Indicates whether the correct SNMP credential is used by Discovery. Will usually show simple Passed or Failed status.
- **SNMP Collection Status:** Indicates whether managed device information has been successfully collected from the device. If the current device shows an **SNMP Credential Status** of Failed, this field will remain blank.
- **Fingerprint Status:** Shows the status of discovery of the device's OS through fingerprinting.
- **Last Update:** Timestamp showing the conclusion of the last data update for the current device.
- **First Seen:** Timestamp showing the initial Discovery event.
- **Last Seen:** The date and time when the device was last successfully polled by discovery.
- **Last Action:** The last action performed by discovery upon the device after the Discovery took place. Hover the mouse over this field to obtain details.

Visible columns can be changed in the Discovery Status window. At the top of any column header, click the down arrow tool, and choose **Columns – > Edit Columns**.

See the topic [Using Discovery Diagnostics](#) for more information on checking and diagnosing discovery behavior for devices listed in the status table.

USING DISCOVERY DIAGNOSTICS

Note: Opening Discovery Diagnostics for viewing requires Superuser permissions under NIOS.

You can apply a diagnosis to help determine why a specific device is presenting difficulties in discovery. For example, a given device may be reachable but show an Overall Status of Failed in the Discovery Status dialog. The Discovery Diagnosis feature steps through a complete discovery process based on the configuration present on the Probe member to which the device is assigned. The diagnosis runs the gamut from fetching SNMP object ID information, to ARP table reading, to ICMP pings and traceroutes. The Diagnostics page runs a full discovery procedure against the specified IP address.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
If you do not have the device IP address for diagnosis, open the Discovery Status page and locate it there (see [Viewing Discovery Status](#) on page 538 for more information).
2. Expand the Toolbar and click the **Discovery Diagnostics** button.
3. Enter the IPv4 or IPv6 address of the device.
4. If there are multiple network views that have discovery Probe members, you can select a network view and NIOS will pick the optimal Probe member for that Network View. Note that a Probe appliance can only operate against one network view. However, multiple Probe appliances can be associated with the same network view.
5. You can enter a community string instead of using the global list of credentials. If you possess the Read community string for the device and need to use it to help enable diagnostics, enter it in the **Community String** field. If you do not enter any value, the global list will be used.
6. To force a complete test against the device, choose **Yes**, for the **Force Test** setting.
7. To select all the log output text, click **Select All**. The contents can then be copied and pasted to the Clipboard and a text editor.

The output log shows the attempt for the complete Discovery process.

VIEWING DISCOVERED INTERFACE INFORMATION

Discovery provides complete accounting for all network interfaces on discovered devices. Interfaces are detected whether loopbacks, unnumbered, or numbered with one or more IP addresses. Interfaces may be listed for either managed or unmanaged devices.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the Action icon for a chosen device and choose **Interfaces** from the popup menu.

Note: You can also simply click the device name to display the Interfaces list.

The Interfaces page appears, displaying a table of all ports for the chosen device. Every listed interface provides its own Action icon, from which you can choose **Show Assets**. This link provides a list of all end-host devices or neighboring devices linked to the current interface. **Show Assets** is only available for switched L3 interfaces with no IP Address. NIOS filters the asset list for the device by the interface name.

Click **Devices Home** to return to the main **Devices** page.

Data points include the following (some data may appear for some device types and not for others):

- **Name:** Detected name of the interface, if any.
- **IP Address:** Detected IP address (IPv4 or IPv6), if any.
- **MAC Address:** The hardware address associated with the interface.
- **VLAN Name/VLAN ID:** The VLAN(s) to which the interface is bound, if applicable.
- **Port Speed:** Interface speed, in Mbps or Gbps.
- **Port Type:** Type of interface as detected by NIOS Discovery. Examples include **ethernet-csmacd**, **propVirtual**, **propPointToPoint Serial**, **l2vlan**, **mplsTunnel**, **tunnel**, and others.
- **Admin Status:** Shows whether the interface is administratively Up or administratively Down—in effect, whether or not the interface is enabled by the administrator.
- **Operation Status:** Shows whether the interface is operationally Up or operationally Down. Indicates whether the port has established an adjacency and is capable of passing packets. A port can be Administratively Up and Operationally Down.
- **Description:** Plain-English note for the interface, generated by the Consolidator from collected information.
- **Trunk Status:** Where applicable, shows the trunking status of the interface.
- **Link Aggregation:** Shows the state of the interface if it is part of a LAG (Link Aggregation Group).
- **Status:** NIOS management status of the IP address associated with the interface.
- **IPAM Type:** The object type that is associated with the IP address for the interface, where applicable. Possible values can be **Unmanaged**, **Lease**, **IPv4 DHCP Range** or **Fixed Address**.
- **Usage:** Indicates whether NIOS has configured the IP address for DNS or DHCP.

The Interfaces table may be sorted by IP address, Name, MAC Address, Port Type, and other values.

Viewing the List of Networks Associated with a Discovered Device

You can view a list of all networks with which any discovered device is associated by clicking the device name on the main **Devices** page.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays the list of all devices currently found and catalogued by discovery.
2. Click the device name from the **Name** column in the table.

The **Network** page appears, displaying a table of all networks to which the chosen device connects.

NIOS sorts the list of networks by IP address. By default, only the **Network IP**, a comment field, and the **Site ID** appear in the Networks table.

Viewing the Management State of IPs in Discovered Networks

You can view the management state for any IP address, in any network, that is associated with any discovered device.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the Action icon for a chosen device and choose **Networks** from the popup menu.
3. Choose a network from the list. NIOS switches to the IPAM page view of the selected network.

The IPAM Home page displays the IP Map for the chosen network. The page shows information in graphical format, indicating elements such as **Used** Addresses, fixed addresses and IP reservations, **Unmanaged** IPs, **Host Not in DNS/DHCP**, and all other objects or information associated with IP management. The user benefits from this view by immediately seeing which IPs in the network contain devices that remain unmanaged by NIOS. These Unmanaged IP values appear in light yellow. Hovering the mouse over any IP address in the graphical table shows the information that has already been determined about the IP address.

Note: An Unmanaged IP cannot be converted to Managed unless the network that contains it, is converted to Managed status. For more information, see the topic [Converting Unmanaged Networks to Managed Networks](#) on page 542.

Viewing the List of IP Addresses Associated with a Discovered Device

You can view the complete list of discovered IP addresses bound to all interfaces for any device, Discovered and Managed devices alike. One useful trick for interfaces is to pick out an interface from the **Interfaces** page that has multiple IPs and open the **IP Addresses** tab; or sort the IP addresses table by its **IP Address** column, and locate the interface name bearing multiple IPs. Frequently, an interface with multiple addresses will have IPv4 and IPv6 addresses bound to it. Loopbacks are another example.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click on the device name. The IP Addresses page appears, sorting its table by IP address value. Data points include the following:
 - **Action:** An icon providing a single menu choice, **Device Details**.
 - **IP Address:** Detected IP address (IPv4 or IPv6).
 - **Interface Name:** the name of the device port. The Interface Name provides a link drilling down to the IP panel. (See [Converting Unmanaged Networks to Managed Networks](#) on page 542 for more information).
 - **MAC Address:** The hardware address associated with the interface.
 - **VLAN Name/VLAN ID:** The VLAN(s) to which the interface is bound, if applicable. In most cases, you see both the VLAN name and the VLAN ID as two values in the same field. Multiple VLAN entries may be present for an interface or IP Address.
 - **Operation Status:** Shows whether the interface is operationally Up or operationally Down—effectively, whether the interface is successfully connected to the network.

The list can be sorted by any of the displayed columns.

3. Click the IP Address link for any interface to open the Related Objects page for the chosen port.

Viewing the List of Assets Associated with a Discovered Interface

In Discovery, NIOS classifies end hosts and any other devices connected to switchport interfaces as “Assets” directly associated with each discovered interface. At a glance, the **Assets** page shows all devices associated with the chosen interface, including switchports supporting dozens of end hosts. Displaying the Assets list applies only to enterprise L2 Ethernet switches. In practice, most Asset tables will show end hosts and devices that populate Ethernet network segments.

The Assets table lists all managed end hosts and application servers detected through discovery and identity resolution by NIOS, that are connected to each network infrastructure device. The records listed in this table date from the Last Seen discovery time stamp of each end host.

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Name** link for the device you want to inspect. The **Interfaces** tab appears for the chosen device.
3. Click the Action icon for an interface in the table, and choose **Show Assets**. (Applies only to switched interfaces that do not have an IP address.)

Values listed in the Assets table include the following:

- **Name:** The device's name on the network as discovered by NIOS. If the Name is that for another infrastructure device, you may click on it to see its associated Assets.
- **Interface Name:** The name of the interface (typically a switched interface) associated with the discovered device.
- **IP Address:** The IP Address for each discovered end host as managed by NIOS and IPAM. The IP address is a link to the home IPAM page for the interface.
- **MAC Address:** The hardware MAC address associated with the asset.
- **VLAN Name/VLAN ID:** The VLAN identifier from which the asset is reachable.
- **Operation Status:** Will normally read Up or Down. Asset records may appear as Down because they are disconnected from the network or being rebooted.

In the Interfaces tab, if you select an interface for a switch that is only connected to a neighboring switch, router, or switch-router, and then choose **Show Assets**, the Assets page displays only the neighboring device interface that is reachable from the chosen port.

CONVERTING UNMANAGED NETWORKS TO MANAGED NETWORKS

The IPAM main page lists all Discovered networks as Unmanaged, highlighted in yellow. Administrators cannot apply services or Objects to IP addresses in unmanaged networks until they are converted to Managed. You can explore Unmanaged networks through IPAM's IP Map and IP List views, but many operations cannot be carried out on unmanaged networks, including editing, splitting, resizing, permissions changes and many other operations.

Unmanaged IP addresses that are part of an unmanaged network cannot be independently converted to a managed IP address.

You can choose to Discover the network after it is converted, or to keep discovery disabled and execute it at another time.

1. From the **Data Management** tab, select the **IPAM** tab.
2. Select one or more Unmanaged networks from the IPAM list.
3. Expand the Toolbar and click the **Convert** button.
4. If necessary, click the Discovery tab and **Enable Discovery** to start discovery on the network immediately after it is converted to Managed state. You can elect not to discover the network.
 - a. Click **Select Member** to choose the Probe member by which the network is discovered.
 - b. If necessary, click **Override** under **Polling Options**, and modify the device discovery polling options for the network.
5. Click **Save & Close** to make the conversion.

The IPAM main page shows **Yes** as the value for the network under the **Managed** column.

Most conversion operations for networks and individual IP addresses are managed under IPAM and are described in the section [Managing IPv4 and IPv6 Addresses](#) on page 487 of this Guide.

ADDING DISCOVERY DEVICE SUPPORT

Note: Adding Device Support Bundles, viewing them and deleting them requires Superuser permissions under NIOS.

Infoblox frequently provides support files for additional network devices that may not previously be supported by discovery, and updates to support new operating system versions of existing devices.

To add device support updates to the NIOS system:

1. From the **Grid** tab, select the **Device Support** tab.
2. Expand the Toolbar and click the **Add** button.
3. Click **Select** and navigate to the file you want to upload.
4. Select the file, and then click **Upload**.

The Device Support table shows its installed library of files with the following data points:

- **Name:** The descriptive device name for the device support file.
- **Version:** The version of the currently active device support file.
- **Author:** The developer of the device support file.

DISCOVERY DATA MANAGEMENT

You can separately back up and restore Discovery data. Procedures for doing so are straightforward and are based upon the standard NIOS file backup and file restore features at the Grid level.

NIOS backs up the discovery database in XML format and compresses the data in a TAR GZip file. The same file can be restored to a NIOS appliance.

For more information, see the topics under [Backing Up and Restoring Configuration Files](#) on page 423.

PERFORMING VM (VIRTUAL MACHINE) DISCOVERY

A VM discovery retrieves information about vSphere servers, ESX servers, and the virtual entities running on each server. You can add more than one vSphere server or ESX server to the discovery session. When you disable specific servers, the appliance excludes them from the VM discovery.

You ensure that NIOS can communicate with any VMware vCenter server and the ESX servers used to house virtual machine instances. ESX servers and the VMs they house may already be discovered by the Discovery service. Running a VM Discovery may determine additional data.

When you perform a VM discovery, the appliance communicates with the specified vCenter and/or ESX and ESXi servers to collect vSphere-specific data. VM discovery will discover only the networks specified in the Discovery dialog.

Executing a VM Discovery

1. From the **Data Management** tab, select the **IPAM** tab.
2. Check the check box for each network for which you wish to run VM Discovery.
3. Expand the Toolbar and click **VM Discovery → Discover Now**.
4. In the **General** tab, click **Select Member** and choose any Grid member (it does not have to be a Probe member) by which the VMs will be discovered and catalogued.
 - If necessary, ensure the **Merge the Discovered Data with Existing Data** and/or **Update discovered data for managed objects** check boxes are enabled. NIOS enables both check boxes by default.

5. In the **Networks** tab, ensure the desired networks for discovery are checked.
6. In the **Servers** tab, click the Add icon. Complete the following in the **Add vSphere Server** section:
 - **Server:** Enter the FQDN or IP address of the vSphere server.
 - **Protocol:** Select the protocol that is used to connect to the vSphere server. The default is HTTPS.
 - **Port Number:** Enter the number of the port the appliance uses to communicate with the vSphere server. The default is 443.
 - **Username:** Enter the username the appliance uses to log in to the vSphere server. The user account on the vSphere server should have at least read-only permission.
 - **Password:** Enter the password of the vSphere server account.
7. Click **Test** to test the settings before adding them to the table.
8. Click **Add** to add the vSphere server to the table. You can also do the following in the table:
 - Click the Add icon again to add more vSphere servers.
 - Select the **Disable** check box in the table to exclude a specific vSphere server from the VM discovery. The appliance keeps the server configuration when you disable the server. All servers in the list are included in the VM discovery by default.
 - Select a server and click the Edit icon to modify its configuration.
 - Select a server and click the Delete icon to delete the server.
 - Click the Export icon to export the data in CSV format.
9. Add any further records as needed.
10. Click **Save**, and then click **Start** to begin the VM Discovery.

Scheduling a VM Discovery Session

After you configure a VM discovery, you can start the discovery process immediately or schedule a single discovery for a later time (using the **Discover Now** option) or schedule it for a later date or use all available scheduling options for a recurring discovery (using the **Schedule Discovery** option).

A recurring VM discovery is a VM discovery that repeats at regular intervals. When you start a discovery immediately, or schedule it for a later date and time after you define it, the discovery happens only once and it will not be repeated.

You can configure a VM discovery to execute repeatedly based on a defined schedule. When you select a network and choose **VM Discovery → Schedule Discovery** from the Toolbar on the IPAM page, you can schedule a VM discovery process to recur on an hourly, daily, weekly, or monthly basis.

You can also use the VM Discovery widget on the NIOS Dashboard Status page to define a scheduled or recurring VM discovery process.

You can configure the **Discover Now** and **Schedule Discovery** jobs independent of each other and each one contains a specific set of networks and discovery settings.

Note: You can pause and resume all discovery tasks.

For the complete procedure on how to configure a recurring VM discovery, see [Configuring a Recurring Discovery](#) on page 506.



PART 4 DNS

This section describes how to configure the Grid to provide DNS services. It includes the following chapters:

- [Chapter 15, *Infoblox DNS Service*](#), on page 547
- [Chapter 16, *Configuring DNS Services*](#), on page 555
- [Chapter 17, *DNS Views*](#), on page 601
- [Chapter 18, *Configuring DNS Zones*](#), on page 615
- [Chapter 19, *DNS Resource Records*](#), on page 655
- [Chapter 20, *Configuring DDNS Updates from DHCP*](#), on page 689
- [Chapter 21, *DNSSEC*](#), on page 733
- [Chapter 22, *Configuring IP Routing Options*](#), on page 757



Chapter 15 Infoblox DNS Service

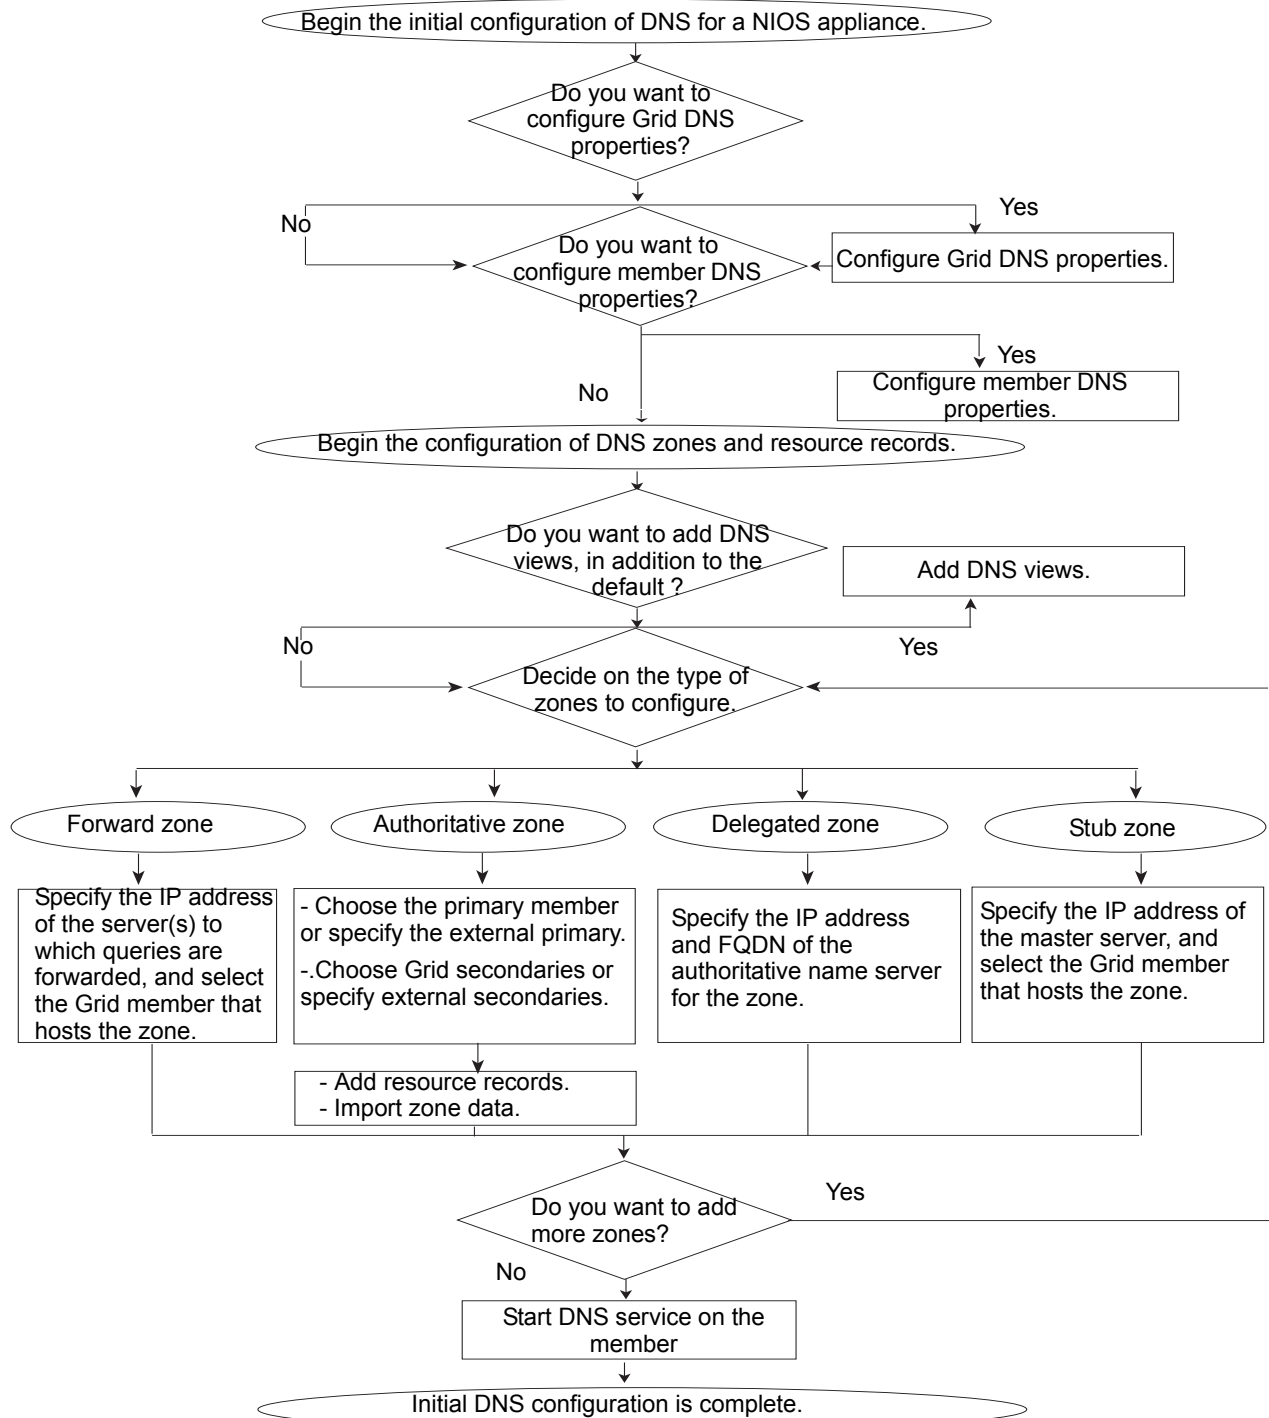
The NIOS appliance uses a standard, BIND-based DNS protocol engine. It interoperates with any other name server that complies with the DNS RFCs (see [DNS RFC Compliance](#) on page 1313).

This chapter provides an overview of the DNS configuration tasks. It includes the following sections:

- [Configuring DNS Overview](#) on page 548
 - [DNS Configuration Checklist](#) on page 549
- [About Inheriting DNS Properties](#) on page 550
 - [Overriding DNS Properties](#) on page 551
- [Understanding DNS for IPv6](#) on page 552
 - [Configuring IPv6 on a Grid Member](#) on page 553
 - [Configuring DNS for IPv6 Addressing](#) on page 554

CONFIGURING DNS OVERVIEW

An overview of the DNS configuration process is outlined in the following diagram, illustrating the required steps for preparing a NIOS appliance for use:



DNS Configuration Checklist

The following checklist includes the major steps for configuring DNS:

Table 15.1 DNS Configuration Checklist

Step	For more information
Decide if you want to configure DNS properties for the Grid and for individual members	<ul style="list-style-type: none"> • Chapter 15, Infoblox DNS Service, on page 547
Decide if you want to create a new DNS view, in addition to the default DNS view	<ul style="list-style-type: none"> • Chapter 17, DNS Views, on page 601
Decide which type of DNS zone you want to configure	<ul style="list-style-type: none"> • Chapter 18, Configuring DNS Zones, on page 615
Add hosts and resource records	<ul style="list-style-type: none"> • Chapter 19, DNS Resource Records, on page 655
Import zone data	<ul style="list-style-type: none"> • Importing Zone Data on page 631
Enable DNS service on the member	<ul style="list-style-type: none"> • Starting and Stopping the DNS Service on page 565

ABOUT INHERITING DNS PROPERTIES

You can configure DNS properties at the Grid, member, zone, and resource records level. The NIOS appliance applies the properties hierarchically, with the Grid at the top of the hierarchy. Grid settings apply to all members in the Grid, unless you override them at the member, zone, or resource record level. When you set DNS properties for a particular member, these properties override the Grid properties and apply to all zones served by that member. When you set properties for a specific zone, they override the member properties and apply to the resource records in the zone. You can also override the zone properties and set properties for specific resource records.

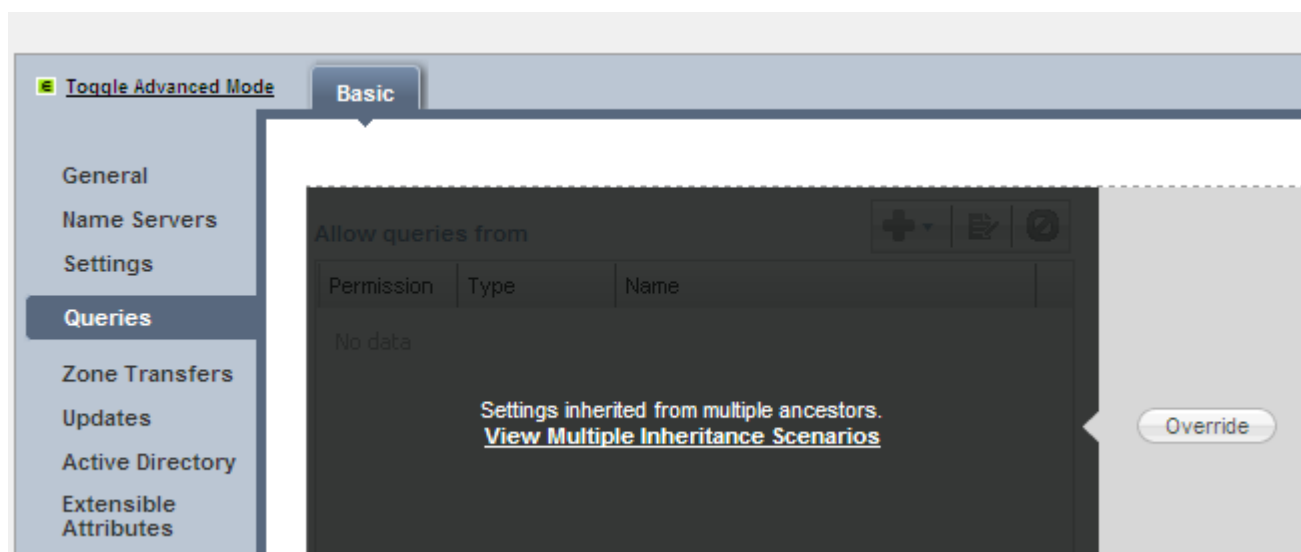
When you configure DNS properties that contain inherited values, the appliance displays the information based on the inheritance sources. There may be times when an object can inherit properties from different sources with different settings. The following table summarizes what the appliance can display:

When you see...	it means...
Inherited From <i><object></i>	the DNS property has a definite value from an inheritance source.
Inherited From Upper Level	the appliance cannot yet determine the inherited value or inheritance source for the DNS property.
Inherited From Multiple	the DNS property has the same value that it inherits from multiple sources.
Settings Inherited from Multiple Ancestors, View Multiple Inheritance Scenarios	the DNS property has different values that it inherits from multiple sources, and you can view the values and their corresponding sources by clicking the View Multiple Inheritance Scenarios link.

Based on the information provided, you can then decide whether to override or keep the inherited values. You must have read/write permissions to the DNS resources to override inherited values. You can only view inherited values and paths if you have at least read-only permissions.

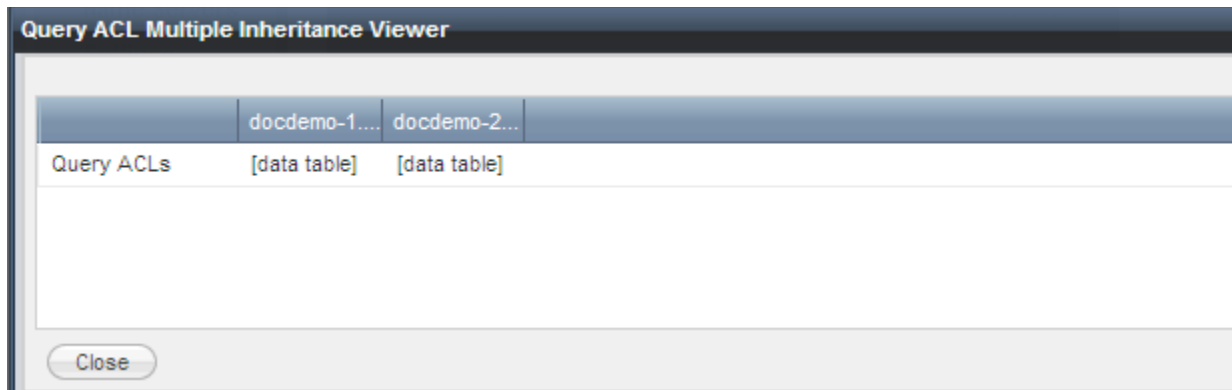
In the example in [Figure 15.1](#), the DNS zone is served by members with different query settings.

Figure 15.1 DNS Zone with Different Inherited Settings



The Multiple Inheritance Viewer indicates that the two servers have different query ACLs, as shown in [Figure 15.2](#). You can then view the Query properties of each member and edit them, or override the setting and specify values that apply to the zone only.

Figure 15.2 Multiple Inheritance Viewer



Overriding DNS Properties

DNS properties configured at the Grid level apply to the entire Grid. You can choose to keep the inherited properties or override them when you configure the properties for a member, zone, or resource record.

To override an inherited value:

1. In a wizard or editor, click **Override** next to a property to enable the configuration. The **Override** button changes to **Inherit**.
2. Enter a new value to override the inherited value.

UNDERSTANDING DNS FOR IPV6

You can configure NIOS appliances to provide DNS services over IPv4 (Internet Protocol version 4) and IPv6 (Internet Protocol version 6) networks. You can configure the Grid member as a dual-mode name server, capable of sending and receiving IPv4 and IPv6 queries and responses. It can serve DNS data in response to both IPv4 and IPv6 queries. The appliance supports authoritative forward-mapping zones containing AAAA records mapping host names to IPv6 addresses, as well as authoritative reverse-mapping zones with PTR records mapping IPv6 addresses to host names. Configuring a Grid containing an IPv4 primary server and IPv6 secondary servers is not supported. You must enable IPv6 on both the primary and secondary servers within the Grid to enable them to communicate with each other. Infoblox highly recommends that you enable IPv6 on your Grid appliances before configuring IPv6 authoritative zones.

The NIOS appliance supports one IPv6 address per Grid member. Infoblox integrates IPv6 address management into many of the same places where IPv4 addresses are entered. Data validation occurs on all IP address fields and automatic validation is done to ensure proper entry of either an IPv4 address or an IPv6 address.

The NIOS appliance supports the following DNS functions for IPv6:

- **AAAA records**—You can import, serve queries, display, add, delete, and modify AAAA records on the appliance. An AAAA record is equivalent to an IPv4 A record, relying upon a forward-mapping zone to map a hostname to an IPv6 address. A single forward-mapping zone can map names to both IPv4 and IPv6 addresses. The appliance autogenerates AAAA records for any of its interfaces that have IPv6 addresses.
- **Hosts**—You can configure IPv4 and IPv6 addresses for hosts. For information, see [Adding Host Records](#) on page 462.
- **ip6.arpa**—A specific domain for IPv6 is used for DNS reverse lookups called ip6.arpa. This domain maps an IPv6 address to a hostname. When you specify an IPv6 network, the appliance automatically creates the appropriate zone under ip6.arpa.
- **PTR records**—Import, serve queries, display, add, delete, and modify PTR records within an ip6.arpa reverse zone. The PTR record returns a domain name corresponding to an IPv6 address contained in the ip6.arpa zone. The appliance does not autogenerate PTR records; the user must configure PTR records manually.
- **DDNS**—The appliance supports AAAA and PTR records for DDNS (Dynamic DNS).

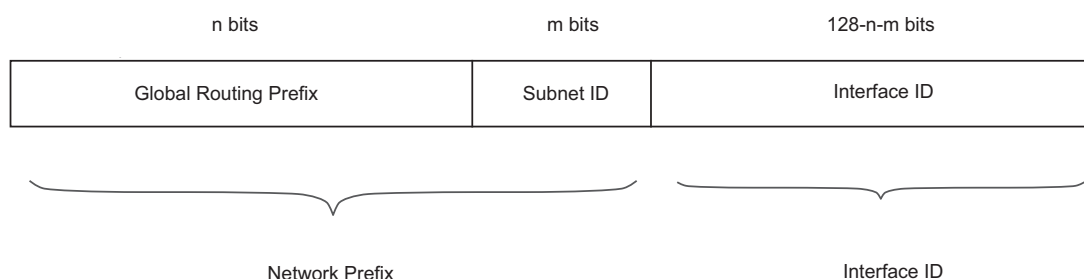
For more information about DNS for IPv6, see RFC 3596, *DNS Extensions to Support IP Version 6*.

Address Structures

IPv4 uses a 32-bit, 4-octet (each octet separated by decimals) addressing structure to designate sources and destinations within a network. Since there are 32 bits that make up the address, IPv4 can support up to 4 billion unique addresses.

An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight groups of four hexadecimal digits separated by colons (example: 12ab:0000:0000:0123:4567:89ab:0000:cdef). Since there are 128 bits that make up the address, IPv6 can support up to 3.4×10^{38} unique addresses. The increase in the number of unique IPv6 addresses is one of the biggest advantages of an IPv6 implementation.

Figure 15.3 IPv6 Address Structure



The IPv6 address structure consists of the following:

- **Global Routing Prefix**—Global routing prefix is a (typically hierarchically-structured) value assigned to a site.
- **Subnet ID**—Subnet ID is an identifier of a link within the site.
- **Interface ID**—Interface Identifier. This portion of the address identifies the interface on the subnet. This is equivalent to the host identifier for IPv4 addresses.

When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered.

Configuring IPv6 on a Grid Member

You can configure a Grid member to support both IPv4 and IPv6 connections by configuring an IPv6 address on the member, in addition to the standard IPv4 address.

When you enable IPv6 on a member, you can manually enter the IPv6 gateway address or enable the member to automatically acquire the address from router advertisements. Routers periodically send router advertisements that contain link-layer addresses and configuration parameters. A NIOS appliance that supports IPv6 can listen for router advertisements and obtain the default gateway IP address and link MTU (maximum transmission unit). The link MTU is the maximum packet size, in octets, that can be conveyed in one transmission unit over a link. Thus you can set parameters on a router once and automatically propagate it to all attached hosts.

To configure the member to support IPv6:

1. From the **Grid** tab, select the **Grid Manager** tab -> *Grid_member* check box -> Edit icon.
2. Select the **Network** -> **Basic** tab of the *Grid Member Properties* editor.
3. Click the Add icon of the Additional Ports and Addresses table, select IPv6 and complete the following:
 - **Address:** Type the IPv6 address for the Grid member on the interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 12ab:0000:0000:0123:4567:89ab:0000:cdef).
 - **Subnet Mask:** Choose the CIDR netmask for the subnet to which the VIP address connects. The prefix length can range from 0 to 128, due to the larger number of bits in the IPv6 address.
 - **Gateway:** Do one of the following:
 - Type the IPv6 address of the default gateway of the subnet to which the VIP address connects.
 - Type **auto** to enable the appliance to acquire the IP address of the default gateway and the link MTU from router advertisements.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring DNS for IPv6 Addressing

Configuring the appliance to manage DNS services for IPv6 connections is similar to configuring DNS services for IPv4 connections. For simplicity, the IPv6 procedures are located in the same location as the corresponding procedures for IPv4 in this chapter. In most cases, the key difference within the procedure involves selecting an IPv6 mapping zone instead of an IPv4 mapping zone. You can configure the following tasks:

Table 15.2 IPv6 DNS Configuration Checklist

Step	For more information
Create primary or secondary name servers and specify an IPv6 root server.	<ul style="list-style-type: none"> • About Authoritative Zones on page 616 • Specifying a Primary Server on page 623 • Specifying a Secondary Server on page 626 • Creating a Root Zone on page 620
Configure the IPv6 zones.	<ul style="list-style-type: none"> • Creating an Authoritative Forward-Mapping Zone on page 617 • Creating an Authoritative Reverse-Mapping Zone on page 618
Configure IPv6 resource records	<ul style="list-style-type: none"> • Managing AAAA Records on page 662 • Managing PTR Records on page 664



Chapter 16 Configuring DNS Services

This chapter provides general information about DNS service properties. The topics in this chapter include:

- [*Configuring DNS Service Properties*](#) on page 557
 - [*Configuring DNS Access Control*](#) on page 557
 - [*About Time To Live Settings*](#) on page 557
 - [*Adding an Email Address to the SOA Record*](#) on page 560
 - [*Notifying External Secondary Servers*](#) on page 561
 - [*Enabling the Configuration of RRset Orders*](#) on page 561
 - [*Specifying Port Settings for DNS*](#) on page 562
 - [*Specifying Minimal Responses*](#) on page 565
 - [*Starting and Stopping the DNS Service*](#) on page 565
- [*About DNS Cache*](#) on page 565
 - [*Clearing DNS Cache*](#) on page 565
 - [*Clearing Cache for DNS Views*](#) on page 566
 - [*Clearing Domain Names from Cache*](#) on page 566
 - [*Viewing DNS Configuration*](#) on page 567
 - [*Viewing DNS Cache Details*](#) on page 568
 - [*Viewing Statistics*](#) on page 568
- [*Using Forwarders*](#) on page 569
 - [*Specifying Forwarders*](#) on page 569
- [*Controlling DNS Queries*](#) on page 570
 - [*Specifying Queriers*](#) on page 570
- [*Enabling Recursive Queries*](#) on page 571
 - [*Enabling Recursion*](#) on page 571
 - [*Restricting Recursive Clients*](#) on page 572
- [*Controlling AAAA Records for IPv4 Clients*](#) on page 573
 - [*Enabling AAAA Filtering*](#) on page 573
- [*About NXDOMAIN Redirection*](#) on page 574
 - [*About NXDOMAIN Rulesets*](#) on page 575
 - [*NXDOMAIN Redirection Guidelines*](#) on page 577
 - [*Configuring NXDOMAIN Redirection*](#) on page 577
 - [*Creating Rulesets*](#) on page 577
 - [*Enabling NXDOMAIN Redirection*](#) on page 578

- [*About Blacklists*](#) on page 579
 - [*About Blacklist Rulesets*](#) on page 580
 - [*Blacklist Guidelines*](#) on page 581
 - [*Configuring the Blacklist Feature*](#) on page 581
 - [*Enabling Blacklisting*](#) on page 582
- [*Enabling Zone Transfers*](#) on page 583
 - [*Configuring Zone Transfers*](#) on page 584
 - [*Configuring Concurrent Zone Transfers*](#) on page 586
- [*About Root Name Servers*](#) on page 587
 - [*Specifying Root Name Servers*](#) on page 587
- [*About Sort Lists*](#) on page 588
 - [*Defining a Sort List*](#) on page 588
- [*Configuring a DNS Blackhole List*](#) on page 590
 - [*Defining a DNS Blackhole List*](#) on page 590
- [*Specifying Hostname Policies*](#) on page 592
 - [*Defining Grid Hostname Policies*](#) on page 592
 - [*Defining Hostname Restrictions*](#) on page 593
 - [*Obtaining a List of Invalid Record Names*](#) on page 594
- [*About DNS64*](#) on page 594
 - [*Configuring DNS64*](#) on page 595
 - [*About Synthesis Groups*](#) on page 595

CONFIGURING DNS SERVICE PROPERTIES

You can configure general DNS service properties and change some default values. The DNS service is disabled by default. To enable the member to provide DNS service, you must start the DNS service. For information about how to start and stop the DNS service, see [Starting and Stopping the DNS Service](#) on page 565.

The following sections describe the DNS service properties that you can configure:

- [Configuring DNS Access Control](#)
- [About Time To Live Settings](#)
- [Adding an Email Address to the SOA Record](#) on page 560
- [Notifying External Secondary Servers](#) on page 561
- [Specifying Port Settings for DNS](#) on page 562
- [Specifying Minimal Responses](#) on page 565
- [Starting and Stopping the DNS Service](#) on page 565

Configuring DNS Access Control

You can add ACEs (access control entries) or use a named ACL (access control list) to determine which hosts can perform specific DNS tasks. For information about how to define a named ACL, see [Defining Named ACLs](#) on page 307. When you add ACEs or a named ACL to Grid DNS properties, the configuration overrides member and object access control for DNS zone transfers, dynamic DNS updates, DNS queries and recursive queries, blackhole lists, and AAAA filtering. For a full list of operations that support access control, see [Operations that Support Access Control](#) on page 306.

To configure DNS access control:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**, and select one of the following tabs for specific DNS tasks:
 - **Updates** tab: Define ACEs or a named ACL to control Grid level dynamic DNS updates, as described in [Enabling DNS Servers to Accept DDNS Updates](#) on page 706.
 - **Queries** tab: Define ACEs or a named ACL to control Grid level DNS queries, recursive queries, and AAAA filtering, as described in [Controlling DNS Queries](#) on page 570, [Enabling Recursive Queries](#) on page 571, and [Controlling AAAA Records for IPv4 Clients](#) on page 573.
 - **Zone Transfers** tab: Define ACEs or a named ACL to control Grid level DNS zone transfers, as described in [Enabling Zone Transfers](#) on page 583. This does not apply to zone transfers for Microsoft servers. For information about Microsoft servers, see [Setting Zone Properties](#) on page 971.
 - **Blackhole** tab: Configure ACEs or a named ACL to define IP addresses and networks that you do not want to include during the DNS resolution process, as described in [Configuring a DNS Blackhole List](#) on page 590.
 - **DNS64** tab: Configure ACEs or a named ACL for clients to which the appliance sends synthesized AAAA records DNS64 groups, as described in [Setting DNS64 Group Properties](#) on page 597.
3. Save the configuration.

You can override the Grid settings at the member and object levels.

About Time To Live Settings

You can specify TTL (time to live) settings for Infoblox host records and resource records. TTL is the time that a name server is allowed to cache data. After the TTL expires, the name server is required to update the data. Setting a high TTL reduces network traffic, but also renders your cached data less current. Conversely, setting a low TTL renders more current cached data, but also increases the traffic on your network.

You can specify global TTL settings at the Grid level, for individual zones, or resource records. When you configure TTL settings for auto-generated records, the following conditions apply:

- NS records that are auto-generated for delegated name servers use TTL settings from their delegated zones.
- Auto-generated glue A and AAAA records use TTL settings from a delegated zone if the name of the name server is below the delegation point and does not belong to an authoritative child zone.
- All other auto-generated NS, A, and AAAA records continue to use TTL settings from their parent zones.
- Auto-generated PTR records do not inherit TTL settings from delegated zones. They use TTL settings from their parent zones.

When you have an RRSET (resource record set) that contains different TTL settings for each record, Grid Manager displays the actual TTL values for these records. However, in DNS responses, the appliance takes the least of the values and returns that as the TTL setting for all resource records in the RRset.

For recursive DNS servers, you can specify the maximum cache TTL value that establishes the time limit for the name server to cache positive responses. You can also specify the maximum negative cache TTL value that specifies the time limit for the name server to cache negative responses. For information about how to configure these settings, see [Specifying Max Cache TTL and Max Negative Cache TTL Settings](#) on page 559.

Specifying TTL Settings for a Grid

To specify global TTL settings for resource records hosted by Grid members:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the **Basic** tab of the **General** section of the *Grid DNS Properties* editor, modify the following values as necessary:
 - **Refresh:** This interval tells secondary servers how often to send a message to the primary server for a zone to check that their data is current, and retrieve fresh data if it is not. The default is three hours.
 - **Retry:** This interval tells secondary servers how long to wait before attempting to recontact the primary server after a connection failure between the two occurs. The default is one hour.
 - **Expire:** If the secondary fails to contact the primary for the specified interval, the secondary stops giving out answers about the zone because the zone data is too old to be useful. The default is 30 days.
 - **Default TTL:** Specifies how long name servers can cache the data. The default is eight hours.
 - **Negative-caching TTL (Time to Live):** Specifies how long name servers can cache negative responses, such as NXDOMAIN responses. The default is 15 minutes.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying TTL Settings for a Zone

To specify TTL settings for host and resource records in a zone:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* -> *zone* check box -> Edit icon.
2. In the *Authoritative Zone* editor, click **Settings**.
3. Click **Override** and complete the fields as described in the preceding section, [Specifying TTL Settings for a Grid](#).

Specifying the TTL of a Host or Resource Record

To specify the TTL setting for an Infoblox host or resource record:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* -> *zone* -> *resource_record*.
2. The **TTL** tab of the resource record editor displays the TTL setting the resource record inherited from the Grid or zone. Click **Override** and enter a value. The setting is in hours by default. You can change it to seconds, minutes, days or weeks.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying TTL Settings for a Lame Server

Servers that are marked as authoritative, but do not respond as authoritative servers are called lame servers. You can specify the number of seconds to cache a lame delegation or lame server indication through the Lame TTL option. Lame TTL usually indicates the amount of time your name server remembers information about the remote name server that is not authoritative for a zone, which is delegated to it.

A domain or sub-domain that is delegated to a server that is not authoritative for the domain is called lame delegation. It indicates that a zone file does not exist for the domain on the server.

The lame time-to-live cache value can be defined at the Grid DNS, Member DNS, or DNS view level.

To specify the Lame TTL cache value for a lame delegation or lame server:

1. **Grid:** From the **Grid** tab -> **Grid Manager** tab, select the **DNS** tab, click the **Services** tab -> *member* check box, expand the **Toolbar** and click **Edit** -> **Grid DNS Properties**. In the *Grid DNS Properties* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**).

Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> *Edit* icon. In the *Member DNS Properties* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**).

DNS View: From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns_view* check box -> *Edit* icon. In the *DNS View* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**).

2. In the *Grid DNS Properties*, *Member DNS Properties*, or the *DNS View* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**) and then complete the following:

Lame TTL: Specify the duration of time to cache a lame delegation or lame server. The default value is 600 seconds (ten minutes) and the maximum value is 1800 seconds (thirty minutes). The appliance displays a warning message when you specify a value equal to 0 (zero). The value 0 (zero) disables lame caching and is not recommended. The appliance displays an error message when you specify a value greater than 1800 seconds.

The **Lame TTL** cache value is inherited from the Grid by the member and DNS view levels and this field is disabled, by default. To override the **Lame TTL** cache value, click **Override**. You can override the value at the member and DNS view levels. To retain the same **Lame TTL** value as the Grid, click **Inherit** at the member and DNS view level.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying Max Cache TTL and Max Negative Cache TTL Settings

You can specify the maximum duration of time for which your name server caches positive responses using the Max Cache TTL settings. The Max Cache TTL indicates the time limit for which the name server retains records in the cache. When the Max Cache TTL for a record expires, the name server deletes the record from the cache.

You can also specify the maximum duration of time for which your name server caches negative responses through the Max Negative Cache TTL settings. The Max Negative Cache TTL sets the time limit for which the name server retains negative responses (NXDOMAIN/NXRRSET responses) in the cache. The name server deletes a negative response from the cache when the Max Negative Cache TTL period for the entry expires.

You can define the Max Cache TTL value and the Max Negative Cache TTL value at the Grid DNS, Member DNS, and DNS view levels.

To specify the Max Cache TTL and the Max Negative Cache TTL:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar** and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> *Edit* icon.
DNS View: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* check box -> *Edit* icon.
2. In the *Grid DNS Properties*, *Member DNS Properties*, or the *DNS View* editor, click **Toggle Advanced Mode** if the editor is in the basic mode.

3. Click the **Advanced** subtab of the **General** tab and then complete the following:
 - **Max Cache TTL:** Specify the maximum duration of time for which the name server caches positive responses. Select the time period in minutes, hours, or days from the drop-down list. The default value is one week (7 days), and the maximum value is 49710 days, 1193046 hours, or 71582788 minutes. The appliance displays an error message when you enter a value greater than the maximum value. Note that setting the Max Cache TTL value to 0 (zero) will disable the name server from caching any data, and it is not recommended.
 - **Max Negative Cache TTL:** Specify the maximum duration of time for which the name server caches negative responses. Select the time period in minutes, hours, or days from the drop-down list. The default value is three hours, and the maximum value is 7 days, 168 hours, or 10080 minutes. The appliance displays an error message when you enter a value greater than the maximum value. Note that setting the Max Negative Cache TTL value to 0 (zero) will disable the name server from caching negative responses, and it is not recommended.

The Max Cache TTL value and the Max Negative Cache TTL value are inherited from the Grid at the member and DNS view levels. To override the inherited values, click **Override** and specify the new value. To retain the Grid values, click **Inherit**.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Adding an Email Address to the SOA Record

If the primary name server of a zone is a Grid member, you can add an administrator email address to the SOA record to help admins determine who to contact about this zone.

Adding an Email Address for SOA Records in the Grid

If all zones hosted by the Grid members have the same administrator, you can add the email address once for the Grid. The appliance then adds the email address to the RNAME field of the SOA records of the zones.

To add an email address to the SOA records at the Grid level:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the **General** -> **Basic** tab of the *Grid DNS Properties* editor, enter the email address in the **E-mail Address (for SOA RNAME field)** field.

Note: The appliance does not support IDN for the **E-mail Address (for SOA RNAME field)** field at the Grid level. You can add an email address containing IDN for the SOA records at the zone level.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Adding an Email Address for the Zone SOA Record

To add an email address to the SOA record of a zone:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns_view*-> *zone* check box -> Edit icon.
2. In the *Authoritative Zone* editor, click **Settings**.
3. Click **Override** beside the **Email address (for SOA RNAME field)** field and enter the email address of the zone administrator.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: The appliance supports IDN for the host name of the **Email address (for SOA RNAME field)** field. For example, you can create `admin@инфоблокс.рф` but not `админ@инфоблокс.рф.com`.

Notifying External Secondary Servers

Grid members can use database replication to maintain up-to-date zone data sets, so the secondary servers in the Grid can keep their zone data synchronized even if the primary server fails. Any external secondary servers can fall out of sync, however, if they rely only on the primary server to send notify messages when there is new zone data. Therefore all authoritative name servers in a Grid (all primary and secondary servers) send notify messages to external secondary servers by default. This ensures that an external secondary name server receives notify messages when its master is a secondary name server in a Grid. However, it also increases the number of notify messages.

Infoblox recommends that you do not configure a large number of external secondary servers in stealth mode. To ensure that these secondary servers receive notifications about zone updates, you can allow zone transfers for these IP addresses and then enable the appliance to add them to the also-notify statement. For information about how to configure this feature, see [Configuring Zone Transfers](#) on page 584.

To specify whether secondary name servers in the Grid are to send notify messages to external secondary name servers:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. Complete the following:
 - **Enable Grid secondaries to notify external secondaries:** This option is enabled by default.
 - **Notify Delay:** Specify the number of seconds that the Grid secondary servers delays sending notification messages to the external secondaries. The default is five seconds.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

For the external secondary servers to accept notify messages from the secondary name servers in the Grid and then request zone transfers from them, you must configure the external secondary servers to use the Grid secondary servers as the source of the zone transfers. This ensures that the external secondary servers continue to receive notify messages, even if the primary server is unavailable.

Enabling the Configuration of RRset Orders

You can use the Infoblox GUI to configure the order that the appliance uses to return the A and AAAA records associated with an Infoblox host. This feature is useful when you want the appliance to return the A and AAAA records of a host in a specific order. For example, if you want the management address to appear first on a list of IP addresses associated with a network device, you can configure the order of the IP addresses so the management address is always returned first on the list when you look up the name of the device. For information about using the Infoblox API to configure RRset order (resource record order) of a host, refer to the *Infoblox API Documentation*.

To specify an RRset order of a host record, you must first enable the feature at the Grid level. When you enable this feature and there are multiple IP addresses associated with the host record, you can specify one of the following RRset orders through the *Host Record* wizard and editor:

- **Fixed:** The A and AAAA records of the host are returned in the order that you specify in the IPv4 and IPv6 address tables.
- **Random:** The A and AAAA records of the host are returned in a random order.
- **Cyclic:** The A and AAAA records are returned in a round robin pattern.

For information about specifying RRset order of a host record, see [Adding Host Records](#) on page 462.

Note that when you configure an order type for the IP addresses associated with a host record, the order type applies to both the A and AAAA records of the host. It does not apply to any non-host A or AAAA records that may have the same owner name as the host record. By default, the appliance returns resource records in a cyclic or round robin order. The return order of non-authoritative data retrieved from a recursion is not affected by the host RRset order, and that remains cyclic.

To enable the configuration of RRset order:

1. From the **Data Management** tab -> **DNS** tab, expand the Toolbar, and then click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. Complete the following:
 - **Enable setting RRset order for hosts with multiple addresses:** Select this check box to enable the configuration of RRset order for a host record. After you enable this feature, you can configure the RRset order in the *Host Record* wizard or editor. For information, see [Adding Host Records](#) on page 462.
 - **Preserve host RRset order for Grid secondaries that use DNS zone transfers:** This is enabled only when you have enabled the setting of RRset order for host records. When you select this check box, the RRset order that you configure for a host record applies to the resource records of the Grid secondaries that are in the DNS transfer mode.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying Port Settings for DNS

When requesting zone transfers from the primary server, some secondary DNS servers use the source port number (the primary server used to send the notify message) as the destination port number in the zone transfer request. If the primary server uses a random source port number when sending the notify message—that the secondary server then uses as the destination port number when requesting a zone transfer—zone transfers can fail if there is an intervening firewall blocking traffic to the destination port number.

Specifying a source port number for recursive queries ensures that a firewall allows the response. If you do not specify a source port number, the NIOS appliance sends these messages from a random port number.

When performing recursive queries, the NIOS appliance uses a random source port number above 1024 by default. The queried server responds using the source port number in the query as the destination port number in its response. If there is an intervening firewall that does not perform stateful inspection and blocks incoming traffic to the destination port number, the recursive query fails.

You can specify a source port number for notify messages to ensure the firewall allows the zone transfer request from the secondary server to the primary server. If you do not specify a source port number, the NIOS appliance sends messages from a random port number above 1024.

You can limit If you have configured anycast and non-anycast IP addresses on the loopback interface, you must enable the appliance to provide DNS services on them. You can also configure the appliance to listen for DNS queries on a specific IP address that you configure on the loopback interface, by separating the source port for DNS queries from the port for notify messages and zone transfer requests. For information about the loopback interface and anycast addressing, see [Using the Loopback Interface](#) on page 758.

Specifying Source Ports

To specify port numbers and settings for queries, notify messages and zone transfer requests:

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. You can change the port settings as follows:
 - **Listen on these additional IP addresses:** Click the Add icon to add an anycast or non-anycast address you configure on the loopback interface. You must add all IP addresses you configure on the loopback interface so the appliance can provide DNS services on them. Adding source ports for listening supports both IPv4 and IPv6 interfaces. For information about adding IP addresses on the loopback interface, see [Using the Loopback Interface](#) on page 758.
 - **Send queries from:** If you want to improve the DNS service performance, you can separate the DNS queries from the notify messages and zone transfer requests. From the drop-down list, select the source port of the DNS queries that the Grid member sends. The appliance lists all physical interfaces and the non-anycast IPv4 and IPv6 addresses you configure on the loopback interface. For more information, see [Configuring IP Addresses on the Loopback Interface](#) on page 759.
 - **Send notify messages and zone transfer requests from:** From the drop-down list, select the source port of the notify messages and zone transfer requests that the Grid member sends. You can only select a physical interface, you cannot select IP addresses on the loopback interface.
 - **Notify Delay:** Specify the number of seconds that the Grid secondary servers delays sending notification messages to the external secondaries. The default is five seconds.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying Static Source Ports

To specify static source ports:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Complete the following:
 - **Set static source UDP port for queries (not recommended):** This is disabled by default. To override the value that has been inherited from the Grid, click **Override**. Select this check box to enable it and enter the UDP port number. To retain the same value as the Grid, click **Inherit**.
 - **Set static source UDP port for notify messages:** This is disabled by default. To override the value that has been inherited from the Grid, click **Override**. Select this check box to specify a source port for notify messages to ensure that the firewall allows the zone transfer request from the secondary server to the primary server. If you do not specify a source port, the appliance sends messages from a random port with a number above 1024. To retain the same value as the Grid, click **Inherit**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Using Extension Mechanisms for DNS (EDNS0)

The NIOS appliance supports EDNS0 (Extension Mechanisms for DNS), which allows DNS clients to expand and advertise up to 4096 bytes of UDP packets for certain DNS parameters. EDNS0 facilitates the transfer of UDP packets beyond the original restricted packet size of 512 bytes. As defined in RFC 2671, EDNS0 provides extended UDP packet size that supports additional DNS functionality, such as DNSSEC. When EDNS0 is supported, the DNS client adds information to the additional data section of a DNS request in the form of an OPT pseudo-RR (resource record). An OPT RR does not contain actual DNS data and its contents pertain to the UDP transport layer message only. An OPT RR is not cached, forwarded, or stored. For more information about EDNS0, refer to *RFC 2671 Extension Mechanisms for DNS (EDNS0)*.

EDNS0 is enabled on the NIOS appliance by default, which means all outgoing recursive queries are set to have a maximum UDP packet size of 4096 bytes. Typically, when the appliance receives a DNS request that contains an OPT RR, it assumes the DNS client supports EDNS0 and thus scales its response accordingly. When the appliance is used as a forwarder or a resolver for recursive queries and communicates with a client that does not support EDNS0, the appliance sends three queries starting with one that contains EDNS0 and DNSSEC support messages and is set to a maximum UDP packet size of 4096 bytes. When the first query fails, the appliance sends another query that contains only the EDNS0 support message. If the second attempt fails too, the appliance sends a third query that indicates a standard 512-byte query. Note that when EDNS0 is not used, DNS packets may be sent over TCP. For DNS service to function properly at this stage, ensure that you configure your firewall accordingly.

The following information demonstrates how the appliance responds when EDNS0 is enabled by default and the end server does not support EDNS0:

Packet 0954: 08:19:38.925 - query for www.google.com from Infoblox to forwarder (with EDNS0 support by setting the Extended Label Type to '01' and DNSSEC OK bit to '1')

Packet 1138: 08:19:47.927 - query for www.google.com from Infoblox to forwarder (with EDNS0 support by setting the Extended Label Type to '01' and DNSSEC OK bit to '0')

Packet 1504: 08:19:58.929 - query for www.google.com from Infoblox to forwarder (without EDNS0 and DNSSEC support by sending a standard 512-byte query)

Packet 1505: 08:19:30.960 - query response for www.google.com from forwarder to Infoblox

To ensure that end servers that do not support EDNS0 can respond to recursive queries from the NIOS appliance and to improve DNS performance, you can disable EDNS0 for the Grid and override the Grid settings for individual members. Note that you cannot configure the maximum UDP packet size, which is set for 4096 bytes by default. When you disable EDNS0, the appliance does not include OPT RRs for all outgoing recursive DNS queries. Thus remote end servers that do not support EDNS0 can still respond to the queries. This feature is useful when your NIOS appliance is used as a forwarder or a resolver for recursive queries, and the end servers in the configuration do not support EDNS0.

WARNING: WHEN YOU DISABLE EDNS0, ALL OUTGOING DNSSEC QUERIES TO ZONES WITHIN TRUSTED ANCHORS WILL FAIL EVEN IF DNSSEC VALIDATION IS ENABLED. THIS IS DUE TO THE RESTRICTION OF THE UDP PACKET LENGTH WHEN YOU DISABLE EDNS0. FOR INFORMATION ABOUT DNSSEC, SEE [ABOUT DNSSEC](#) ON PAGE 734.

To disable EDNS0:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **General** tab -> **Advanced** tab, and complete the following:
 - **Disable EDNS0:** This check box is deselected and EDNS0 is enabled by default. To override the value inherited from the Grid, click **Override**. To retain the same value as the Grid, click **Inherit**. Select this check box to disable EDNS0. When you disable EDNS0, the appliance does not include OPT RRs for all outgoing recursive DNS queries and all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying Minimal Responses

A NIOS appliance returns a minimal amount of data in response to a query, by default. It includes records in the authority and additional data sections of its response only when required, such as in negative responses. This feature speeds up the DNS services provided by the appliance.

To disable returning minimal responses:

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
2. In the *Member DNS Configuration* editor, click **General** -> **Basic** tab.
3. Clear the **Return minimal responses** check box
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Starting and Stopping the DNS Service

The DNS service is disabled by default. After you complete the DNS configuration, you can start DNS service on a member. You can also disable the DNS service on any Grid member. Be aware that disabling the DNS service on a member removes the NS records from it. If you later re-enable DNS service for this member, the NS records are then restored.

To start DNS service on a member:

1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* check box.
2. Expand the Toolbar and click **Start**.
3. In the *Start Member DNS Service* dialog box, click **Yes**.

Grid Manager starts the DNS service on the selected member.

You can stop DNS service on a member by selecting the member check box and click **Stop** from the Toolbar.

ABOUT DNS CACHE

The NIOS appliance allows you to clear certain information from the DNS cache. You can do the following:

- [Clearing DNS Cache](#) on page 565
- [Clearing Cache for DNS Views](#) on page 566
- [Clearing Domain Names from Cache](#) on page 566

Clearing DNS Cache

You can clear all the entries that are saved in the DNS cache. When you clear DNS cache on the NIOS appliance, entire BIND recursive cache is cleared.

To clear DNS cache:

1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* check box.
2. Expand the Toolbar, click **Clear** -> **Clear DNS Cache**.
3. Click **Yes** in the confirmation dialog box to clear DNS cache.

Clearing Cache for DNS Views

You can configure the NIOS appliance to clear cache of a specific DNS view. This feature clears cache entries of a specific DNS view that is associated with the selected member.

To clear cache of a DNS view:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **Clear** -> **Clear View's Cache**.
3. Specify the following in the *Clear View's Cache* dialog box:
 - **Member:** Click **Select Member** to select a member. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
 - **DNS View:** Select a **DNS View** from the drop-down list. This list box appears only when there are multiple DNS views in the network view.
 - Click **Clear Cache** to clear the cache entries of the corresponding DNS View.

Note: The entire name server recursive cache is cleared, if you do not specify a DNS view when you clear cache using **Clear View's Cache** and **Clear Domain Name** features on the NIOS appliance.

Clearing Domain Names from Cache

You can clear specific domain names from the DNS cache. The selected domain name entry is deleted from the BIND recursive cache.

To clear a domain name:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **Clear** -> **Clear Domain Name**.
3. Specify the following in the *Clear Domain Name from Cache* dialog box:
 - **Domain Name:** Enter a domain name you want to delete.
 - **Member:** Click **Select Member** to select a member. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
 - **DNS View:** Select a **DNS View** from the drop-down list. This list box appears only when there are multiple DNS views in the network view.
 - Click **Clear Domain Name** to clear the domain name from the cache.

VIEWING DNS CACHE ENTRIES

The NIOS appliance allows you to view certain information that is stored in the DNS cache. You can do the following:

- [Viewing DNS Configuration](#) on page 567
- [Viewing DNS Cache Details](#) on page 568
- [Viewing Statistics](#) on page 568

Viewing DNS Configuration

The NIOS appliance supports **View Configuration** feature that enables you to view DNS configuration details. You can view the configuration details through a browser.

To view DNS configuration:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **View** -> **View Configuration**.

Sample Output

```
include "/infoblox/var/named_conf/tsig.key";

options {
zone-statistics yes;
directory "/infoblox/var/named_conf";
version none;
hostname none;
recursion yes;
listen-on { 127.0.0.1; 10.34.1.18; };
query-source address 10.34.1.18 port *;
notify-source 10.34.1.18 port *;
transfer-source 10.34.1.18;
minimal-responses yes;
max-cache-size 536870912;
infoblox-top-query yes;
infoblox-top-query-log-interval 60;
infoblox-top-query-client 500;
infoblox-top-query-name 500;
infoblox-top-query-rr-type 500;
infoblox-top-query-nxdomain 500;
infoblox-top-query-servfail 500;
infoblox-top-query-rpz 99;
infoblox-top-query-rpz-items-per-client 100;
    lame-ttl 600;
# for service restart: allow_bulkhost_ddns = Refusal
allow-transfer { any; };
forwarders { 10.32.0.177; };
avoid-v4-udp-ports { 2114; 2113; 2115; 8000; 8089; 9997; 2222; 7911; 7912; 8000; 8089; 9997;
8080; 9000; 9999; 9004; 2022; 3374; 3115; 1194; };
transfer-format many-answers;
};

# Worker threads: default

# Bulk Host Name Templates:
#Four Octets: "-$1-$2-$3-$4" (Default)
#One Octet: "-$4"
#Three Octets: "-$2-$3-$4"
#Two Octets: "-$3-$4"

include "/infoblox/var/named_conf/dhcp_updater.key";
```

```
include "/infoblox/var/named_conf/rndc.key";

controls {
    inet 127.0.0.1 port 953
    allow { 127.0.0.1; } keys { "rndc-key"; };
};

logging {
    channel ib_syslog {
        syslog daemon;
        severity info;
    };
    category default { ib_syslog; };
    category rpz { null; };
};

acl all_dns_views_updater_keys { key DHCP_UPDATER_default; key DHCP_UPDATER1; key
DHCP_UPDATER3; };
```

Viewing DNS Cache Details

You can view data stored in cache for the DNS views that are configured in the NIOS appliance. You can view the details through a browser.

To view cache details:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **View** -> **View Cache**.

Sample Output

```
;; Start view _default
;;; Cache dump of view '_default' (cache _default)
;$DATE 20121018180555
; authanswer
a.test.com.23876IN A4.4.4.4
;; Address database dump
;; Dump complete
```

Viewing Statistics

The View Statistics feature enables you to view DNS Statistics of a Grid member. You can view statistics through a browser.

To view statistics:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **View** -> **View Cache**.

You can view statistics in the *DNS Statistics for Member* dialog box.

USING FORWARDERS

A forwarder is essentially a name server to which all other name servers first send queries that they cannot resolve locally. The forwarder then sends these queries to DNS servers that are external to the network, avoiding the need for the other name servers in your network to send queries off-site. A forwarder eventually builds up a cache of information, which it uses to resolve queries. This reduces Internet traffic over the network and decreases the response time to DNS clients. This is useful in organizations that need to minimize off-site traffic, such as a remote office with a slow connection to a company's network.

You can select any Grid member to function as a forwarder. You must configure your firewall to allow that Grid member to communicate with external DNS servers. You can also configure the NIOS appliance to send queries to one or more forwarders.

You can define a list of forwarders for the entire Grid, for each Grid member, or for each DNS view.

Specifying Forwarders

To configure forwarders for a Grid, member, or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns_view* check box -> Edit icon.
 Note that if there is only one DNS view—for example, the predefined default view—you can just click the Edit icon beside it.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Click the **Forwarders** tab.
3. Click the Add icon.
4. Enter an IP address in the text field. The field supports entry for both IPv4 and IPv6 values.
 - To remove a forwarder, select the IP address from the Forwarders list, and then click the Delete icon.
 - To move a forwarder up or down on the list, select it and click the Up or Down arrow.
5. To use only forwarders on your network (and not root servers), select the **Use Forwarders Only** check box.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

CONTROLLING DNS QUERIES

By default, the NIOS appliance responds to DNS queries from any IP address. You can create a list of queriers to which the appliance is allowed to respond; restricting it to specific networks, IP addresses, and remote servers that present specified TSIG (transaction signature) keys. When using TSIG keys, it is important that the appliances and servers involved with the authentication procedure use NTP (Network Time Protocol) for their time settings (see [Using NTP for Time Settings](#) on page 313).

In addition, you can also configure the appliance to respond to recursive queries. A recursive query requires the appliance to return requested DNS data, or locate the data through queries to other servers. Recursion is disabled by default. If you enable this feature, you can also create a list of allowed recursive queriers. For information about allowing recursion, refer to [Enabling Recursive Queries](#) on page 571.

You can create a list of allowed queriers for the Grid and for individual Grid members.

Specifying Queriers

To configure a list of allowed queriers for the Grid or for a member:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode**, select the **Queries** tab.
3. In the Allow queries from section, select one of the following:
 - **None:** Select this if you do not want to configure access control for DNS queries. The appliance allows queries from all clients. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows clients that have the **Allow** permission to send and receive DNS queries. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the remote querier. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server. This name must match the name of the same TSIG key on other name servers.
 - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.

- **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
- **Any Address/Network:** Select to allow or deny queries from any IP addresses.
After you have added access control entries, you can do the following:
 - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

4. Save the configuration.

ENABLING RECURSIVE QUERIES

You can enable the appliance to respond to recursive queries and create a list of allowed networks, IP addresses, and remote servers that present specified TSIG (transaction signature) keys. When using TSIG keys, it is important that the appliances and servers involved with the authentication procedure use NTP (Network Time Protocol) for their time settings (see [Using NTP for Time Settings](#) on page 313).

A recursive query requires the appliance to return requested DNS data, or locate the data through queries to other servers. When a NIOS appliance receives a query for DNS data it does not have and you have enabled recursive queries, it first sends a query to any specified forwarders. If a forwarder does not respond (and you have disabled the **Use Forwarders Only** option in the **Forwarders** tab of the *Member DNS Properties* editor), the appliance sends a non-recursive query to specified internal root servers. If no internal root servers are configured, the appliance sends a non-recursive query to the Internet root servers. For information on specifying root name servers, see [About Root Name Servers](#) on page 587.

You can enable recursion for a Grid, individual Grid members, and DNS views. For information about enabling recursion in a DNS view, see [Configuring DNS Views](#) on page 604.

Enabling Recursion

To enable recursion and create a list of recursive queriers:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab → **Members** tab → *member* check box → Edit icon.
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode**, select the **Queries** tab.
3. Click **Allow recursion**, and then in the *Allow recursive queries from* section, select one of the following:
 - **None:** Select this if you do not want to configure access control for recursive queries. When you select **None**, the appliance allows recursive queries from all clients. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows clients that have the **Allow** permission to send and receive recursive DNS queries. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.

- **IPv4 Address** and **IPv6 Address**: Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the remote querier. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network**: In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address**: Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission**: Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network**: In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address**: Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission**: Select **Allow** or **Deny** from the drop-down list.
 - **TSIG Key**: In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name**: Enter a meaningful name for the key, such as a zone name or the name of the remote name server. This name must match the name of the same TSIG key on other name servers.
 - **Key Algorithm**: Select either **HMAC-MD5** or **HMAC-SHA256**.
 - **Key Data**: To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
 - **Any Address/Network**: Select to allow or deny queries from any IP addresses.
- After you have added access control entries, you can do the following:
- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

4. Save the configuration.

Restricting Recursive Clients

By default, the appliance is allowed to serve up to 1,000 concurrent clients that send recursive queries. You can change this default according to your business needs.

From the **Data Management** tab, select the **DNS** tab and click the **Members** tab → *member* check box → Edit icon.

1. In the *Member DNS Properties* editor, click **Toggle Advanced Mode**.
2. When the additional tabs appear, click the **Advanced** subtab of the **Queries** tab.
3. Select the **Limit number of recursive clients to** option and enter a number.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

CONTROLLING AAAA RECORDS FOR IPv4 CLIENTS

By default, the NIOS appliance returns resource records, including AAAA records, in response to DNS queries. You can enable the appliance to filter and remove AAAA records in response to queries received over IPv4 for each name server and DNS view. This feature is useful in a configuration where a client issues a DNS query over IPv4 when it does not have the ability to use an IPv6 address. When a response returns an IPv6 address however, the client that sends the query over an IPv4 transport would lose connectivity. By enabling AAAA filtering, you can configure your name server not to return AAAA records to clients that request queries over an IPv4 transport. Presumably, these clients then re-query the name server for A records for the same domain name.

Depending on your configuration, the appliance can remove AAAA records for all queries over IPv4 (even when DNSSEC is enabled), or only for queries that are not DNSSEC-signed. You can also create a list of IPv4 networks and addresses to which the appliance applies AAAA filtering and vice versa. You can enable and configure AAAA filtering for the Grid, members, and DNS views.

To control whether you want the appliance to return AAAA records for queries sent over IPv4, you must first enable AAAA filtering, and then create a list of IPv4 networks and addresses that allow or deny AAAA filtering from the appliance, as described in [Enabling AAAA Filtering](#).

Note: An AAAA record is filtered only when there is also an A record for the same domain name. In this case, the appliance still sends a response, but without any AAAA or A record in it. When a client queries for an AAAA record and there is no corresponding A record for it, the appliance returns the AAAA record even if you have enabled AAAA filtering for this client.

Enabling AAAA Filtering

To enable AAAA filtering and configure a list of IPv4 networks and addresses:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, click the **DNS** tab -> **Zones** tab -> *dns_view* check box -> Edit icon.
 To override the Grid settings, click **Override** and complete the appropriate fields.

2. In the editor, click the **Queries** tab and complete the following:

Enable AAAA Filtering: From the drop-down list, select one of the following:

- **Break DNSSEC:** Select this to remove AAAA records in response to queries sent over IPv4, including those that are signed by DNSSEC.

Note: Be aware that when you select this option, DNSSEC configuration will no longer be in effect.

- **No:** Select this to disable AAAA filtering for queries over IPv4. When you select this, the appliance returns AAAA records in response to all DNS queries issued over IPv4. This is selected by default.
 - **Yes:** Select this to enable AAAA filtering for queries over IPv4. When you select this, the appliance removes AAAA records in response to all DNS queries issued over IPv4, except for DNSSEC-signed requests.
3. In the AAAA Filtering section, select one of the following:
 - **None:** Select this if you do not want to configure access control for AAAA filtering. The appliance allows all clients to issue DNS queries over IPv4 when they do not have the ability to use IPv6 addresses. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows clients that have the **Allow** permission can filter AAAA responses. You can click **Clear** to remove the selected named ACL.

- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address:** Select this to add an IPv4 address. Click the **Value** field and enter the IP address of the client. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list. When you select **Allow**, the appliance applies AAAA filtering and removes AAAA records in response to queries sent by the specified IPv4 address. When you select **Deny**, the appliance does not apply AAAA filtering and thus returns AAAA records.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **Any Address/Network:** Select to allow or deny AAAA filtering from any IP addresses.
- After you have added access control entries, you can do the following:
- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

Note: Note that if you do not enter any addresses or networks in the table, the appliance applies AAAA filtering to all IPv4 clients. In other words, the appliance removes AAAA records in responses to all queries sent over IPv4.

ABOUT NXDOMAIN REDIRECTION

When a DNS member with recursion enabled receives a recursive query for data for which it is not authoritative, it locates the data through queries to other servers. If the query is for a non-existent domain name, the DNS member receives an NXDOMAIN response from the authoritative name server, which the member then forwards to the DNS client. An NXDOMAIN response contains a “Name Error” RCODE, signifying that the domain name referenced in the query does not exist. (For information, you can refer to *RFC 1035, Domain Names — Implementation and Specification*.)

You can install a Query Redirection license on a recursive DNS member to control its response to queries for A records of non-existent domain names and other domain names that you specify. After the license is installed, Grid Manager displays the **NXDOMAIN Rulesets** tab where you can create rules that specify how a DNS member responds to queries for A records for certain domain names and non-existent domain names. Each rule contains a domain name specification, and the action of the DNS member when the domain name in the query matches that in the rule. After you create the rules, you then enable the NXDOMAIN redirection feature and list the IP addresses that are included in the synthesized responses.

Recursive DNS members can redirect responses to queries for A records only. DNS members resolve queries for all other records as they normally would.

In addition, you can enable DNS members to log queries that match rules with an action of “Redirect” or “Modify”. You can view the logs in the *Syslog* viewer. The logs include the queried domain name, source IP address, the pattern of the matched rule, and the name of the corresponding ruleset.

When DNSSEC is enabled on the Infoblox DNS server, it does not redirect DNS clients that request DNSSEC data for a non-existent domain name. Instead, it returns an authenticated negative response in the form of an NSEC or NSEC3 RR. (For information about DNSSEC, see [Chapter 21, DNSSEC](#), on page 733.) If DNSSEC is not enabled, the appliance ignores the request for DNSSEC data and redirects the clients.

You can enable NXDOMAIN redirection at the Grid, member, and DNS view levels. Only recursive DNS servers can redirect DNS clients. Non-recursive DNS members do not redirect DNS clients. For information on enabling recursion on a DNS member, see [Enabling Recursive Queries](#) on page 571.

Note that if both NXDOMAIN redirection and the blacklisting feature are enabled, the DNS member applies the blacklist rulesets before the NXDOMAIN rulesets. For information about blacklisting domain names, see [About Blacklists](#) on page 579.

About NXDOMAIN Rulesets

An NXDOMAIN ruleset is a list of rules that a DNS member uses to determine its response to recursive queries for A records it does not have. Each rule consists of a domain name specification or pattern, and an associated action.

Domain names can contain any printable character. You can use certain metacharacters to create domain name patterns that are used to match the domain names in DNS queries. Pattern matching is case-insensitive. Patterns support the following metacharacters:

- Use the caret character (^) to indicate the beginning of a pattern. For example, **^foo** matches **foo.com** but not **barfoo.com**. The caret character has a special meaning only if it is specified at the beginning of a pattern.
- Use the dollar sign character (\$) to indicate the end of a pattern. The dollar sign character has a special meaning, only if it is specified at the end of the pattern. For example, **.com\$** matches **corp100.com** but not **corp100.com.net**.

When the pattern contains a \$ at the end, NIOS automatically adds a period (.) before the \$. For example, if you enter **.com\$**, NIOS saves it as **.com.\$**. The period indicates that the pattern specifies a complete domain name that ends with the root label.

- Use the asterisk character (*) as a wildcard that can match zero or more characters in one or more labels of a domain name. For example, **xf*oy** matches **xfooy.com**, but not **xfoobary**.

A pattern that contains a single asterisk (*) (or an equivalent expression, such as **“^*\$”**) matches any domain name.

- Use the backslash character (\) with one of the metacharacters (\$, ^, * and \) to remove their special meaning. If \ is followed by any other character, that character is taken as an ordinary character, as if \ is not present. For example, **foo\\bar** matches **foo\bar**, and ***** matches a literal asterisk in a domain name.

No other characters have any special meaning. Note in particular that the period character (“.”) only matches a period used as a separator in a domain name.

The action specifies how the DNS member responds when a domain name in a query matches a pattern. The action can be one of the following: **Pass**, **Modify** or **Redirect**.

- **Pass:** The DNS member resolves the query and forwards the response to the DNS client, even if it is an NXDOMAIN response.
- **Modify:** The DNS member resolves the query and forwards the response to the DNS client, only if it is not an NXDOMAIN response. But if the member receives an NXDOMAIN response, it sends the client a synthesized response that includes predefined IP addresses.
- **Redirect:** The DNS member does not resolve the query. Instead, it sends the client a synthesized response that includes predefined IP addresses.

You can configure multiple rulesets. The DNS member applies the rulesets and their rules in the order in which they’re specified in the configuration. If multiple rulesets contain rules with duplicate patterns, the DNS member applies the first rule it encounters and ignores the other rules.

Examples

The following example illustrates how the appliance applies NXDOMAIN rulesets.

Ruleset 1:

Pattern	Action
a1.corp100.com	PASS
*.corp100.com	REDIRECT

- If the DNS member receives a query for a1.corp100.com, it resolves the query and forwards the response, even if it is an NXDOMAIN response, to the client. Note that if the order of the rules was switched, the DNS client would have been redirected immediately, because the domain name a1.corp100.com matches the *.corp100.com pattern.
- If the DNS member receives a query for b1.corp100.com, the member immediately redirects the DNS client to the specified IP address because the domain name in the query matches the second rule.
- If the DNS member receives a query for b1.corp200.com, it resolves the query because the domain name does not match any rule. If the DNS member receives an A record from an authoritative server, the member forwards the response to the client. However, if the member receives an NXDOMAIN response, it redirects the DNS client to the specified IP address.

In the following example, the rules redirect queries for dotted domain names that do not have “.com” As shown in the example, an explicit PASS rule is required at the end.

Ruleset 2:

Pattern	Action
*.com	PASS
.\$	MODIFY
*	PASS

- If the DNS member receives a query for corp100.com which matches the pattern “*.com”, the member resolves the query and forwards the response, even if it is an NXDOMAIN response, to the client.
- If the DNS member receives a query for corp100.org, which matches the pattern “*.*\$”, the member resolves the query. If the member receives an NXDOMAIN response, it redirects the client to the specified IP address. If the member receives a non-NXDOMAIN response, it forwards the response to the client.
- If the DNS member receives a query for corp200, the member resolves the query and forwards the response to the client.

NXDOMAIN Redirection Guidelines

The following summarizes how a DNS member responds to a query for an A record when the NXDOMAIN feature is enabled:

- If there are no rulesets configured, the DNS member queries other name servers.
 - If the DNS member receives a non-NXDOMAIN response from an authoritative server, it forwards the response to the DNS client.
 - If the DNS member receives an NXDOMAIN response from an authoritative server, it redirect the DNS client.
- If rulesets are configured, the DNS member tries to match the domain name in the query with a domain name in the rules.
 - If the DNS member finds a match, it perform the action specified in the rule.
 - If the action is “Redirect”, the DNS member redirect the DNS client.
 - If the action is “Pass”, the DNS member queries other name servers and forwards the response to the DNS client.
 - If the action is “Modify”, the DNS member queries other name servers. If it receives a non-NXDOMAIN response, it forwards the response to the DNS client; if it receives and NXDOMAIN response, it redirects the DNS client.
 - If the DNS member does not find a match, the DNS member queries other name servers.
 - If the DNS member receives a non-NXDOMAIN response, it forwards the response to the DNS client.
 - If the DNS member receives an NXDOMAIN response from an authoritative server, it redirects the DNS client.

Note that if an A record with a dotted hostname is added to an authoritative zone through a dynamic DNS update, and that A record should actually belong in an existing delegation, the appliance may not redirect a query for that A record according to the Blacklist and NXDOMAIN guidelines.

Configuring NXDOMAIN Redirection

To enable NXDOMAIN redirection and configure its properties:

1. Configure NXDOMAIN rulesets. You can create NXDOMAIN rulesets through Grid Manager, as described in [Creating Rulesets](#). You can also specify the rulesets in a CSV file and import the file to the Grid, as described in [About CSV Import](#) on page 86.
2. Enable this feature and specify the redirection IP addresses, as described in [Enabling NXDOMAIN Redirection](#) on page 578.

Creating Rulesets

To create a ruleset:

1. From the **Data Management** tab -> **DNS** tab -> **NXDOMAIN Rulesets** tab, click the Add icon.
2. In the **NXDOMAIN Ruleset** wizard, complete the following and click **Next**:
 - **Name:** Enter a name for the ruleset.
 - **Comment:** You can enter additional information.
 - **Disable:** You can disable this ruleset for use later on. The appliance ignores disabled rulesets.
3. Click the Add icon to add a rule to the ruleset table.
 - In the **Pattern** column, enter a domain name or pattern, using the guidelines specified in [About NXDOMAIN Rulesets](#).
 - In the **Action** column, select **PASS**, **REDIRECT** or **MODIFY**.
 - In the **Order** column, NIOS automatically displays the number of the entry in the list.

The appliance applies the rules in the order they are listed. You can order the list as follows:

- Use the up and down arrows to move rules up or down on the list.
- Use the go-to-top or go-to-bottom arrow to move a rule to the top or bottom of the list.
- Change the Order number of a rule to move it to the desired location.
- Delete a rule by selecting it and clicking the Delete icon.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing NXDOMAIN Rulesets

To view NXDOMAIN rulesets, navigate to the **Data Management** tab -> **DNS** tab -> **NXDOMAIN Rulesets** tab. The panel lists the configured rulesets and their associated comments. You can also display the Disabled column which indicates which rulesets are disabled. From this panel, you can do the following:

- Add more rulesets, as described in the preceding section, [Creating Rulesets](#).
- Edit a ruleset, by clicking its check box and clicking the Edit icon. You can set the following in the NXDOMAIN Ruleset editor:
 - In the **General Basic** tab, you can change entries in any of the fields.
 - In the **Rules** tab, you can do the following:
 - Add a rule by clicking the add icon and specifying the pattern and action.
 - Change the pattern or action of a rule, by clicking in the appropriate row.
 - Delete a rule by clicking its check box and clicking the Delete icon.
 - Move rules up and down, by using the arrows.
 - In the **Permissions** tab, you can set admin permissions for the ruleset. For information about admin permissions, see [Chapter 4, Managing Administrators](#), on page 149.
- Delete a ruleset, by clicking its check box and clicking the Delete icon.

Enabling NXDOMAIN Redirection

Only DNS members with recursion enabled can support NXDOMAIN redirection.

You can enable this feature at the Grid level, and override it for a member or DNS view with recursion enabled. You must specify at least one IP address as the redirection destination. You can specify different redirection IP addresses and rulesets for each Grid member or DNS view, and you can also define members that do not provide redirection. This is useful when you want to define a set of “opt out” servers for DNS clients that do not want to be redirected.

You can also enable the DNS member to log queries that match rules with an action of “Redirect” or “Modify”. The logs include the queried domain name, source IP address, the pattern of the matched rule, and the name of the corresponding ruleset. The DNS member does not log queries that matched rules with an action of “Pass”.

To enable NXDOMAIN redirection:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. If the *Grid DNS Properties* or *Member DNS Properties* editor is in basic mode, click **Toggle Advanced Mode**.
3. Click **NXDOMAIN** and complete the following:
 - **Enable NXDOMAIN redirection (recursive members only):** Select this option to enable recursive DNS members to synthesize their responses to DNS queries for A records.
 - **Redirect to:** Click the Add icon and enter the IP addresses that the DNS server includes in its synthesized response. You must specify at least one IP address. You can add up to 12 IP addresses.
 - **TTL:** Specify how long the DNS client caches the A record with the redirected IP address.

- **Rulesets:** Click the Add icon to add an NXDOMAIN ruleset. Use the up and down arrows to move rulesets up and down in the list. The appliance applies them in the order they are listed.
 - **Log redirected queries:** Select this check box to log the redirected queries to syslog.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

ABOUT BLACKLISTS

Your organization can prevent customers or employees from accessing certain Internet resources, particularly web sites, by prohibiting a recursive DNS member from resolving queries for domain names that you specify.

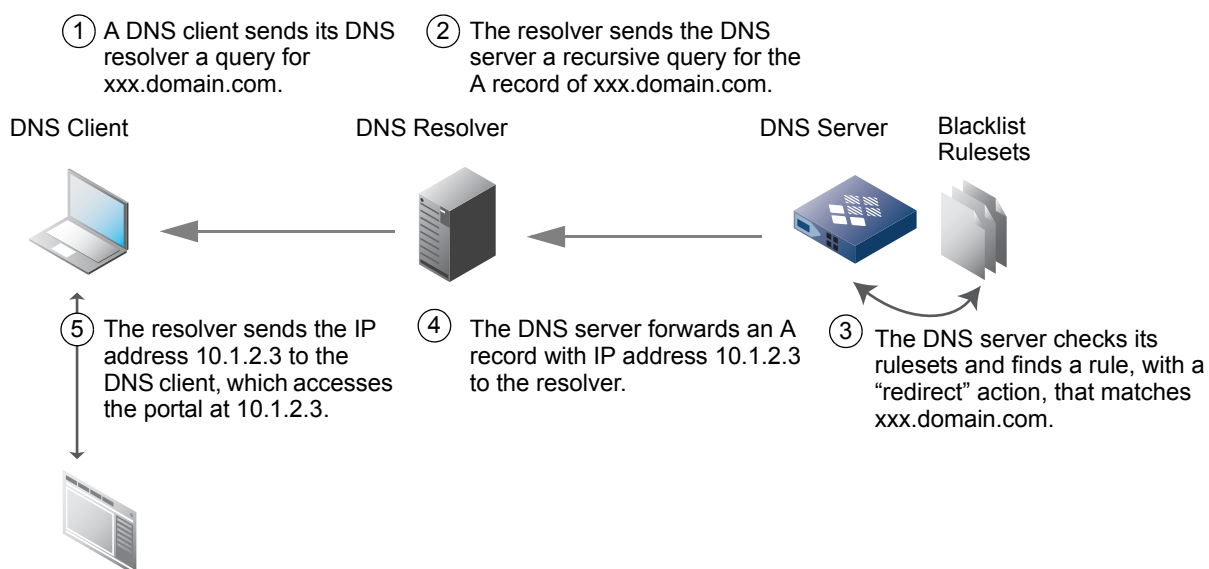
You can create blacklist rules that specify how a DNS member responds to recursive queries for data for which it is not authoritative. Each rule specifies a domain name and the action of the DNS member when the domain name in the query matches that in the rule. Instead of resolving the query, the DNS member can redirect the DNS client to predefined IP addresses or return a REFUSED response code indicating that resolution is not performed because of local policy.

When the DNS member receives a query for data for which it is not authoritative, it first tries to match the domain name in the query with a domain name in any of its rules. If it finds a match, it responds according to the action specified in the rule. If it does not find a match and the NXDOMAIN feature is enabled, the DNS member checks the NXDOMAIN rulesets for a match and responds accordingly. If the NXDOMAIN feature is not enabled, the DNS member resolves the query. (For information about the NXDOMAIN feature, see [About NXDOMAIN Redirection](#) on page 574.

Infoblox DNS members can modify their responses to queries for A records only. Therefore, if the matched query is for a record other than an A record, including a query with a type of “ANY”, the DNS member sends a REFUSED response if the matched rule has an action of “Redirect”.

In [Figure 16.1](#), a DNS client opens a web browser and tries to access xxx.domain.com. When the DNS member receives the query for xxx.domain.com, it checks its blacklist rulesets and finds xxx.domain.com in a rule with an action of “Redirect”. The DNS client is redirected to the configured redirection destination IP address 10.1.2.3.

Figure 16.1 Blacklist



This feature supports queries for data in IPv4 and IPv6 reverse-mapping zones, as well as forward-mapping zones. Note that when a user with a Windows DNS client with IPv6 installed tries to access a domain name, the Windows client sends queries for AAAA records before queries for A records. After the DNS member sends a Refused response to the query for the AAAA record, the DNS client then sends a query for the A record. The DNS member then responds according to the blacklist rules.

When DNSSEC is enabled on the Infoblox DNS server, it does not redirect DNS clients that request DNSSEC data. (For information about DNSSEC, see [Chapter 21, DNSSEC](#), on page 733.) If DNSSEC is not enabled and the query includes a request for DNS data, the appliance ignores the request for DNSSEC data and redirects the clients.

You can enable the blacklist feature at the Grid, member, and DNS view levels. Note that only recursive DNS servers can support this feature. For information on enabling recursion on a DNS member, see [Enabling Recursive Queries](#) on page 571.

About Blacklist Rulesets

A blacklist ruleset is a list of rules that a DNS member uses to determine its response to recursive queries for certain domain names. When you enable the blacklist feature, you must define at least one rule in a ruleset. Each rule consists of a domain name and an associated action. The DNS member matches the domain names in the rules with the entire domain name in the query, including its suffix. The domain name in the rule can contain any printable character. Domain name matching is case-insensitive. Unlike the NXDOMAIN rules, blacklist rules do not support metacharacters in domain names.

The action in a rule is either “Pass” or “Redirect”.

- **Pass:** The DNS member resolves the query and forwards the response to the DNS client.
- **Redirect:** The DNS member does not resolve the query. The DNS member redirects the client to the predefined IP addresses or sends a REFUSED response, depending on your configuration. Note that the DNS member can redirect the client only if the query is for an A record. If the query is for another resource record, the DNS member sends a REFUSED response.

You can use the Blacklist wizard, described in [Adding a Blacklist Ruleset](#), to add blacklist rulesets, but not rules. You can only add rules by importing them in a CSV file, as described in [About CSV Import](#) on page 86. Note that if a blacklist ruleset contains duplicate domain names, the DNS member loads the first rule in the ruleset and discards the other rules.

The following example illustrates how the DNS member applies blacklist rules.

Ruleset 1:

Pattern	Action
a1.foo.com	PASS
foo.com	REDIRECT/BLOCK

- If the DNS member receives a recursive query for a1.foo.com, it resolves the query and forwards the response to the client.
- If the DNS member receives a recursive query for the A record of b1.foo.com, it redirects the DNS client to the specified IP address. If the query is for another record type, such as an MX record, the member sends a REFUSED response to the client.

Blacklist Guidelines

The following summarizes how a DNS member responds to a DNS client when the blacklist feature is enabled:

- If the domain name in the query matches a domain name in a rule, the member does the following:
 - If the query is for an A record, the member performs the action specified in the rule.
 - If the action is “Redirect”, the member performs the action specified in the Blacklist wizard.
 - If the action in the wizard is to redirect, the DNS member redirects the client to the listed IP addresses.
 - If the action in the wizard is to return a REFUSED response, the DNS member sends a REFUSED response to the DNS client.
 - If the action in the rule is “Pass”, the DNS member resolves the query and forwards the response to the DNS client.
 - If the query is for a non-A record, the member performs the action in the rule as follows:
 - If the action is “Redirect”, the DNS member returns a REFUSED response to the DNS client.
 - If the action is “Pass”, the DNS member resolves the query and forwards the response to the DNS client.
- If the domain name in the query does not match a domain name in a rule:
 - If the NXDOMAIN feature is enabled, the DNS member tries to find a match with the NXDOMAIN rules and responds accordingly.
 - If the NXDOMAIN feature is disabled, the DNS member resolves the query and forwards the response to the DNS client.

Note that if an A record with a dotted hostname is added to an authoritative zone through a dynamic DNS update, and that A record should actually belong in an existing delegation, the appliance may not redirect a query for that A record according to the Blacklist and NXDOMAIN guidelines.

Configuring the Blacklist Feature

To configure the blacklist feature:

1. Add blacklist rulesets, as described in [Adding a Blacklist Ruleset](#) on page 581.
2. Create one or more CSV files that contain the rules for each ruleset and import the files to the Grid. For information about importing CSV files, see [About CSV Import](#) on page 86.
3. Enable blacklisting, as described in [Enabling Blacklisting](#) on page 582.

Adding a Blacklist Ruleset

To add the name of a blacklist ruleset:

1. From the **Data Management** tab -> **DNS** tab -> **Blacklist Rulesets** tab, click the Add icon.
2. In the *Blacklist* wizard, complete the following:
 - **Name:** Enter a name for the ruleset.
 - **Comment:** You can enter additional information.
 - **Disable:** You can disable this ruleset for use later on. The appliance ignores disabled rulesets.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can then use the CSV Import feature to import the rules for each ruleset.

Managing Blacklist Rulesets

To view rulesets, navigate to the **Data Management** tab -> **DNS** tab -> **Blacklist Rulesets** tab. The panel lists the configured rulesets and their associated comments. You can also display the Disabled column which indicates which rulesets are disabled. From this panel, you can do the following:

- Add more rulesets, as described in the preceding section, [Adding a Blacklist Ruleset](#) on page 581.
- Edit a ruleset, by clicking its check box and clicking the Edit icon. You can set the following in the Blacklist Ruleset editor:
 - In the **General Basic** tab, you can change entries in any of the fields.
 - In the **Permissions** tab, you can set admin permissions for the ruleset.
- Delete a ruleset, by clicking its check box and clicking the Delete icon.
- View the rules that were imported in each ruleset by selecting it. For each rule, the panel displays the following:
 - Domain name
 - The action of the recursive DNS member when the domain name in a query matches the domain name in the rule.

To delete or edit rules in a ruleset, you must delete the ruleset from this panel, edit the CSV file and re-import it.

Enabling Blacklisting

Only DNS members with recursion enabled can support this feature. You can enable this feature at the Grid level and override it for a member or DNS view with recursion enabled.

You can also enable the DNS member to log queries that matched blacklist rules. The logs include the queried domain name, source IP address, the pattern of the matched rule, and the name of the corresponding ruleset.

To enable blacklisting:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. If the *Grid DNS Properties* or *Member DNS Properties* editor is in basic mode, click **Toggle Advanced Mode**.
3. Click **Blacklist** and complete the following:
 - Enable Domain Name Blacklist:** Select this check box.
 - Blacklist Rulesets:** To add a ruleset, click the Add icon. If there are multiple rulesets, select one from the *Select Ruleset* dialog box. Use the up and down arrows to move rulesets up and down in the list. The appliance applies rulesets in the order they are listed.
 - For blacklisted domain names, return:** Select the action of the appliance when it receives a query for a record that matches a rule with an action of Redirect/Block.
 If you selected **This list of IP addresses**, add an IP address to the **Redirect to** table by clicking the Add icon and entering the address. The addresses are listed in round robin fashion in the synthesized response of the DNS member. You can enter up to 12 IP addresses.
 - Blacklist TTL:** Specify how long the DNS client caches the A record with the redirected IP address.
 - Log queries for blacklisted domain names:** Select this option to enable the appliance to log queries for blacklisted domain names, including the source IP address of the query.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

ENABLING ZONE TRANSFERS

A zone transfer is the process of sending zone data across a network from one name server to another. When the primary server detects a change to its zone data, it notifies the secondary servers. The secondary servers reply by checking to see if the serial number they have for the zone is as large as the serial number for the zone on the primary server. If not, the secondary servers request a zone transfer.

In addition to receiving zone change notifications, a secondary server periodically polls the primary server to see if their zone data is in sync. In response, the primary server can send a DNS message containing just the changed zone data, or the entire data set. The first type of transfer is known as an incremental zone transfer, or IXFR. The second type of transfer is known as a full zone transfer, or AXFR.

A NIOS appliance, acting as the primary name server for a zone, allows zone transfers to secondary name servers by default. This includes all servers listed in the NS records for that zone. (Secondary name servers in a Grid, however, receive updated zone data via database replication by default, as explained later in this section.) You can also specify zone transfers to other name servers, such as when migrating zone data to a new server or to a management system. You can specify one or more destinations to which the local appliance sends zone transfers. You can also specify the security and format of the transfers.

Note that secondary name servers periodically query the primary name server to find out if zone data has been changed. Each query takes a certain amount of time and bandwidth on the network. By default, secondary name servers limit the rate (serial-query-rate) at which these queries are being sent. Thus when the secondary name servers are serving a large number of zones, it may take a long time to detect changes to their zone data. You can configure this value to optimize the query rate on the network. In addition, when you have set up a few secondary name servers for a large number of zones, a delay in zone transfers may occur due to the default zone transfer configuration that limits concurrent zone transfers to 10 per secondary server. You can configure the maximum value of concurrent zone transfers to optimize the zone transfer operation. For information about how to optimize zone transfers, see [Configuring Concurrent Zone Transfers](#) on page 586.

By default, Grid members automatically receive updated zone data via database replication (through an encrypted VPN tunnel). You can change the default behavior to allow Grid members to use zone transfers instead of Grid replication.

Keep in mind that a database replication updates zone data for both the active and passive nodes of an HA member. Therefore, if there is a failover, the new active node (the previous passive node) immediately begins serving zone data with fresh information. In the case of a zone transfer, the passive node does not receive zone data until after a failover, when it becomes an HA master. At that time, it performs a zone transfer. If there is a lot of zone data, the transfer can take up to several minutes, thereby causing a break in the availability of the new HA master.

If you have HA members as secondary servers, zone transfers can result in service interruption when there is a failover. Furthermore, if the primary server is down when the HA member fails over, the new active node cannot receive zone data until the primary server comes back online.

You can use TSIG (transaction signature) keys to authenticate zone transfer requests and replies. The same key name and key value must be on the primary and secondary name servers for TSIG-authenticated zone transfers to occur. When using TSIG, it is important that both appliances involved with the authentication procedure use NTP (Network Time Protocol) for their time settings (see [Using NTP for Time Settings](#) on page 313).

You can control zone transfers at the Grid, member, and zone levels. This enables you to specify a different set of name servers for a Grid, member, and zone, if necessary. You can also control which external secondary servers should receive notifications about zone updates by adding their IP addresses to the also-notify statement for each authoritative zone that is served by a Grid member. Infoblox recommends that you use this feature to notify hidden external secondary servers about zone updates, instead of putting them in stealth mode, especially when you plan to configure a large number of them. For information about how to add IP addresses to the also-notify statement, see [Configuring Zone Transfers](#) on page 584.

Configuring Zone Transfers

To configure zone transfers, you identify the servers to which zone data is transferred and optionally, servers to which data must not be transferred. For example, you can allow transfers to a network, but not to a specific server in the network.

You can specify a different set of servers for specific Grid members and zones. For example, if certain Grid members are primary servers for a zone, then you can specify the secondary servers to which that member is allowed to transfer zones.

You can also enable the appliance to add all IPv4 and IPv6 addresses for which you allow zone transfers to the also-notify statement for each authoritative zone that is served by a Grid member. The also-notify statement defines a list of addresses that receive notifications about zone updates, in addition to the IP addresses listed in the NS records for the zone. Use this feature to configure a large number of external secondary servers, instead of putting them in stealth mode.

To configure zone transfer properties:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* check box -> Edit icon.
Zone: From the **Data Management** tab, select the **DNS** tab, click the **Zones** tab -> *zone* check box, and then click the Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**, select the **Zone Transfers** tab.
3. In the Allow zone transfers to section, select one of the following:
 - **None:** Select this to deny all clients for DNS zone transfers. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows remote name servers that have the **Allow** permission to send and receive zone transfer data. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the remote name server. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server with which the local server authenticates zone transfer requests and replies. This name must match the name of the same TSIG key on other name servers that use it to authenticate zone transfers with the local server.
 - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.

- **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
- **DNSone 2.x TSIG Key:** Select this when the other name server is a NIOS appliance running DNS One 2.x code. The appliance automatically populate the value of the key in the **Value** field. The **Permission** column displays **Allow** by default. You cannot change the default permission.
- **Any Address/Network:** Select to allow or deny the local appliance to send zone transfers to any IP address.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
4. Optionally, select the **Add allowed IP addresses to also-notify** check box to add all IPv4 and IPv6 addresses listed in the “Allow zone transfers to” table to the also-notify statement for each authoritative zone served by a Grid member. When you enable this, all external secondary servers that are not defined for the zone and are allowed zone transfers will receive notifications about zone updates, in addition to name servers assigned to the zone. Infoblox recommends that you do not configure a large number of external secondary servers in stealth mode. To ensure that these secondary servers receive notifications about zone updates, add their addresses to the “Allow zone transfers to” table and grant them the “Allow” permission, and then select this check box.

Note: The appliance includes only IPv4 and IPv6 addresses. It does not include network addresses, TSIG keys, and denied addresses. When you configure a named ACL, all allowed IPv4 and IPv6 addresses in the named ACL are added to the also-notify statement.

5. Optionally, you can:
- Modify an item on the list by selecting it and clicking the Edit icon.
 - Remove an item from the list by selecting it and clicking the Delete icon.
 - Move an item up or down the list. Select it and drag it to its new position, or click the up or down arrow.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

Specifying a Zone Transfer Format

The zone transfer format determines the BIND format for a zone transfer. This provides tracking capabilities for single or multiple transfers and their associated servers.

To specify a zone transfer format:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **Zone Transfers** tab to specify the zone transfer format. Select one of the following options from the **Default Zone Transfer Format** drop-down menu:
 - **many-answers** (Secondaries run BIND 8/9): includes as many records as the packet size allows
 - **one-answer** (Secondaries run BIND 4): includes one record per packet

4. To exclude servers, click the Add icon in the **Zone Transfer Format Exceptions** table and enter the IP address of the server in the Addresses field.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring Concurrent Zone Transfers

The default number of zone transfers that are allowed is set at the Grid or member level. However, you can override the default value and configure the required concurrent zone transfers. Note that when you increase the number of concurrent zone transfers, there will be an impact on CPU and memory usage.

Note: The tcp-client value is unconditionally set to 200 to control the total number of simultaneous TCP connection, which caps the maximum inbound and maximum outbound transfer plus any DNS request made with the TCP. The tcp-client value specifies the maximum number of simultaneous DNS clients that can be handled with TCP connections and does not account for UDP connections. The UDP connection accounts for the regular DNS requests and TCP is used only for AXFR and rare DNS requests that don't fit in a UDP connection.

To specify concurrent zone transfers:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Independent appliance: From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> **Edit**.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. You can change the zone transfer settings as follows:
 - **Maximum inbound concurrent zone transfers:** The maximum number of inbound zone transfers that can be performed concurrently. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 10. Make sure that you specify a value from 10 to 100. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
 - **Maximum outbound concurrent zone transfers:** The maximum number of outbound zone transfers that can be performed concurrently. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 10. Make sure that you specify a value from 1 to 100. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
 - **Maximum concurrent outbound zone transfers per remote name server:** The maximum number of zone transfers that can be performed concurrently from a given remote name server. This configuration can be done on a per server basis. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 2. Make sure that you specify a value from 2 to 100. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
 - **Maximum concurrent SOA queries:** The maximum number of concurrent queries a secondary name server sends to the primary server to find out if the zone serial numbers have been changed. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 20. Make sure that you specify a value from 20 to 1000. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Click this Add icon to enter a new source IP address.

ABOUT ROOT NAME SERVERS

Root name servers contain the root zone file which lists the names and IP addresses of the authoritative name servers for each top-level zone. When a root name server receives a query for a domain name, it provides at least the names and addresses of the name servers that are authoritative for the top-level zone of the domain name.

You can configure the NIOS appliance to use Internet root name servers or custom root name servers. If you enable recursive queries and the appliance receives a recursive query it cannot resolve locally, it queries specified forwarders (if any) and then queries any root name servers you configure. If you do not specify internal root name servers and the appliance can access the Internet, it queries the Internet root name servers.

You can specify root name servers for the Grid, individual members, and user-defined DNS views. You can specify root name servers for all DNS views except the default view. The default view uses either the member level root name servers (if specified) or the Grid level root name servers.

Every Grid member has a default view. If you want to specify root name servers for a default view, override the Grid root name server setting at the member level and the default view can use the member-level setting.

Specifying Root Name Servers

To specify root name servers for a Grid, member, or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, select the **DNS** tab, click the **Zones** tab -> *dns_view* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* and *Member DNS Properties* editors, you must click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Root Name Servers**.
4. Select one of the following options:
 - **Use Internet root name servers:** This option is selected by default.
 - **Use custom root name servers:** Click the Add icon and enter the following information when a new row appears:
 - **Name:** Enter a name for the root name server.
 - **Address:** Enter the IP address of the root name server. The feature supports IPv4 or IPv6 values.
5. Optionally, you can:
 - Select a server from the list and click the Edit icon, to modify its information.
 - Select a server from the list and click the Delete icon.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

ABOUT SORT LISTS

A sort list prioritizes A and AAAA records on certain networks when those records are included in responses, sorting them to the beginning of the list in the response. For example, you can define a sort list when a server has two interfaces and you want the DNS clients to prefer one interface because it has a faster link.

When you define a sort list on the NIOS appliance, you specify the following:

- The IP address or network of the source of the query
- The IP addresses or networks that the appliance lists first in its response when it receives a query from the corresponding source address

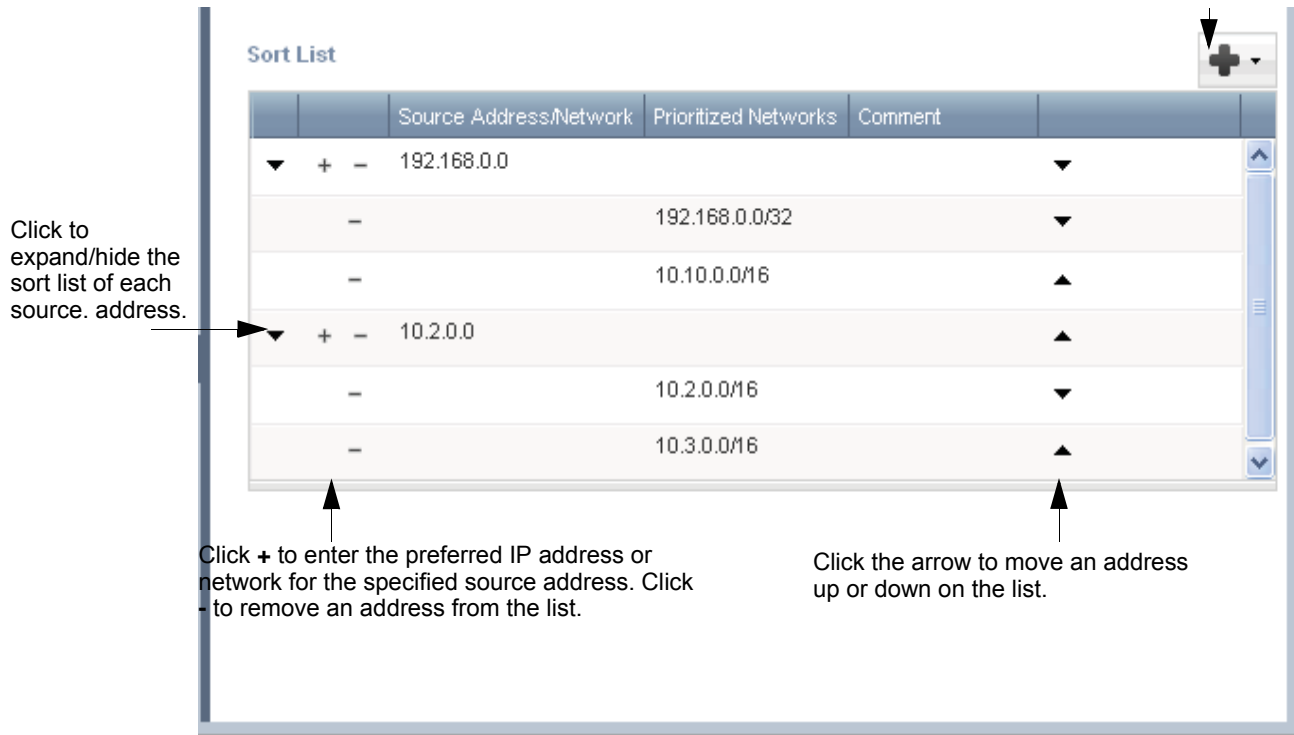
When the NIOS appliance receives a query from the specified IP address or network and the DNS lookup produces a response with multiple addresses, the NIOS appliance sorts the addresses so that those in the sort list are at the beginning of its response.

Defining a Sort List

To define a sort list for a Grid, member, or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, click the **DNS** tab -> **Zones** tab -> *dns_view* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Sort List**.
4. Click the Add icon and select either **Any** to define a sort list for any address and network, or **Address/Network** to define a sort list for a particular source IP address or network.
5. Do the following in the new row:
 - If you selected **Address/Network**, enter the IP address or network of the source of the query. The feature supports IPv4 or IPv6 values.
 - Click the Add icon beside the source IP address to add the preferred IP addresses or networks for the source. You can add as many IP addresses as necessary. When you add multiple IP addresses, you can change the order of the IP addresses. Select an IP address and drag it to its new position, or click the up or down arrow, as show in [Figure 16.2](#).

Figure 16.2 Sort List



- Enter the IP address or network in the **Prioritized Networks** field. You can add as many IP addresses as necessary. When you add multiple IP addresses, you can change the order of the IP addresses. Select an IP address and drag it to its new position, or click the up or down arrow.
- You can add additional information about the source address or network in the **Comment** field.
- To add another source IP address or network, click the Add icon again. You can create a separate sort list for each source IP address or network.
- Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING A DNS BLACKHOLE LIST

The DNS blackhole feature provides the ability to specify IP and network addresses of network devices that you do not want to use in the DNS resolution process. The DNS blackhole feature is disabled by default. When enabled, the NIOS appliance does not accept queries from IP addresses in the blackhole list and does not use them to resolve queries. For example, you can add the IP addresses of name servers that are using DNS incorrectly to prevent the NIOS appliance from accepting their queries and from using them as resolvers. You can also use this feature to fix temporary network issues. For example, you can add the IP addresses of delegated servers, configured forwarders, and DHCP servers that have temporary DNS-related issues.

You can create a DNS blackhole list for the entire Grid or create a separate list for each Grid member. For example, if one of your Grid members is behind a firewall, you might need to configure a different DNS blackhole list for this member because the clients that can access it might be mapped differently.

The appliance accepts queries from addresses and networks that are excluded from the blackhole list and uses these addresses and networks as resolvers. To add an IP address to the blackhole list, enter it and set its permission to **Include**. You can also add an IP address to the blackhole list and set its permission to **Exclude** so its not in the blackhole list, effectively allowing the NIOS appliance to respond to queries from that address and to use it as a resolver.

When you add a network to a DNS blackhole list, all the IP addresses in the network are not used in the DNS resolution process. If you want to allow some IP addresses within the network, add these addresses to the list and set their permission to “Exclude.” Ensure that you list these IP addresses before the network address because the appliance applies permissions to the addresses in the order they are listed. For example, when you add the network 10.10.0.0/24 to a DNS blackhole list, all 256 IP addresses in the network are put on the blackhole list. To allow DNS traffic to the IP addresses 10.10.0.55 and 10.10.0.88, add these two addresses before the network address in the DNS blackhole list, and then set their permissions to **Exclude**.

You can define ACEs or a named ACL to determine the IPv4 and IPv6 addresses and networks that you want to include in or exclude from a blackhole list.

Defining a DNS Blackhole List

To enable the DNS blackhole feature and configure a DNS blackhole list for a Grid or member:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click the **Blackhole** tab and complete the following:
Enable Blackhole: Select this check box to enable the DNS blackhole feature. This is disabled by default.
3. Select one of the following:
 - **None:** Select this if you do not want to configure a blackhole list. The appliance allows all clients to resolve DNS queries. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance uses clients that have the **Exclude** permission in the DNS resolution process. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.

- **IPv4 Address** and **IPv6 Address**: Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the client. The **Permission** column displays **Include** by default. You can change it to **Exclude** by clicking the field and selecting **Exclude** from the drop-down list. When you select **Include**, the appliance adds the IP address to the blackhole list and does not allow DNS queries and DNS resolution for this address. When you select **Exclude**, the appliance excludes the address from the blackhole list and allows DNS queries and resolution for the address.
- **IPv4 Network**: In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address**: Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission**: Select **Allow** or **Deny** from the drop-down list.
- **IPv6 Network**: In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address**: Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission**: Select **Allow** or **Deny** from the drop-down list.
- **Any Address/Network**: Select to include or exclude any IP addresses and networks for the DNS resolution process.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

4. Save the configuration.

SPECIFYING HOSTNAME POLICIES

You can enforce a naming policy for the hostnames of A, AAAA, Host, MX, NS, and bulk host records based on user-defined or default patterns. For MX and NS records, the hostname restrictions apply to the text in the RDATA field (right-hand side) of the resource record name.

Records that you created before you enabled the hostname checking policy need not comply with the hostname restriction that you specify.

You can select one of three preconfigured policies or define your own host naming policy with a POSIX regular expression. The policies Infoblox provides implement standard host naming restrictions according to *RFC 952, DOD Internet Host Table Specification*, and *RFC 1123, Requirements for Internet Hosts -- Application and Support*.

Note: The hostname restriction limits the hostname of A, AAAA, Host, MX, NS, and bulk host records only.

You can define your own hostname restriction policy at the Grid level only. At the member and zone levels, you can select a predefined policy or a policy that was defined at the Grid level. The appliance supports IDNs for DNS zones and resource records. For more information about IDNs, see [Support for Internationalized Domain Names](#) on page 93. You can use UTF-8 characters when you configure your own hostname checking policy.

Defining Grid Hostname Policies

You can define new hostname policies and set the hostname policy for all zones in the Grid as follows:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Host Naming**.

The Host Name Policies section lists the following preconfigured record policies:

- **Strict Hostname Checking:** You can only use hostnames that contain alphanumeric characters and dashes ("-"). You cannot use other special characters, such as underscore ("_"). Note that when you select this policy, the appliance automatically applies the policy to dynamic DNS updates and zone transfers it receives. When you select this, you can enter host names through Grid Manager using punycode, but not IDNs. The appliance stores IDNs that are created through DDNS updates and DNS transfers in punycode. You can monitor non-compliant host names using the Hostname Compliance report. For information, see [Obtaining a List of Invalid Record Names](#) on page 594.
- **Allow Underscore:** You can only use hostnames with alphanumeric characters, dashes, and underscores ("-_" and "_"). This is the default.
- **Allow Any:** You can use any hostname.

Select **Default** from the drop-down list in the Default column to change the Grid default hostname policy.

4. Click **Add** to define your own hostname checking policy.
5. Enter a record policy name and a regular expression string, and click **OK**. See [Appendix D, "Regular Expressions"](#), on page 1303 for definitions of regular expressions.

Note that Grid Manager does not validate the regular expressions that you enter. Therefore, you can inadvertently specify an invalid regular expression that might cause noncompliance errors when you create records.

6. If you select the Strict Hostname Checking policy, the **Apply policy to dynamic updates and inbound zone transfers (requires Strict Hostname Checking setting)** option is enabled by default. It enables the appliance to apply the policy to dynamic DNS updates and zone transfers that it receives. You can then select which action the appliance takes when it encounters names that do not conform to the policy. Select either **Fail** or **Warn**. If you select **Warn**, the appliance allows the dynamic DNS update or zone transfer, but logs a syslog message.

Note: The Strict Hostname Checking policy only allows alphanumeric characters and dashes ("-"). In addition, this policy allows IDNs that are written in punycode. You cannot use other special characters, such as underscore ("_"). Therefore, DDNS updates from Microsoft Active Directory controllers may not be accepted.

7. Save the configuration and click **Restart** if it appears at the top of the screen.

After you specify a hostname restriction policy, if you create a record name that does not comply with this policy and try to save it, an error message appears.

Defining Hostname Restrictions

You can select a hostname restriction policy for an individual Grid member or zone. You can specify hostname restrictions for authoritative forward-mapping zones only. You cannot specify hostname restrictions for forward zones, stub zones, IPv4 reverse-mapping zones, and IPv6 reverse mapping zones.

To select a hostname restriction policy for a Grid member or zone:

1. **Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* check box -> Edit icon.

Zone: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* -> *zone* check box -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the *Member DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Host Naming**.
4. Click **Override**.
5. From the **Host Name Policy** drop-down list, select a predefined policy or a policy that was defined at the Grid level.
6. If you select the Strict Hostname Checking policy, the **Apply policy to dynamic updates and inbound zone transfers (requires Strict Hostname Checking setting)** is enabled by default. It enables the appliance to apply the policy to dynamic DNS updates and zone transfers that it receives. You can then select which action the appliance takes when it encounters names that do not conform to the policy. Select either **Fail** or **Warn**. If you select **Warn**, the appliance allows the dynamic DNS update or zone transfer, but logs a syslog message.

Note: The strict hostname checking policy only allows alphanumeric characters and dashes. It does not allow for the use of other special characters, such as underscore ("_"). Therefore, DDNS updates from Microsoft Active Directory controllers might not be accepted.

7. Save the configuration and click **Restart** if it appears at the top of the screen.

Obtaining a List of Invalid Record Names

You can retrieve a list of all record names that do not comply with the current hostname checking policy of a zone. These could be records that were created before the current host naming policy was set. In addition, if you selected the Strict Hostname Checking policy and allowed illegal hostnames in DDNS updates and inbound zone transfers with a warning, those records are listed in this report as well.

To display the Hostname Compliance report:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns_view*-> *zone* check box.
2. Click **Hostname Compliance**.

The Hostname Compliance Report for the zone displays. It lists the record name, type, value, and comment for all records that do not comply with the hostname restriction policy of the zone.

From the report, you can select a record and do the following:

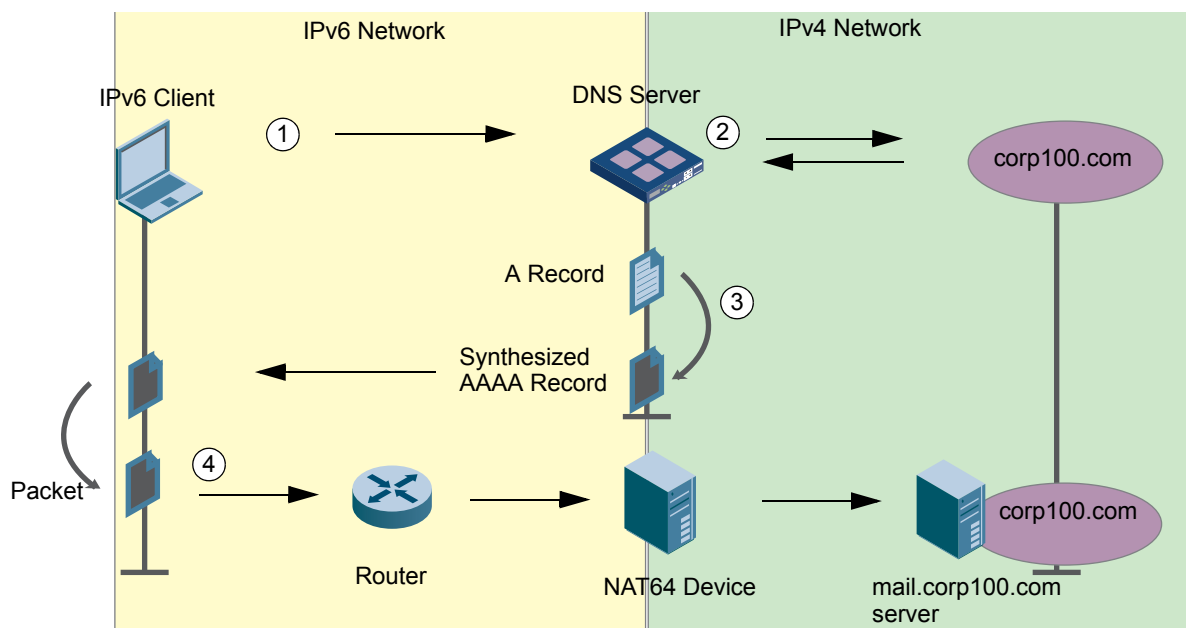
- Click the Edit icon to open the record editor.
- Click the Delete icon to move it to the Recycle Bin.

ABOUT DNS64

To support the increasing number of IPv6 and dual-stack networks, Infoblox DNS servers now support DNS64, a mechanism that synthesizes AAAA records from A records when no AAAA records exist. When you enable DNS64 on an Infoblox DNS server, it can operate with a third-party NAT64 device so IPv6-only nodes can communicate with IPv4-only nodes without any changes to either of the devices.

As illustrated in [Figure 16.3](#), when an IPv6-only host requests the AAAA record of an IPv4-only server and none exists, a DNS64-enabled server can retrieve the A record of the IPv4 server and synthesize an AAAA record. The IPv6-only host can then use the synthesized AAAA record, which contains the IPv6 proxy address for the IPv4 address in the original A record, to initiate communication with the IPv4 host.

Figure 16.3



Following are the steps illustrated in [Figure 16.3](#):

1. An IPv6-only host sends a recursive query for the AAAA record of the IPv4 server mail1.corp100.com.

2. The Infoblox DNS server attempts to resolve the request for the AAAA record, and determines that an AAAA record for mail1.corp100.com does not exist. The DNS server then performs a query for the A record of mail1.corp100.com.
3. The DNS server creates a synthetic AAAA resource record from the information in the A record, and returns the synthesized AAAA record to the requesting IPv6 host.
4. The host receives the synthetic AAAA record and sends a packet to the destination address specified in the synthetic AAAA record. The packet is routed to the IPv6 interface of the NAT64 device, which translates the packet from IPv6 to IPv4 and forwards it to the server, mail1.corp100.com.

Infoblox DNS servers can return synthesized AAAA records to both IPv4 and IPv6 clients when the client explicitly requests an AAAA record and none exists for the requested host. If a host has multiple A records, the DNS server synthesizes an AAAA record for each A record.

Infoblox DNS servers can also synthesize records for reverse-mapping zones. When a DNS server receives a query for a PTR record in the IP6.ARPA domain whose address matches a configured DNS64 prefix, the server synthesizes a CNAME record that contains an IPv4 address derived from the IPv6 address in the query. The server then sends a query for the PTR record so it can resolve the IPv4 address to the hostname.

For example, if a DNS server that is configured to synthesize records for the prefix 2001:db8::/96 receives a query for the PTR record of 2001:db8::0102:0304, it synthesizes a CNAME record that contains the IPv4 address 4.3.2.1.in-addr.arpa. The server then resolves the PTR record of the IPv4 address 4.3.2.1.in-addr.arpa.

If the server obtains the PTR record, then it sends the synthesized CNAME record and the PTR record to the client. If the zone exists, but there is no PTR record, then the server sends the synthesized CNAME record only. If the zone does not exist, then the server responds with a SERVFAIL with no answers.

Additionally, Infoblox DNS servers can generate synthesized records for DNSSEC secure zones, but only for non-DNSSEC clients. A DNS client or resolver includes the EDNS OPT pseudo-RR with the DO (DNSSEC OK) bit set to indicate that they are requesting DNSSEC data. DNS servers can generate synthesized AAAA records only when the request does not have the DO bit set. This ensures that DNSSEC clients receive only valid responses.

For additional information about DNS64, refer to the following Internet drafts:

- <http://tools.ietf.org/html/draft-ietf-behave-dns64-11>
- <http://tools.ietf.org/html/draft-ietf-behave-address-format-10>

Configuring DNS64

You can enable DNS64 on both authoritative and recursive DNS servers. You can configure DNS64 at the Grid, member or DNS view level.

To configure DNS64 on Infoblox DNS servers:

1. Create at least one DNS64 synthesis group. A synthesis group specifies the IPv6 prefix of the synthesized AAAA records. For more information, see [Adding a DNS64 Synthesis Group](#) on page 596.
2. Optionally, specify additional parameters for the synthesis group. For more information, see [Setting DNS64 Group Properties](#) on page 597.
3. Enable the DNS64 service and assign a synthesis group to the Grid, a member or a DNS view. For more information, see [Enabling DNS64 Service](#) on page 599.

On the NAT64 device, you must specify the IPv6 prefixes that are configured on the DNS server.

About Synthesis Groups

A synthesis group specifies, among other things, the IPv6 prefix for the synthesized AAAA records. Infoblox DNS servers provide a default DNS64 synthesis group with the well-known prefix 64:ff9b::/96, which is reserved for representing IPv4 addresses in the IPv6 address space. You can keep the default group, change the prefix or delete the group. You can also add a synthesis group for a Network-Specific Prefix (NSP), which is an IPv6 prefix assigned to an organization to create IPv6 representations of IPv4 addresses.

After you create a synthesis group, you can define rules to restrict the synthesis of AAAA records to certain IPv4 addresses and networks, and specify the DNS clients and networks to which the server can send synthesized AAAA records. For more information, see [Setting DNS64 Group Properties](#) on page 597.

Note that though you can control the synthesis of AAAA records, the DNS server always synthesizes CNAME records when it receives a query for an IPv6 PTR record whose address matches a prefix in a DNS64 synthesis group.

Adding a DNS64 Synthesis Group

To add a synthesis group:

1. From the **Data Management** tab, select the **DNS** tab -> **DNS64 Groups** tab, and then click the Add icon.
2. In the *DNS64 Synthesis Group* wizard, complete the following:
 - **Name:** Enter a name for the group.
 - **Prefix:** The IPv6 prefix used for the synthesized AAAA records. The default is the well-known prefix 64:FF9B::/96. The prefix length must be /32, /40, /48, /56, /64, and /96, and all bits beyond the specified length must be zero.
 - **Comment:** Optionally, enter additional information about the group.
 - **Disabled:** Select this check box if you would like to disable the group at this time. Note that you cannot disable the group if it is the only group that is used by a Grid, member or DNS view that has DNS64 enabled.
3. Click **Next** to define extensible attributes for the synthesis group. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration

Viewing DNS64 Synthesis Groups

To view synthesis groups, from the **Data Management** tab, select the **DNS** tab -> **DNS64 Groups** tab. This tab displays the following information about each group:

- **Name:** The group name.
- **Prefix:** The IPv6 prefix that is assigned to the group.
- **Comment:** The comment that was entered for the group.
- **Site:** The value of this attribute, if specified.

You can display the following additional column:

- **Disabled:** Indicates whether the group is disabled.

You can do the following:

- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Edit the properties of a synthesis group.
 - Select the synthesis group, and then click the Edit icon.
- Move a synthesis group to the Recycle Bin.
 - Select the synthesis group, and then click the Delete icon. Note that you cannot delete a synthesis group that is assigned to a Grid, member or DNS view.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Export the synthesis groups to a .csv file.
 - Click the Export icon.
- Print the list of synthesis groups.
 - Click the Print icon.

Setting DNS64 Group Properties

After you create a DNS64 synthesis group, you can specify the following:

- The IPv4 and IPv6 DNS clients and networks to which the DNS server is allowed to send synthesized AAAA records with the specified IPv6 prefix.
- The IPv4 addresses and networks for which the DNS server can synthesize AAAA records with the specified prefix.
- IPv6 addresses or prefix ranges that cannot be used by IPv6 only hosts, such as IP addresses in the ::ffff:0:0/96 network. When the DNS server retrieves an AAAA record that contains an IPv6 address that matches an excluded address, it does not return the AAAA record. Instead, it synthesizes an AAAA record from the A record. Note that a DNS server synthesizes the AAAA record of a host that has both A and AAAA records when all the IPv6 addresses in the AAAA records match the excluded addresses. If the host has multiple AAAA records and some of them contain excluded IPv6 addresses, then the server returns the remaining AAAA records.

You can add individual access control entries (ACEs) or use a named access control list (ACL) to define these clients. For information about how to define named ACLs, see [Defining Named ACLs](#) on page 307.

To configure DNS64 group properties:

1. From the **Data Management** tab, select the **DNS** tab -> **DNS64 Groups** tab -> *group* check box -> Edit icon.
2. In the **General** tab of the *DNS64 Synthesis Groups* editor, you can do the following:
 - Modify the name, prefix or comment.
 - Select the **Disabled** check box, if you want to disable the group at this time.

Perform DNS64 synthesis for these clients: Specify IPv4 and IPv6 hosts and networks to which Infoblox DNS servers can send synthesized AAAA records. The default is to allow any IPv4 and IPv6 address and network.

Select one of the following:

- **None:** Select this if you do not want to define specific addresses or networks to which the appliance sends synthesized AAAA records. When you select this, the appliance sends synthesized AAAA records to all clients. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance sends synthesized AAAA records to the clients that have the **Allow** permission in the list. You can click **Clear** to remove the selected named ACL.
- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **Any Address/Network:** Select this to allow or deny any IP addresses to which the appliance sends synthesized AAAA records.

Mapped IPv4 Addresses: Specify IPv4 addresses and networks for which the DNS server synthesizes AAAA records. The default is to allow the DNS server to synthesize AAAA records for any IPv4 address in any network. Select one of the following:

- **None:** Select this if you do not want to define specific IPv4 addresses or networks for which the DNS server synthesizes AAAA records. The appliance synthesizes AAAA records for all IPv4 clients. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance synthesizes AAAA records for the clients that have the **Allow** permission in the list. You can click **Clear** to remove the selected named ACL.
- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address:** Select this to add an IPv4 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **Any Address/Network:** Select this to allow or deny any IPv4 addresses for which the appliance synthesizes AAAA records.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

Exclude IPv6 addresses: Specify IPv6 addresses of AAAA records that the appliance treats as nonexistent. The DNS server does not return the AAAA record of an address from this list. Instead, it synthesizes an AAAA record from the A record.

- **None:** Select this if you do not want to define specific IPv6 addresses or networks of AAAA records that the appliance treats as nonexistent. The appliance treats all IPv6 addresses as nonexistent. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance synthesizes AAAA records from A records for the clients that have the **Allow** permission in the list. You can click **Clear** to remove the selected named ACL.
- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv6 Address:** Select this to add an IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.

- **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
- **Any Address/Network:** Select this to allow or deny any IP addresses of AAAA records that the appliance treats as nonexistent.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
- **Extensible Attributes:** You can modify the attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#) on page 160.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Enabling DNS64 Service

You can enable DNS64 at the Grid, member, and DNS view level. At least one DNS64 synthesis group must be configured before you can enable DNS64.

To enable DNS64 and apply DNS64 synthesis groups:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* check box -> Edit icon.
DNS View: From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns_view* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid* and *Member DNS Properties* editor, click **Toggle Advanced Mode**, and then click **DNS64**. In the *View DNS Properties* editor, just click **DNS64**.
3. Do the following in the **DNS64** tab:
 - **Enable DNS64:** Select this check box.
 - **Synthesis Groups:** Click the Add icon and select a synthesis group.
4. Save the configuration and click **Restart** if it appears.



Chapter 17 DNS Views

DNS views enable the NIOS appliance to serve different versions of DNS data based on the host accessing it. The topics in this chapter include:

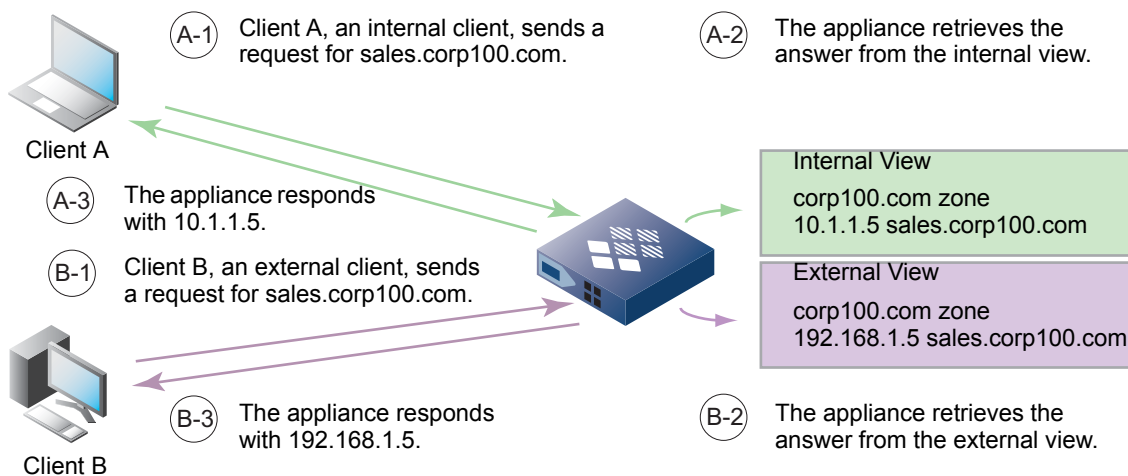
- [*Using Infoblox DNS Views*](#) on page 602
 - [*About DNS Views and Network Views*](#) on page 604
- [*Configuring DNS Views*](#) on page 604
 - [*Adding a DNS View*](#) on page 605
 - [*Defining Match Clients Lists*](#) on page 605
 - [*Defining a Match Destinations List*](#) on page 607
 - [*Managing the DNS Views of a Grid Member*](#) on page 609
 - [*Managing Recursive DNS Views*](#) on page 609
 - [*Managing the Order of DNS Views*](#) on page 610
 - [*Managing DNS Views*](#) on page 611
- [*Configuration Example: Configuring a DNS View*](#) on page 612

USING INFOBLOX DNS VIEWS

DNS views provide the ability to serve one version of DNS data to one set of clients and another version to another set of clients. With DNS views, the NIOS appliance can provide a different answer to the same DNS query, depending on the source of the query.

In [Figure 17.1](#), the appliance has two views: an Internal and an External DNS view. When the appliance receives queries from DNS clients, it responds with data from either the Internal or External DNS view, depending on the source IP address. When the appliance receives a query from Client A and determines that it can resolve the query from data in the Internal view, the appliance responds with the IP address of the site in the Internal view. When the appliance receives a query from Client B and determines that it can resolve the query from data in the External view, it responds with the IP address in the External view.

Figure 17.1 Internal and External Views



You can configure both forward and reverse mapping zones in DNS views and provide DNS services, such as name resolution, zone transfers and dynamic DNS updates. For information about these services, see [Configuring DNS Services](#) on page 555.

You can provide multiple views of a given zone with a different set of records in each DNS view. In [Figure 17.2](#), both views contain the corp100.com zone and the sales.corp100.com zone. The finance.corp100.com zone is only in the internal DNS view, and only internal users are allowed to access records in that zone. Resource records can also exist in multiple zones. In the example, the A records for serv1.sales.corp100.com and serv2.sales.corp100.com are in the sales.corp100.com zones in both views.

Figure 17.2 Zone Data in Each DNS View

Internal DNS View	corp100.com	sales.corp100.com	finance.corp100.com
	MX rmail.corp100.com	A serv1.sales.corp100.com	A server.finance.corp100.com
	NS dnsoneA.corp100.com	A serv2.sales.corp100.com	A printer.finance.corp100.com
	A host1.corp100.com	A serv3.sales.corp100.com	A fin1.finance.corp100.com
	A host2.corp100.com	A printer.sales.corp100.com	A fin2.finance.corp100.com
		A host1.sales.corp100.com	
		A host2.sales.corp100.com	
External DNS View	corp100.com	sales.corp100.com	
	MX email.corp100.com	A web3.sales.corp100.com	
	A web1.corp100.com	A ftp.sales.corp100.com	
	A web2.corp100.com	A serv1.sales.corp100.com	
		A serv2.sales.corp100.com	

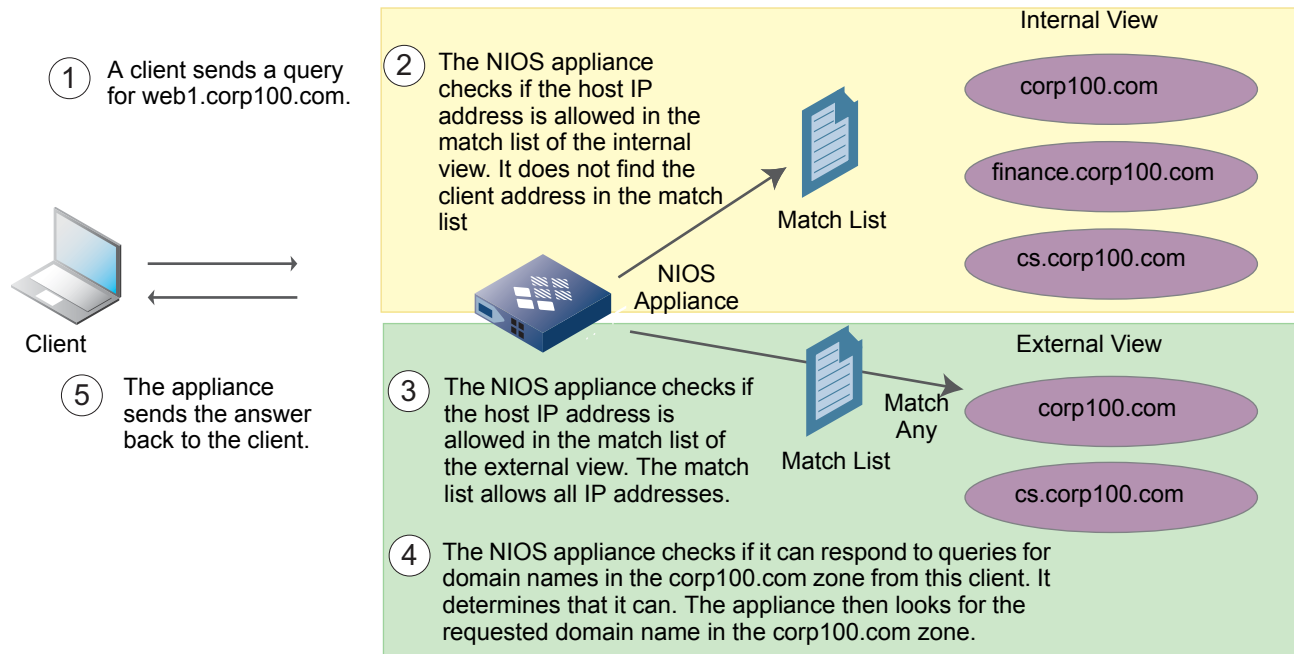
You can control which clients access a DNS view through the use of a match list specifying IP addresses and/or TSIG (transaction signature) keys. When the NIOS appliance receives a request from a client, it tries to match the source IP address and/or TSIG key with its match list when determining which DNS view, if any, the client can access. After the appliance determines that a client can access a DNS view, it checks the zone level settings to determine if it can provide the service that the client is requesting.

For information on TSIG keys or defining zone transfer settings, see [Enabling Zone Transfers](#) on page 583. For more information on match lists, see [Defining Match Clients Lists](#) on page 605. For information on defining query settings, refer to [Controlling DNS Queries](#) on page 570.

[Figure 17.3](#) illustrates how the NIOS appliance resolves a query for a domain name in a zone of a DNS view. In the example, the internal DNS view is listed before the external DNS view. Therefore, when the appliance receives a query, it checks the match list of the internal DNS view first. If it does not find the source address in the match list of the internal DNS view, it checks the match list of the external DNS view. The match list of the external DNS view allows all IP addresses.

Next, the NIOS appliance checks the zone level settings to determine if it is allowed to resolve queries from the client for domain names in that zone. After the appliance determines it is allowed to respond to queries from this client, it resolves the query and sends back the response to the client.

Figure 17.3 Query Resolution



When you create more than one DNS view, as shown in [Figure 17.3](#), the order of the views is important. View order determines the order in which the NIOS appliance checks the match lists. In [Figure 17.3](#), the internal DNS view is listed before the external DNS view. If the views were reversed, no hosts would receive DNS replies from the internal DNS view because the match list of the external DNS view allows replies to clients with any IP address. For information on how to order views, see [Managing the DNS Views of a Grid Member](#) on page 609.

In a Grid, each Grid member can host its own set of views. A Grid member can serve as the primary or secondary server for multiple views of a particular zone. For information about specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#) on page 623.

About DNS Views and Network Views

The NIOS appliance provides one default DNS view, which is always associated with the default network view. You can create additional network and DNS views. A network view is a single routing domain with its own networks. For information about network views, see [Configuring DHCP for IPv4](#) on page 843.

The default DNS view initially allows all IP addresses access, and has the same recursion setting as the Grid. You can change these properties and rename the default DNS view, but you cannot delete it. When you upgrade or migrate from a name server, or an earlier version of software that does not support DNS views, the appliance places all the zones defined in the older release in the default DNS view. You can then create additional views and organize the zones in each view.

When you create a network view, the appliance automatically creates a corresponding DNS view with “default.” prepended to the name of the network view. You can rename the system-defined DNS view and configure its properties.

If an appliance contains only one network view, all DNS views are associated with that network view. If there are multiple network views, the appliance lists the network views so you can select one from the list. The appliance displays the network views only when there are multiple network views configured.

A DNS view can be in one network view only, but a network view can have multiple DNS views. If you enable dynamic DNS updates, you must specify which DNS view receives the updates. In a network view, only one DNS view can receive the dynamic DNS updates. For information, see [Sending DDNS Updates to a DNS Server](#) on page 709.

CONFIGURING DNS VIEWS

Following are the tasks to configure a DNS view:

1. Add a DNS view, as described in [Adding a DNS View](#) on page 605.
2. Add zones to the DNS view. You can add authoritative forward-mapping and reverse-mapping zones, as well as delegated, forward, and stub zones. For information about configuring each type of zone, see [Configuring Authoritative Zones](#) on page 616 and [Configuring Delegated, Forward, and Stub Zones](#) on page 638.

You can optionally do the following:

1. Define a Match Clients list and a Match Destination list to restrict access to the DNS view. For more information, see [Defining Match Clients Lists](#) on page 605 and [Defining a Match Destinations List](#) on page 607.
2. Copy resource records from one zone to another. This is useful when different DNS views have the same zone and have multiple resource records in common. For information, see [Managing DNS Views](#) on page 611.
3. Create resource records in a group and share the group among multiple zones. For information, see [About Shared Record Groups](#) on page 680.
4. Specify which interface IP address is published in the glue A record of the DNS view. For information, see [Changing the Interface IP Address](#) on page 609.
5. Manage recursive views. For information, see [Managing Recursive DNS Views](#) on page 609.
6. Manage the order of the DNS views, as this determines the order in which the NIOS appliance checks the Match Clients list. For information, see [Managing the Order of DNS Views](#) on page 610.
7. Configure forwarders for a DNS view. For more information, see [Using Forwarders](#) on page 569.
8. Enable AAAA filtering and configure a list of IPv4 networks and addresses for allowing or denying AAAA filtering from the appliance. For information, see [Controlling AAAA Records for IPv4 Clients](#) on page 573.

Adding a DNS View

You can add up to 255 DNS views. When you add a DNS view, specify the following:

- The network view in which you are creating the DNS view.
The appliance lists the network views only when there are multiple network views. Otherwise, it automatically associates the DNS view with the default network view.
- A Match Clients list specifying the hosts allowed access to the DNS view.
If you do not define a list, the appliance allows all hosts to access the DNS view. For more information, see [Defining Match Clients Lists](#) on page 605.
- Whether recursive queries are allowed.
When a name server is authoritative for the zones in a DNS view, you can disable recursion since your name server should be able to respond to the queries without having to query other servers.
If you want to allow a Grid member to respond to recursive queries from specific IP addresses, you can create an empty DNS view, that is, one that has no zones in it, and enable recursion. For information, see [Configuration Example: Configuring a DNS View](#) on page 612.

Note: This setting overrides the recursion setting at the Grid and member levels.

To configure a new DNS view:

1. If there is more than one network view in the Grid, select the network view in which you are creating the DNS view.
2. From the **Data Management** tab -> **DNS** tab, expand the Toolbar and click **Add** -> **Add DNS View**.
3. In the *Add DNS View* wizard, complete the following fields:
 - **DNS View:** Enter the name of the DNS view. It can be up to 64 characters long and can contain any combination of printable characters. Each DNS view must have a unique name. You cannot create two DNS views with the same name, even if they are in different network views.
 - **Comment:** Optionally, enter information about the DNS view. You can enter up to 256 characters.
 - **Enable Recursion:** This field's initial default state is inherited from the Grid. It is inactive and greyed out until you click **Override**. After you click override, you can select or clear the check box to define a setting that applies to the DNS view only.
Note that a DNS view actually inherits its recursion setting from the Grid members that serve its zones. When you first create a DNS view though, it does not have any zones and therefore inherits its setting from the Grid. After you create zones in the DNS view, Grid Manager can then determine the associated members and display the resulting inheritance. If a DNS view has multiple zones served by multiple members with different recursion settings, you can view the different settings in the Multi-Inheritance viewer.
You can click **Inherit** to have the DNS view inherit its recursion setting from the Grid.
 - **Disable:** Select this check box to disable this DNS view.
4. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define a Match Clients list. For information, see [Defining Match Clients Lists](#) on page 605.
or
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Defining Match Clients Lists

When you configure a DNS view, you can create a Match Clients list to identify source IP addresses and TSIG keys that are allowed or denied access to the DNS view. The NIOS appliance determines which hosts can access a DNS view by matching the source IP address or TSIG key with its Match Clients list. After the appliance determines that a host can access a DNS view, it checks the zone level settings to determine whether it can provide the service that the host is requesting for that zone.

If you do not configure a Match Clients list, then all devices are allowed access to the DNS view. However, if you configure a Match Clients list, then only those devices in the list with “Allow” permission can access the DNS view. All other devices are denied access, including Grid members. Therefore, to allow a primary server of a zone to receive dynamic DNS updates from member DHCP servers, you must add the members to the Match Clients list as well. Note that if you “Deny” permission to certain IP addresses or networks, you must add the “Allow Any” permission at the end of the Match Clients list to ensure that all other IP addresses and networks that are not in the “Deny” list are allowed access to the DNS view. You can add individual ACEs (access control entries) or a named ACL (access control list) to the Match Clients list. For information about named ACLs and how to define them, see [Defining Named ACLs](#) on page 307.

Defining a Match Clients List for a DNS View

You can define a Match Clients list for a DNS view when you add a new DNS view (second step of the Wizard) or when you edit an existing DNS view. For information about adding a DNS view, see [Adding a DNS View](#) on page 605. To define a Match Clients list for an existing DNS view:

1. From the **Data Management** tab, click the **DNS** tab -> **Zones** tab -> *dns_view* check box -> Edit icon. Or, if there is only one DNS view, for example the predefined default view, you can just click the Edit icon beside it.
2. In the *DNS View* editor, click **Toggle Advanced Mode**, select the **Match Clients** tab, and select one of the following:
 - **None:** Select this if you do not want to configure a Match Clients list. The appliance allows all clients to access the DNS view. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance allow access to the DNS view from sources that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the client or Grid member. This name must match the name of the same TSIG key on other name servers.
 - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
 - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
 - **DNSone 2.x TSIG Key:** Select this when the other name server is a NIOS appliance running DNS One 2.x code. The appliance automatically populate the value of the key in the **Value** field. The **Permission** column displays **Allow** by default. You cannot change the default permission.

- **Any Address/Network:** Select this to allow or deny any IP addresses to access the DNS view.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
3. Save the configuration and click **Restart** if it appears at the top of the screen. You can also click the Schedule icon at the top of the editor to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone.

Defining a Match Destinations List

You can define a Match Destinations list that identifies destination addresses and TSIG keys that are allowed access to a DNS view. When the NIOS appliance receives a DNS request from a client, it tries to match the destination address or TSIG key in the incoming message with its Match Destination list to determine which DNS view, if any, the client can access. After the appliance determines that a host can access a DNS view, it checks the zone level settings to determine whether it can provide the service that the host is requesting for that zone.

You can define a Match Destination list when you edit an existing DNS view as follows:

1. From the **Data Management** tab, click the **DNS** tab -> **Zones** tab -> *dns_view* check box -> Edit icon. Or, if there is only one DNS view, for example the predefined default view, you can just click the Edit icon beside it.
2. In the *DNS View* editor, click **Toggle Advanced Mode**, select the **Match Destinations** tab, and select one of the following:
 - **None:** Select this if you do not want to configure a Match Destinations list. The appliance allows all destination addresses to access the DNS view. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance allows access to the DNS view from the destination addresses that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:

- **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the client or Grid member. This name must match the name of the same TSIG key on other name servers.
- **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
- **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
- **DNSone 2.x TSIG Key:** Select this when the other name server is a NIOS appliance running DNS One 2.x code. The appliance automatically populate the value of the key in the **Value** field. The **Permission** column displays **Allow** by default. You cannot change the default permission.
- **Any Address/Network:** Select this to allow or deny any IP addresses to access the DNS view.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
3. Save the configuration and click **Restart** if it appears at the top of the screen. You can also click the Schedule icon at the top of the editor to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone.

Copying Zone Records

Different views of the same zone may have a number of records in common. If this is the case, you can copy zone records between views and zones.

Note: You cannot copy shared records and records that the NIOS appliance automatically creates, such as NS records and glue A records.

To copy zone records between DNS zones and views:

1. From the **Data Management** tab -> **DNS** tab, click **Copy Records** from the Toolbar.
2. In the *Copy Records* dialog box, Grid Manager displays the last selected zone or the zone from which you are copying zone records in the **Source** field. Complete the following to copy records:
 - **Destination:** Click **Select Zone** to select the destination zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one. After you select the zone, Grid Manager displays the associated DNS view.
 - **Copy All records:** Select this option to copy all the zone records, including those records not created on the NIOS appliance, such as HINFO records.
 - **Copy Specific Records:** Select this option to copy specific types of records. Select a resource record type from the **Available** column and click the right arrow to move it to the **Selected** column.
 - **Copy Options:** Select one of the following:
 - **Delete all records in destination before copying the records:** Select to delete all resource records in the destination zone before the records are copied.
 - **Overwrite existing records:** Select to overwrite existing resource records that have the same domain name owners as the records being copied.
3. Click **Copy & Close**.

Note: When you copy resource records between zones and there are pending scheduled tasks associated with these records, the appliance allows the copying of zone records before it executes the scheduled tasks.

Managing the DNS Views of a Grid Member

A Grid member can serve zones in different DNS views. You can manage the DNS views associated with a Grid member as follows:

- You can specify which interface IP address is published in glue A records in the DNS view, as described in [Changing the Interface IP Address](#) on page 609.
- You can assign an empty recursive view to a member, as described in [Managing Recursive DNS Views](#) on page 609.
- You can control the list of DNS views as described in [Changing the Order of DNS Views](#) on page 611.

Changing the Interface IP Address

By default, a Grid member publishes its LAN address in glue A records in the DNS view. You can change this default for each DNS view associated with a member. You can specify the NAT IP address or another IP address.

To specify the interface IP address for glue A records in a view:

1. From the **Data Management** tab, click the **DNS** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member DNS Configuration* editor, click **Toggle Expert Mode** if the editor is in basic mode, and then select the **DNS Views** tab.
The *Address Of Member Used in DNS Views* table lists the default DNS view and DNS views with zones that are served by the member.
3. To change the address, click the entry in the Interface column of a DNS view, and select one of the following:
 - **NAT IP Address:** Select this to use the member NAT address for glue A records in a Grid setting. Select this when you want to notify the Grid Master that it should expect packets from this member on the NAT address, not the configured interface address. The Grid Master broadcasts this NAT address to all NAT members outside of its NAT group. Do not use this option for an independent appliance serving as a DNS server. Select **Other IP Address** to publish the NAT address for the independent appliance. For information about NAT compatibility, see [NAT Groups](#) on page 226.
 - **Other IP Address:** Select this to specify another address for glue A records, or to publish the NAT address for an independent appliance. Enter the address in the **Address** field.

Note: The 255.255.255.255 limited broadcast address is reserved. The appliance does not automatically create glue A records for this address. You can however create an NS record without the associated glue records.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing Recursive DNS Views

When you add a DNS view that has recursion enabled, the appliance resolves recursive queries from hosts on the Match Clients list of that view. If the DNS view contains zones and you delete those zones, the appliance retains the view in its configuration file, as long as recursion is enabled in the view. Such a view is called an empty recursive DNS view because it does not contain any zones. It enables the appliance to respond to recursive queries from the specified clients.

In a Grid, all members automatically store DNS views that have recursion enabled in their configuration files. If you do not want a Grid member to respond to recursive queries for clients in a particular DNS view, you can remove the view from the member's configuration file.

To delete or retain an empty recursive DNS view in the DNS configuration file of a Grid member:

1. From the **Data Management** tab, click the **DNS** tab -> **Members** tab -> *Grid_member* check box -> Edit icon.
2. In the *Member DNS Configuration* editor, click **Toggle Expert Mode** if the editor is in basic mode, and then select the **DNS Views** tab.

3. The *Recursive Views Assigned to this Member* section lists the empty recursive DNS views. Move a DNS view to the *Selected* column to explicitly assign the view to the Grid member and include it in the DNS configuration file of the member. Move a DNS view to the *Available* column to remove it from the configuration file of the member. Empty recursive DNS views that you retain in the configuration file are automatically listed at the bottom of the list of DNS views. You can move them up on the list when you manually change the order of the DNS views, as described in [Managing the DNS Views of a Grid Member](#) on page 609.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing the Order of DNS Views

When a member receives a query from a DNS client, it checks the Match Client lists in the order the DNS views are listed in the *Order of DNS Views* table of the **DNS Views** tab in the DNS Member editor. The NIOS appliance can order DNS views automatically, or you can order the DNS views manually. If you choose to have the appliance automatically update the order of the DNS views, it does so after each of the following events:

- Adding a DNS view to a member.
- Removing a DNS view from a member.
- Changing the address match list of a DNS view hosted by the member.

About IP Addresses and the Order of DNS Views

NIOS appliances with both IPv4 and IPv6 enabled can contain both types of addresses in the Match Clients list. When you enable IPv6 on the appliance, the order of DNS views in the GUI may be affected. Views are ordered and sorted automatically based on Match Clients lists. Views with IPv6 enabled are sorted as follows:

- If the Match Clients lists of all views contain IPv4 addresses only—The appliance orders views based on IPv4 addresses.
- If the Match Clients lists of all views contain IPv6 addresses only—The appliance orders views based on IPv6 addresses.
- If the Match Clients list of one DNS view has IPv6 addresses and all other views have IPv4 addresses—The appliance orders views based on IPv4 addresses, and the IPv6 address is given lowest priority in the ordering.
- If the Match Clients list of one DNS view has IPv4 addresses and all other DNS views have IPv6 addresses—The appliance orders DNS views based on IPv6 addresses, and the IPv4 address is given lowest priority in the ordering.
- If the Match Clients list of one DNS view has both IPv4 and IPv6 addresses—The appliance orders DNS views based on both IPv4 and IPv6 addresses, but more priority is given to the IPv4 addresses in the ordering.

The DNS views are ordered based on the number of IP addresses that are matched by the Access Control Lists (ACLs). The order of the DNS view is as follows:

- ANY
- Large Network
- Small Network
- Multiple Addresses
- Single Address

The actual precedence of the order of the views is also based on the ACL elements:

- any match: precedence = `UINT_MAX + 1`
- address match: precedence += 1
- TSIG match: precedence += 1
- network match: precedence += 129 - split (BOTH v4 and v6)

Note that views with the same precedence are sorted based on the internal view name. For example, ‘_default’ or ‘0’.

Note: Only superusers can change the order of the views.

Changing the Order of DNS Views

To change the order of DNS views:

1. From the **Data Management** tab, click the **DNS** tab -> **Members** tab-> *Grid_member* check box -> Edit icon.
2. In the *Member DNS Configuration* editor, click **Toggle Expert Mode** if the editor is in basic mode, and then select the **DNS Views** tab.
3. In the Order of DNS Views section, select one of the following:
 - **Order DNS Views Automatically:** Click this to automatically order views after adding a new DNS view, removing a DNS view, or changing the match client list.
 - **Order DNS Views Manually:** This table lists the DNS views that have zones assigned to the Grid member and the empty recursive views associated with the member. Select a DNS view, then click an arrow to move it up or down in the list.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing DNS Views

You can list the DNS views, and then modify, disable, or remove any custom DNS view. You can modify and disable the default DNS view; however, under no circumstances can it be removed.

Listing DNS Views

After you configure additional DNS views, you can list all DNS views by navigating to the **Data Management** tab -> **DNS** tab -> **Zones** panel. This panel lists DNS views only after you modify the default DNS view or add a DNS view. If you never added DNS views or modified the default DNS view, this panel does not display the default DNS view. Instead, it lists the zones in the default DNS view. To display the properties of the default DNS view and edit it, use the Global Search function to locate and edit it.

Note that if you have not used Grid Manager to add a new DNS view, and you import DNS views through the Data Import Wizard or the API, you must log out and log back in to Grid Manager to display the newly imported DNS views.

For each DNS view, this panel displays the following by default:

- **Comment:** Comments that were entered for the DNS view.
- **Site:** Values that were entered for this pre-defined attribute.

You can also display the following column:

- **Disabled:** Indicates if the DNS view is enabled or disabled. Disabled DNS views are excluded from the `named.conf` file.

From this list, you can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- List the zones in a DNS view by clicking a DNS view name.
- Edit information about a DNS view, as described in the next section.
- Delete a DNS view, as described in [Deleting DNS Views](#) on page 612.

Modifying DNS Views

To modify a DNS view:

1. From the **Data Management** tab, click the **DNS** tab -> **Zones** tab-> *dns_view* check box -> Edit icon.
2. In the *DNS View* editor, you can do the following:

- In the **General** tab, you can change any of the information you entered through the wizard. You can also disable a DNS view to temporarily block access to a DNS view. Disabling a DNS view excludes it from the named.conf file. For a description of the fields, see the online Help or [Configuring DNS Views](#) on page 604.
 - In the **Match Clients** tab, define or update a Match Clients list, as described in [Defining Match Clients Lists](#) on page 605.
 - In the **Match Destinations** tab, define or update match destinations, as described in [Defining a Match Destinations List](#) on page 607.
 - In the **Forwarders** tab, configure forwards for the view, as described in [Using Forwarders](#) on page 569.
 - In the **Queries** tab, enable AAAA filtering and configure a list of IPv4 networks and addresses for allowing or denying AAAA filtering, as described in [Enabling AAAA Filtering](#) on page 573.
 - In the **DNSSEC** tab, you can specify parameters for DNSSEC as described in [Configuring DNSSEC on a Grid](#) on page 741.
 - In the **Root Name Servers** tab, you can configure root name servers, as described in [About Root Name Servers](#) on page 587.
 - In the **Sort List** tab, define a sort list for the DNS view, as described in [Defining a Sort List](#) on page 588.
 - In the **Blacklist** tab, define blacklist rulesets, as described in [Enabling Blacklisting](#) on page 582.
 - In the **Extensible Attributes** tab, you can modify the attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - The **Permissions** tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#) on page 160.
3. Save the configuration and click **Restart** if it appears at the top of the screen.
- or
- Click the **Schedule** icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Deleting DNS Views

You can delete a DNS view if it is not the only view associated with a network view and if it is not selected for dynamic DNS updates. You cannot remove the system-defined default DNS view. When you remove a DNS view, the NIOS appliance removes the forward and reverse mappings of all the zones defined in the DNS view.

To delete a DNS view:

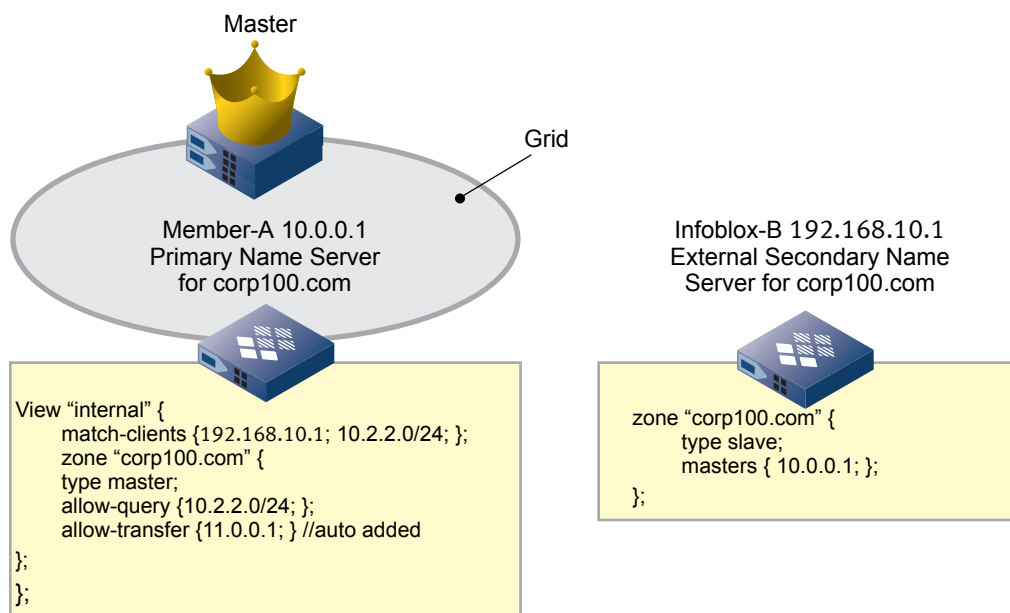
- From the **Data Management** tab, select the → **DNS** tab → **Zones** tab → *dns_view* check box.
To delete the DNS view immediately, click the **Delete** icon, and then click **Yes** to confirm the delete request. To schedule the deletion, click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.
- Grid Manager moves the view to the Recycle Bin, from which you can restore or permanently delete it.

Configuration Example: Configuring a DNS View

In [Figure 17.4](#), Member-A is a member of a Grid. It is the primary name server for the corp100.com zone in the internal DNS view. It allows the IP address 192.168.10.1 and the 10.2.2.0/24 subnet access to DNS zone data in the internal DNS view. At the zone level, it allows transfers to an external secondary server, Infoblox-B, with an IP address of 192.168.10.1. Infoblox-B is a secondary server for the corp100.com zone. The process follows these steps:

1. [Adding an Internal DNS View](#) on Member-A
2. [Adding a Zone to a DNS View](#)
3. [Copying Records Between DNS Zones](#), from the corp100.com zone in the default DNS view to the corp100.com zone in the internal DNS view
4. [Verifying the Configuration](#)

Figure 17.4 Configuring a DNS View



Adding an Internal DNS View

1. Expand the Toolbar and click **Add -> Add DNS View**.
2. In the *Add DNS View* wizard, specify the following, and then click **Next**:
 - **Name:** internal
 - **Comment:** internal DNS view
3. In the *Match Clients* panel, click **Add** and select **IPv4 Network** from the drop-down list.
4. Do the following for IP addresses in the network 10.2.2.0/24:
 - Enter **10.2.2.0/24** in the **Address** field.
 - The **Permission** field displays **Allow** by default. Leave it as is.
 - Click **Add**.

You will have 255 allowed client addresses in the Match Clients list when you are done.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

Adding a Zone to a DNS View

1. Expand the Toolbar and click **Add -> Zone -> Add Auth Zone**.
2. In the *Add Auth Zone* wizard, click **Add an authoritative forward-mapping zone** and click **Next**.
3. Specify the following, and then click **Next**:
 - **Name:** Enter **corp100.com**.
 - **DNS View:** Select **Internal** from the drop-down list.
4. In step 3 of the wizard, do the following:
 - a. Select **Use this set of name servers**.
 - b. Click the Add icon and select **Grid Primary**.
 - c. Click **Select Member** and select **Member A** from the *Select Grid Member* dialog box.
 - d. Click **Add** to add the Grid member to the list of name servers.
 - e. Click the Add icon again and select **External Secondary**.
 - f. Enter the following information, and then click **Add**:
 - **Name:** InfobloxB

- **IP Address:** 192.168.10.1

5. Click **Save & Edit** to display the *Authoritative Zone* editor and continue with the zone configuration.
6. Click **Queries**.
7. Click **Override**, and then click the Add icon and select **IPv4 Network**.
 - Enter **10.2.2.0/8** in the in the **Address** field.
 - The **Permission** field displays **Allow** by default. Leave it as is.
 - Click **Add**.

This allows queries that the appliance answers from its internal DNS view.

8. Save the configuration and click **Restart** if it appears at the top of the screen.

Copying Records Between DNS Zones

1. Navigate to the default DNS view and select the corp100.com zone.
2. Expand the Toolbar and click **Copy Records**.
3. In the **Destination** field, click **Select Zone**, and then select the **corp100.com** zone in the Internal DNS view.
4. Select **Copy all records**, and then click **OK**.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

The records from corp100.com in the default DNS view are copied to corp100.com in the internal DNS view.

Verifying the Configuration

1. In the **DNS** tab, click **Members** and select the **Member-A** check box .
2. Expand the Toolbar and click **View -> View DNS Configuration**.
3. In the *DNS Configuration File* viewer, scroll through the contents of the file.
Verify that the internal DNS view section is similar to the configuration file shown.



Chapter 18 Configuring DNS Zones

This chapter provides general information about DNS zones that you can configure and manage on the Infoblox appliance. The topics in this chapter include:

- [*About Authoritative Zones*](#) on page 616
 - [*Configuring Authoritative Zones*](#) on page 616
 - [*Creating an Authoritative Forward-Mapping Zone*](#) on page 617
 - [*Creating an Authoritative Reverse-Mapping Zone*](#) on page 618
 - [*Creating a Root Zone*](#) on page 620
 - [*Adding an Authoritative Subzone*](#) on page 620
 - [*Locking and Unlocking Zones*](#) on page 621
 - [*Enabling and Disabling Zones*](#) on page 621
- [*About Domains and Zones*](#) on page 622
 - [*IDN Support For DNS Zones*](#) on page 622
- [*Assigning Zone Authority to Name Servers*](#) on page 623
 - [*Specifying a Primary Server*](#) on page 623
 - [*Specifying a Secondary Server*](#) on page 626
- [*Using Name Server Groups*](#) on page 629
 - [*Adding Name Server Groups*](#) on page 629
 - [*Viewing Name Server Groups*](#) on page 630
 - [*Applying Name Server Groups*](#) on page 630
- [*Importing Zone Data*](#) on page 631
 - [*About Importing Data into a New Zone*](#) on page 632
 - [*About Importing Data into an Existing Zone*](#) on page 632
 - [*Importing Data into Zones*](#) on page 632
- [*Configuring Authoritative Zone Properties*](#) on page 633
- [*Removing Zones*](#) on page 634
- [*Restoring Zone Data*](#) on page 636
- [*Configuring Delegated, Forward, and Stub Zones*](#) on page 638
 - [*Configuring a Delegation*](#) on page 638
 - [*Configuring a Forward Zone*](#) on page 640
 - [*Configuring Stub Zones*](#) on page 644
- [*Viewing Zones*](#) on page 653

ABOUT AUTHORITATIVE ZONES

An authoritative zone is a zone for which the local (primary or secondary) server references its own data when responding to queries. The local server is authoritative for the data in this zone and responds to queries for this data without referencing another server.

There are two types of authoritative zones:

- Forward-mapping – An authoritative forward-mapping zone is an area of domain name space for which one or more name servers have the responsibility to respond authoritatively to name-to-address queries.
- Reverse-mapping – A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility to respond to address-to-name queries.

Configuring Authoritative Zones

You can configure and manage authoritative forward-mapping and IPv4 and IPv6 reverse-mapping zones on an Infoblox appliance. In a Grid, an authoritative forward-mapping zone is an area of domain name space for which one or more Grid members have the responsibility to respond authoritatively to name-to-address queries. The Grid members can function as primary or secondary servers for the zone.

Following are the tasks to configure an authoritative zone:

1. Create the zone. The following sections explain how to create authoritative forward-mapping zones, reverse-mapping zones, subzones, and a custom root zone:
 - [Creating an Authoritative Forward-Mapping Zone](#) on page 617
 - [Creating an Authoritative Reverse-Mapping Zone](#) on page 618
 - [Creating a Root Zone](#) on page 620
2. Assign an Infoblox appliance as the primary or secondary server of the zone. For information, see [Assigning Zone Authority to Name Servers](#) on page 623.
3. Import resource records or add resource records manually. The following provides information about resource records:
 - [Managing Resource Records](#) on page 660
 - [Importing Zone Data](#) on page 631
4. Configure additional parameters. For information, see [Configuring Authoritative Zone Properties](#) on page 633.
5. Optionally, associate the zone with one or more networks. This is useful when you want to restrict the A, AAAA and host records to IP addresses from specific networks. For information, see [Associating Networks with Zones](#) on page 813.

Creating an Authoritative Forward-Mapping Zone

An authoritative forward-mapping zone is an area of domain name space for which one or more Grid members have the responsibility to respond authoritatively to name-to-address queries.

Note: A single forward-mapping zone can map names to both IPv4 and IPv6 addresses.

To create an authoritative forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar, and click **Add -> Zone -> Add Auth Zone**.
2. In the *Add Authoritative Zone* wizard, click **Add an authoritative forward-mapping zone** and click **Next**.
3. Specify the following:
 - **Name:** Enter the domain name for the zone. Omit the trailing period (“.”) that signifies the root zone. You can use IDNs as well. For information about IDNs, see [Support for Internationalized Domain Names](#) on page 93.
 - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
 - **Comment:** Enter a descriptive comment about the zone.
 - **Disable:** Click this check box to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#) on page 621.
 - **Lock:** Click this check box to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#) on page 621
4. Save the configuration, or click **Next** to continue to the next steps in the wizard as follows:
 - Define the name servers for the zone. For information on specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#) on page 623. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629.
 - Define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322
5. Click **Restart** if it appears at the top of the screen.

Creating an Authoritative Reverse-Mapping Zone

An authoritative reverse-mapping zone is an area of network space for which one or more name servers—primary and secondary—have the responsibility to respond to address-to-name queries. Infoblox supports reverse-mapping zones for both IPv4 and IPv6 addresses.

Note: When you add an IPv4 reverse-mapping zone, the appliance automatically generates an `in-addr.arpa` space for the network address that you specify. When you add an IPv6 reverse-mapping zone, the appliance automatically generates an `ip6.arpa` space.

Specifying an RFC 2317 Prefix

RFC 2317, Classless IN-ADDR.ARPA delegation is an IETF (Internet Engineering Task Force) document that describes a method of delegating parts of the DNS IPv4 reverse-mapping tree that correspond to subnets smaller than a /24 (from a /25 to a /31). The DNS IPv4 reverse-mapping tree has nodes broken at octet boundaries of IP addresses, which correspond to the old classful network masks. So, IPv4 reverse-mapping zones (and delegation points) usually fall on /8, /16, or /24 boundaries.

With the proliferation of CIDR (Classless Inter-Domain Routing) support for routing, ISPs no longer assign entire /24 networks to customers that only need a handful of IPv4 addresses. In general, IPv4 address assignments no longer fall on classful boundaries. For DNS, a problem comes into play when an ISP gives a customer an address range that is smaller than a /24, but the customer also wants to be delegated the DNS reverse-mapping zone.

If the ISP gives you, for example, a subnet with a 25-bit mask, then you only have half of the /24 address range. If you configure your DNS server to be authoritative for the zone corresponding to a /24 subnet, the DNS server cannot resolve half of the possible reverse-mapping records in the zone. RFC 2317 defines an approach, considered a best practice, which addresses this issue.

In addition to IPv4 reverse-mapping zones, you can also configure IPv4 reverse-mapping delegation zones that have an RFC2317 prefix. For more information about configuring a delegation for a reverse-mapping zone, see [Configuring a Delegation](#) on page 638.

Note: Before enabling RFC 2317 support for zones, disable forwarders for the zone, especially when any sort of delegation (including RFC 2317) is being used. If you do not, reverse lookups may fail. For more information, contact Infoblox Support for the Tech Note on RFC 2317 delegation.

Adding an IPv4 Reverse-Mapping Zone

To add an IPv4 reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Auth Zone**.
2. In the *Add Authoritative Zone* wizard, click **Add an authoritative IPv4 reverse-mapping zone** and click **Next**.
3. Specify the following zone information:
 - Enter one of the following to identify the zone:
 - **IPv4 Network:** Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the **Netmask** drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.
To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the **RFC 2317 Prefix** field. Enter a prefix in the text field. Prefixes can include alphanumeric characters. For information, see [Specifying an RFC 2317 Prefix](#) on page 618.
 - **Name:** Enter the domain name of the reverse-mapping zone.
 - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
 - **Comment:** Optionally, enter additional information about the zone.

- **Disable this zone:** Select this option to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#) on page 621.
 - **Lock this zone:** Select this option to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#) on page 621.
4. Save the configuration, or click **Next** to continue to the next steps in the wizard as follows:
- Define the name servers for the zone. For information on specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#) on page 623. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629.
 - Define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.
5. Click **Restart** if it appears at the top of the screen.

Adding an IPv6 Reverse-Mapping Zone

To add an IPv6 reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Auth Zone**.
2. In the *Add Authoritative Zone* wizard, click **Add an authoritative IPv6 reverse-mapping zone** and click **Next**.
3. Enter the following zone information:
 - Enter one of the following to identify the zone:
 - **IPv6 Network Prefix:** Enter the 128-bit IPv6 address for the address space for which you want to define the reverse-mapping zone. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Choose the network prefix that defines the IPv6 network address space.
 - **Name:** Enter the domain name of the reverse-mapping zone.
 - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
 - **Comment:** Enter a descriptive comment about the zone.
 - **Disable:** Click this check box to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#) on page 621.
 - **Lock:** Click this check box to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#) on page 621.
4. Save the configuration, or click **Next** to continue to the next steps in the wizard as follows:
 - Define the name servers for the zone. For information on specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#) on page 623. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629.
 - Define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Creating a Root Zone

The NIOS appliance allows you to create an internal root zone for your organization. When the appliance receives a query for DNS data that is not in its cache or authoritative data, it can query an internal root server after querying any specified forwarders. If you do not specify an internal root server and the appliance can access the Internet, it queries the Internet root servers. For information on root name server, see [About Root Name Servers](#) on page 587.

To create a root zone, create an authoritative forward-mapping zone as described in [Creating an Authoritative Forward-Mapping Zone](#) on page 617 and specify the following:

- Enter a period (.) in the **Name** field.
- Optionally, enter a comment.
- Select a Grid member as the primary name server for the root zone.

Once created, the root zone automatically becomes the parent of all the zones under the root zone.

Adding an Authoritative Subzone

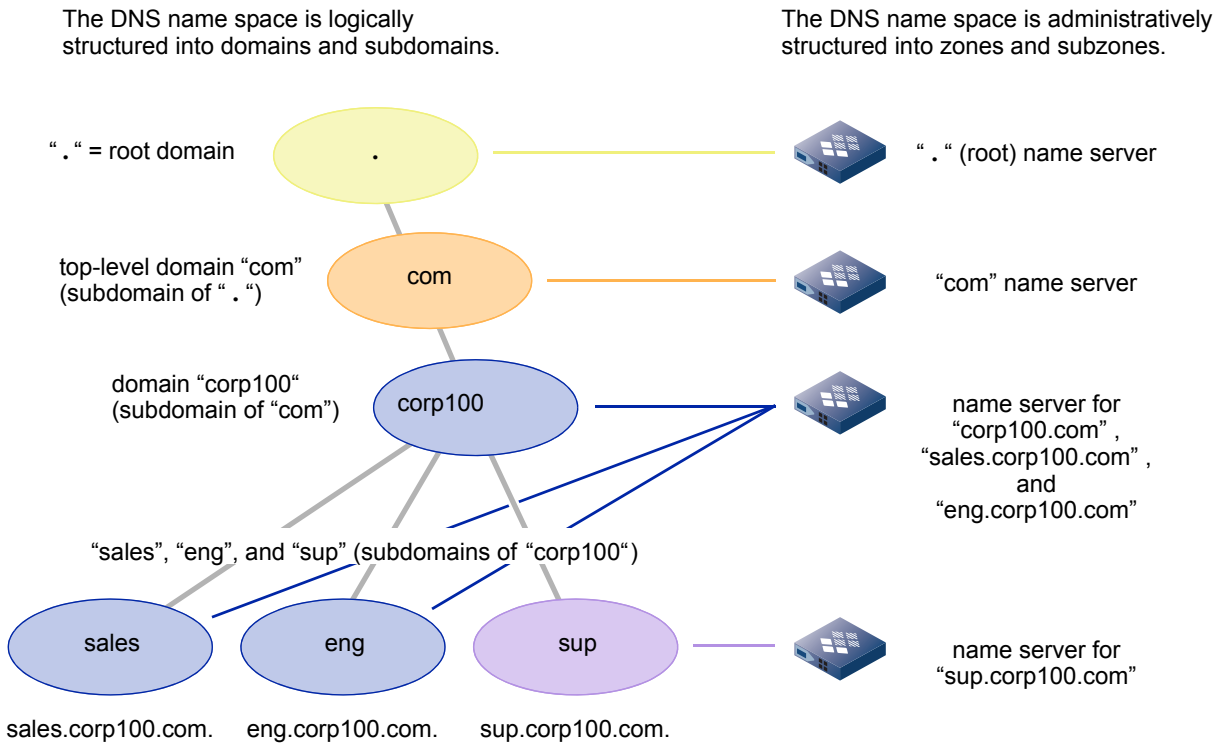
After creating a zone, you can add more zones at the same level, or add subordinate zones (*subzones*). The subzones can be authoritative, delegated, forward, or stub. For simplicity, the zones created in this example are authoritative (as are all zones by default). For information about configuring the other zone types, see [Configuring Delegated, Forward, and Stub Zones](#) on page 638.

You create an authoritative zone when you assign authority for all the resource records of a particular domain to one or more name servers. You create a subzone when you assign authority for all the resource records of a subdomain to name servers. The name servers can be the same as, or different from, the name servers that serve resource records for the parent domain.

The distinction between domains and zones is that domains provide a logical structure to the DNS name space while zones provide an administrative structure. The difference between domains and subdomains and zones and subzones is that the terms *subdomains* and *subzones* reference their relationship to a parent domain or zone. With the exception of the root domain and root zone, all domains are subdomains and all zones are subzones.

You can organize a domain based on logical divisions such as type (.com, .gov, .edu; or sales, eng, sup) or location (.uk, .jp, .us; or hq, east, west). [Figure 18.1](#) on page 621 shows one way to organize the external (public) name space and the internal (private) name space for a corporation with the domain name *corp100.com*. The external name space follows standard DNS conventions. Internally, you create an individual subdomain and corresponding subzone for each department.

Figure 18.1 Domains and Subdomains, and Forward-Mapping Zones and Subzones



Note: Throughout this documentation, the trailing dot indicating the root zone is not shown, although its presence is assumed.

The procedure for adding a subzone is the same as that used to add an authoritative zone. The only difference is that you specify the subzone name in the **Name** field. For information about adding authoritative zones, see [Configuring Authoritative Zones](#) on page 616.

Locking and Unlocking Zones

You can lock a zone when you create or edit it to prevent other administrators from making conflicting changes. When you lock a zone, Grid Manager displays **LOCKED** beside the zone name when you view the records and subzones of the zone in the Zones panel. When other administrators try to make changes to a locked zone, the system displays a warning message that the zone is locked by *admin_name*.

You can perform dynamic updates through mechanisms such as DDNS and nsupdate on a locked zone. The system can also add auto-generated records such as glue A records and NS records to a locked zone. Locks on a zone do not impact its child zones.

Only a superuser or the administrator who locked the zone can unlock it. Locks do not expire; you must manually unlock a locked zone.

Enabling and Disabling Zones

The NIOS appliance allows you to disable and enable a zone when you create or edit it. When you disable a zone, Grid Manager removes it from the DNS configuration file, but not from the database. This feature is especially helpful when you have to move or repair the server for a particular zone. You can easily disable a zone temporarily, and then enable it after the move or repair is completed.

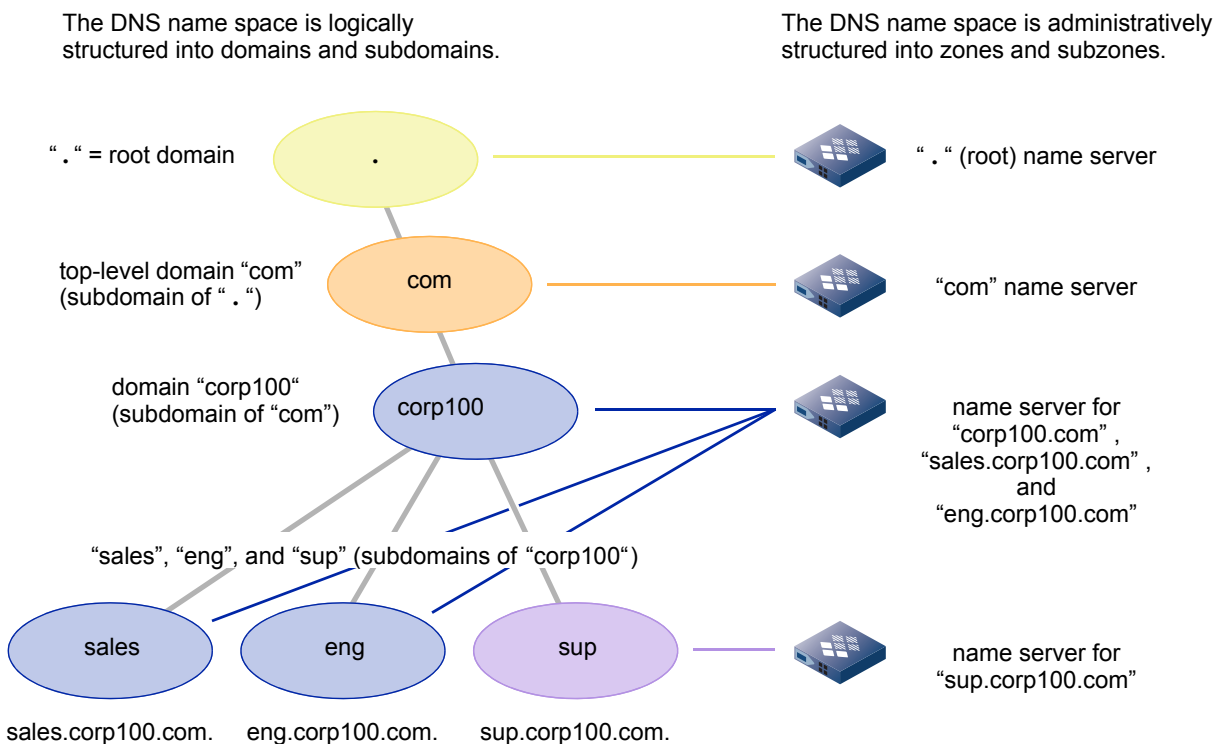
ABOUT DOMAINS AND ZONES

After creating a zone, you can add more zones at the same level, or add subordinate zones (*subzones*). The subzones can be authoritative, delegated, forward, or stub.

The distinction between domains and zones is that domains provide a logical structure to the DNS name space while zones provide an administrative structure. The difference between domains and subdomains and zones and subzones is that the terms *subdomains* and *subzones* reference their relationship to a parent domain or zone. With the exception of the root domain and root zone, all domains are subdomains and all zones are subzones.

You can organize a domain based on logical divisions such as type (.com, .gov, .edu; or sales, eng, sup) or location (.uk, .jp, .us; or hq, east, west). [Figure 18.1](#) on page 621 shows one way to organize the external (public) name space and the internal (private) name space for a corporation with the domain name *corp100.com*. The external name space follows standard DNS conventions. Internally, you create an individual subdomain and corresponding subzone for each department.

Figure 18.2 Domains and Subdomains, and Forward-Mapping Zones and Subzones



On the Infoblox appliance, you can configure and manage DNS zones and subzones.

IDN Support For DNS Zones

Grid Manager supports IDNs for DNS zones and resource records. For information about IDN, see [Support for Internationalized Domain Names](#) on page 93. You can use either IDN or punycode (representation of IDN) to create DNS zones. Even if you use punycode to create a zone, the appliance automatically generates the corresponding IDN and displays the zone name in its native characters. Make sure that you use valid punycode to create a DNS zone. If you specify an invalid punycode, the appliance retains the punycode and does not convert it into IDN. Note that the appliance displays both the IDN and punycode for an IDN zone.

The following table summarizes how the appliance displays IDNs at the DNS zone level:

Input	NIOS Displays...	NIOS DNS Domain (Punycode in the GUI)	Conversion Guidelines
hello.com	hello.com	hello.com	No conversion
прывітанне.com	прывітанне.com	xn--80adk5aaihr3f9e.com	IDN to punycode
xn--80adk5aaihr3f9e.com	прывітанне.com	xn--80adk5aaihr3f9e.com	Punycode to IDN
\xyz format	\xyz format	\xyz format	No conversion

ASSIGNING ZONE AUTHORITY TO NAME SERVERS

Forward-mapping zones answer name-to-address queries, and reverse-mapping zones answer address-to-name queries. When you create an authoritative forward-mapping zone or reverse-mapping zone, you must define a name server as a primary server for that zone. A primary server contains editable zone data, which that server can send to other (secondary) servers through zone transfers. You can also create one or more secondary name servers for a zone. A secondary server for a zone receives read-only zone data from the primary server.

Note: The primary/secondary relationship between name servers is also called “master”/“slave”. You can enter, modify, and remove zone data on the primary (or master) server, which can then send new and modified data in a read-only form to the secondary (or slave) server. Both primary and secondary name servers are authoritative for the zone data they serve. The distinction between them is how they get their zone data.

If a zone is part of an internal DNS structure for a private network, the inclusion of a secondary DNS server is optional, though highly recommended. If a zone is part of an external DNS structure for a public network such as the Internet, then a secondary server in a different subnet from the primary server is required. This requirement provides an additional safeguard against localized network failures causing both primary and secondary name servers for a zone to become inaccessible.

In Grid Manager, you can specify the primary and secondary servers for a zone or you can specify a name server group. A name server group is a collection of one primary server and one or more secondary servers. For information on name server groups, see [Using Name Server Groups](#) on page 629.

Specifying a Primary Server

When you create a zone, the primary server can be a Grid member, an external DNS server that you specify, or a Microsoft DNS server that is managed by a Grid member. For information about managing Microsoft Windows DNS servers, see [Chapter 32, Managing Microsoft Windows Servers](#), on page 953.

Although a zone typically has just one primary name server, you can specify up to ten independent servers for a single zone. When the primary server is a Grid member, however, then only that member can be the primary server.

A primary server can be in stealth mode, which means that it does not respond to queries from other name servers and its NS record is not published among the zone data. Such a server is also called a “hidden primary”.

A hidden primary provides data to its secondary servers, which in turn respond to DNS queries using this data. One of several advantages of this approach is that you can take the primary server offline for administrative or maintenance reasons without causing a disruption to DNS service (within the expiration interval set for the validity of its zone data—the default is 30 days).

When you add an authoritative forward-mapping zone and assign responsibility for the zone to a primary name server whose host name belongs to the name space of the zone, the NIOS appliance automatically generates an NS (name server) record and an A (address) record for the name server. This type of A record is called a glue record because it “glues” the NS record to the IP address (in the A record) of the name server.

In Grid Manager, you can specify the primary server for a zone when you create it using the *Add Authoritative Zone* wizard or when you edit an existing zone using the *Authoritative Zone* editor. For information on how to add a new zone through the wizard, see [Configuring Authoritative Zones](#) on page 616. The following procedure describes how to access the editor of a zone.

To specify a primary server for an existing zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Name Servers**.
3. Select **Use this set of name servers**.
4. Click the Add icon and select one of the following options for a primary server:
 - **Grid Primary:** Choose this option to select a Grid member as the primary server for the zone. See [Specifying a Grid Primary Server on page 624](#).
 - **Microsoft Primary:** Choose this option to select a Microsoft DNS server as the primary server for the zone. See [Specifying a Microsoft Primary Server](#) on page 625.
 - **External Primary:** Choose this option if the appliance is in a Grid and you want to specify a primary server outside the Grid (“external” to the Grid). See [Specifying an External Primary Server](#) on page 625.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Specifying a Grid Primary Server

In the *Add Grid Primary* panel, do the following, and then click **Add** to add the Grid member to the list of name servers for the zone as primary:

- If no member is displayed, click **Select** to specify a Grid member. When there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select a primary name server.
- **Stealth:** Click this to hide the NS record for the primary name server from DNS queries. The NIOS appliance does not create an NS record for the primary name server in the zone data. Clear the check box to display the NS record for the primary name server in responses to queries.

Changing the SOA Name for a Zone

If the primary name server of a zone is a Grid member, the NIOS appliance allows you to change the SOA (start of authority) name that is automatically created when you initially configure the zone. For example, you might want to hide the primary server for a zone. If your appliance is named `dns1.zone.tld`, and for security reasons, you may want to show a secondary server called `dns2.zone.tld` as the primary server. To do so, you would go to `dns1.zone.tld` zone (being the true primary) and change the SOA to `dns2.zone.tld` to hide the true identity of the real primary server.

To change the SOA name for a zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns_view* -> *zone* check box -> Edit icon.
2. In the *Authoritative Zone* editor, click **Settings**.
3. Click **Override** beside the **Primary name server** field and enter the new SOA name. This field supports IDN.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Specifying a Microsoft Primary Server

You can assign a Microsoft server as the primary server of a zone when it is managed by a Grid member in read/write mode. For information, see [Chapter 32, Managing Microsoft Windows Servers](#), on page 953.

When a Microsoft server is the primary server of a zone, the zone can only support standard DNS resource records. It does not support the Infoblox record types host records, bulk host records, and shared record groups. You cannot add any of these records to the zone nor assign a DNS zone with these records to a Microsoft server as the primary server.

In the Add Grid Primary panel, do the following to assign a Microsoft primary server:

1. Complete the following:
 - Select **Use this set of name servers**.
 - Click the *Add* icon and select **Microsoft Primary**.
2. In the *Add Microsoft Primary* panel, do the following, and then click **Add** to add the Microsoft primary server to the list of name servers for the zone:
 - If no server is displayed, click **Select Server** to specify a Microsoft server. When there are multiple servers, Grid Manager displays the *Server Selector* dialog box from which you can select a Microsoft server. Grid Manager lists Microsoft servers that are managed in read/write mode. It does not include Microsoft servers managed in read-only mode.
 - **Information to create NS record:** Grid Manager automatically creates the NS record. After you select a server, Grid Manager populates the **Name** and **IP Address** fields. Grid Manager uses this information when it creates the NS record, unless you select **Stealth**. You can specify a different FQDN or IP address for the NS record; for example, for a multihomed server.
 - **Store the zone in Active Directory (AD Integrated Zone):** This is enabled and selected by default only if the Microsoft server is a domain controller. Note that you can enable Active Directory integration only after the Microsoft server has been synchronized at least once because its AD ability is not known before the synchronization. This is disabled when the Microsoft server is not a domain controller.
 - **Stealth:** Select this option to hide the NS record for the primary name server from DNS queries. Grid Manager does not create an NS record for the primary name server in the zone data. Clear this option to display the NS record for the primary name server in responses to queries. Note that this option is not available for AD-integrated zones.

Specifying an External Primary Server

In the Add External Primary panel, do the following, and then click **Add** to add the external primary server to the list of name servers for the zone:

- **Name:** Type a resolvable domain name for the external primary server.
- **Address:** Type the IP address of the external primary server.
- **Use TSIG:** To authenticate zone transfers between the local appliance and the external primary server using a TSIG (transaction signature), select this check box. Infoblox TSIGs use HMAC-MD5 hashes. These are keyed one-way hashes for message authentication codes using the Message Digest 5 algorithm. For details, see *RFC 1321, The MD5 Message-Digest Algorithm*, and *RFC 2104, HMAC: Keyed-Hashing for Message Authentication*.
 - **Key name:** Type or paste the name of the TSIG key you want to use. This must be the same name as that of the TSIG key on the external primary server.
 - **Key Data:** Type or paste a previously generated key. This key must also be present on the external primary server. You can generate a TSIG key, or obtain the TSIG key name and key from the external name server, either by accessing the server yourself or by requesting the server administrator to deliver them to you through some out-of-band mechanism. Then type or copy-and-paste the name and key into the appropriate fields.
- **Use 2.x TSIG:** If you want to use TSIG authentication and the external primary name server is a NIOS appliance running DNS One 2.x code, select this check box. The local appliance generates the required TSIG key for authenticating DNS messages to and from appliances running DNS One 2.x code.

Note: On the appliance you configure as a secondary server for a zone, you must associate a TSIG key for each primary server to which the secondary server requests zone transfers. On the appliance you configure as a primary server for a zone, you can set a TSIG key at the Grid, member, or zone level. Because the secondary server requests zone transfers, it must send a specific key in its requests to the primary server. Because the primary server responds to the requests, it can have a set of TSIG keys from which it can draw when responding. As long as the primary server can find the same TSIG key that the secondary sends it, it can verify the authenticity of the requests it receives and authenticate the responses it sends. Use NTP to synchronize the time on both name servers that use TSIG-authenticated zone transfers.

Specifying a Secondary Server

A secondary name server is as authoritative for a zone as a primary server. Like a primary server, a secondary server answers queries from resolvers and other name servers. The main difference between a secondary and primary server is that a secondary server receives all its data from a primary server, or possibly from another secondary server that relays zone data it receives. The zone data passes from a primary to a secondary server (and possibly from that secondary server on to another secondary server). This process is called a zone transfer.

The advantage of using primary and secondary name servers is that you enter and maintain zone data in one place—on the primary server. The data is then distributed to the one or more secondary servers.

Secondary servers can be Grid members, external DNS servers or Microsoft DNS servers that are managed by Grid members.

In Grid Manager, you can specify the secondary server for a zone when you create it using the *Add Authoritative Zone* wizard and when you edit an existing zone using the *Authoritative Zone* editor. For information on how to add a new zone through the wizard, see [Configuring Authoritative Zones](#) on page 616. The following procedure describes how to access the editor of a zone.

To specify a secondary server for an existing zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* check box, and then click the Edit icon.
 2. In the *Authoritative Zone* editor, click **Name Servers**.
 3. Select **Use this set of name servers**.
 4. Click the Add icon and select one of the following options:
 - **Grid Secondary:** Selects the local appliance as the secondary server (or if the appliance is deployed in a Grid and you want to make a different member the secondary server). See [Adding Grid Secondaries](#) on page 627.
 - **Microsoft Secondary:** Select this option if you want to specify a managed Microsoft DNS server as a secondary server. See [Specifying Microsoft Secondary Servers](#) on page 627.
 - **External Secondary:** Select this option if the appliance is in a Grid and you want to specify a secondary server outside the Grid (“external” to the Grid), or if the appliance is deployed independently from a Grid. See [Specifying External Secondaries](#) on page 628.
 5. Save the configuration and click **Restart** if it appears at the top of the screen.
- or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Adding Grid Secondaries

In the Add Grid Secondary panel, enter the following, and then click **Add** to add the Grid secondary server to the list of name servers for the zone:

- If no member is displayed, click **Select** to specify a Grid member. When there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select a secondary name server.
- **Stealth:** This setting applies only if the primary server is a Grid member or Microsoft server. Select this to hide the NS record for the secondary name server from DNS queries. The NIOS appliance does not create an NS record for this name server in the zone data. Select the check box again to display the NS record for the secondary name server in responses to queries. A secondary server in stealth mode is also known as a “hidden secondary”.
For example, you can configure a hidden secondary when a secondary server is at a branch office with a slow connection to the rest of corporate network. Configure local hosts at the branch office to send DNS queries to the secondary server, but keep it hidden from other name servers on the rest of the network so that they do not send it queries. Instead, they use a server located in a different part of the network that has faster connection speeds.
- **Lead Secondary:** This option becomes available only after you specify the primary name server as external. When a primary server is external to a Grid whose members are secondary servers, you can select this check box to designate one member as a lead secondary. The primary server sends zone transfers to the lead secondary, which distributes the zone data to the other secondary servers in the Grid using zone transfers (not the Grid data replication mechanism). After you designate a Grid member as a lead secondary for a zone, you do not have to configure members to use the lead secondary server. All other Grid members acting as secondary servers for the zone automatically use the lead secondary to get zone data. Using a lead secondary simplifies the addition, modification, and removal of other secondary servers in the Grid. As long as the lead secondary remains unchanged, you need not update intervening firewall policies or the external primary server whenever you make changes to non-lead secondary Grid members. This approach also reduces the amount of traffic between primary and secondary servers.
- **Update Zones Using:** This option becomes available only after you specify a Grid member as the primary server.
 - **Grid Replication (recommended):** Select this check box to use Grid replication to move zone data from the primary to secondary servers.
 - **DNS Zone Transfers:** Select this check box to use the DNS zone transfer process to move zone data from the primary to secondary servers.

Specifying Microsoft Secondary Servers

You can assign a Microsoft server as the primary server of a zone when it is managed by a Grid member in read/write mode. For information, see [Chapter 32, Managing Microsoft Windows Servers](#), on page 953.

Since Microsoft servers cannot replicate data from the Grid, when a DNS zone is defined as a secondary on a Microsoft server, the Microsoft server obtains the content of the zone only through DNS zone transfers.

- In the Add Microsoft Secondary panel, do the following:
 - If no server is displayed, click **Select Server** to specify a Microsoft server. When there are multiple servers, Grid Manager displays the *Server Selector* dialog box from which you can select a Microsoft server. Grid Manager lists Microsoft servers that are managed in read/write mode. It does not include Microsoft servers managed in read-only mode.
 - **Information to create NS record:** Grid Manager automatically creates the NS record. After you select a server, Grid Manager populates the **Name** and **IP Address** fields. Grid Manager uses this information when it creates the NS record, unless you select Stealth.
 - **Stealth:** This setting applies only if the primary server is a Grid member or a Microsoft server. Select this option to hide the NS record for the secondary name server from DNS queries. Grid Manager does not create an NS record for this name server in the zone data. Clear this option to display the NS record for the secondary name server in responses to queries.

Specifying External Secondaries

In the Add External Secondary panel, enter the following, and then click **Add** to add the external secondary server to the list of name servers for the zone:

- **Name:** Enter a resolvable domain name for the external secondary server.
- **Address:** Enter the IP address of the external secondary server.
- **Stealth:** This setting applies only if the primary server is a Grid member or a Microsoft server. Click this check box to hide the NS record for the secondary name server from DNS queries. The NIOS appliance does not create an NS record for the secondary name server in the zone data. Select the check box again to display the NS record for the secondary name server in response to queries.

Note: To avoid an impact on your database performance, Infoblox recommends that you do not configure a large number of external secondary servers in stealth mode. To ensure that these secondary servers receive notifications about zone updates, you can allow zone transfers for these IP addresses and then enable the appliance to add them to the also-notify statement. For information about how to configure this feature, see [Configuring Zone Transfers](#) on page 584.

- **Use TSIG:** To authenticate zone transfers between the local appliance and the external secondary server using a TSIG (transaction signature), select this check box. Infoblox TSIGs use HMAC-MD5 hashes. These are keyed one-way hashes for message authentication codes using the Message Digest 5 algorithm. For details, see RFC 1321, *The MD5 Message-Digest Algorithm*, and RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.
 - **Key name:** Type or paste the name of the TSIG key you want to use. This must be the same name as that of the TSIG key for this zone on the external secondary server.
 - **Key:** Type or paste a previously generated key. On the external secondary server, this key must also be present and associated with this zone. You can generate a TSIG key, or you can obtain the TSIG key name and key from the external name server, either by accessing the appliance yourself or by requesting the appliance administrator to deliver them to you through some out-of-band mechanism. Then, type or copy-and-paste the name and key into the appropriate fields.
- **Use 2.x TSIG:** Select this check box to use TSIG authentication and the external secondary name server is a NIOS appliance running DNS One 2.x code. The local appliance generates the required TSIG key for authenticating DNS messages to and from appliances running DNS One 2.x code.

Note: On the appliance you configure as a secondary server for a zone, you must associate a TSIG key for each primary server to which the secondary server requests zone transfers. On the appliance you configure as a primary server for a zone, you can set a TSIG key at the Grid, member, or zone level. Because the secondary server requests zone transfers, it must send a specific key in its requests to the primary server. Because the primary server responds to the requests, it can have a set of TSIG keys from which it can draw when responding. As long as the primary server can find the same TSIG key that the secondary sends it, it can verify the authenticity of the requests it receives and authenticate the responses it sends. Use NTP to synchronize the time on both name servers that use TSIG-authenticated zone transfers.

USING NAME SERVER GROUPS

A name server group is a collection of one primary DNS server and one or more secondary DNS servers. Grouping a commonly used set of primary and secondary DNS servers together simplifies zone creation by enabling you to specify a single name server group instead of specifying multiple name servers individually. After you create a name server group, you can then assign it to serve authoritative forward-mapping and reverse-mapping zones.

Note: Only superusers can create and manage name server groups.

Adding Name Server Groups

To add a name server group:

1. From the **Data Management** -> **DNS** tab, do one of the following:
 - Click the **Name Server Groups** tab -> Add icon -> **Group** -> **Name Server Group**.
 - From the Toolbar, click the Add icon -> **Group** -> **Name Server Group**
2. In the *Name Server Group* wizard, do the following:
 - **Name:** Type a name that provides a meaningful reference for this set of servers.
 - **Name Servers:** Click the Add icon and select one of the following options for every server that you are adding to the NS group:
 - **Grid Primary:** Choose this option to select a Grid member as the primary server for the zone. See [Specifying a Grid Primary Server on page 624](#).
 - **Grid Secondary:** Choose this option to select a Grid member as a secondary server for the zone. See [Adding Grid Secondaries](#) on page 627.
 - **External Primary:** Choose this option if the appliance is in a Grid and you want to specify a primary server outside the Grid (“external” to the Grid). See [Specifying an External Primary Server](#) on page 625.
 - **External Secondary:** Choose this option if the appliance is in a Grid and you want to specify a secondary server outside the Grid (“external” to the Grid), or if the appliance is deployed independently from a Grid. See [Specifying External Secondaries](#) on page 628.
 - **Default NS Group:** Select this to specify this name server group as the default.
 - **Comment:** Optionally, enter additional information about the NS group.
3. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

A newly created name server group appears in the **Name Server Groups** tab. You can then associate it with forward-mapping and reverse-mapping zones.

Viewing Name Server Groups

You can view the configured name server groups by navigating to the **Data Management** tab -> **DNS** tab -> **Name Server Groups** tab.

The panel displays the following information about each name server group:

- **Name:** The name of the name server group.
- **Comment:** Comments that were entered for the name server group.
- **Site:** Values that were entered for this pre-defined attribute.

You can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Edit the properties of a name server group.
 - Click the check box beside a name server group, and then click the Edit icon.
- Delete a name server group.
 - Click the check box beside a name server group, and then click the Delete icon.
- Export the list of Grid members to a .csv file.
 - Click the Export icon.
- Print the list of Grid members.
 - Click the Print icon.

Applying Name Server Groups

In Grid Manager, you can assign a name server group to a zone when you first create it using the *Add Authoritative Zone* wizard and when you edit an existing zone using the *Authoritative Zone* editor. For information on creating a zone using the wizard, see [Configuring Authoritative Zones](#) on page 616. The panels used to assign a name server group to a zone are the same in the wizard or and editor. The only difference is the way you access it. The following procedure describes how to specify a name server group when editing a forward-mapping zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* check box, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Name Servers**.
3. Select **Use this name server group**, and then select the name server group from the drop-down list.

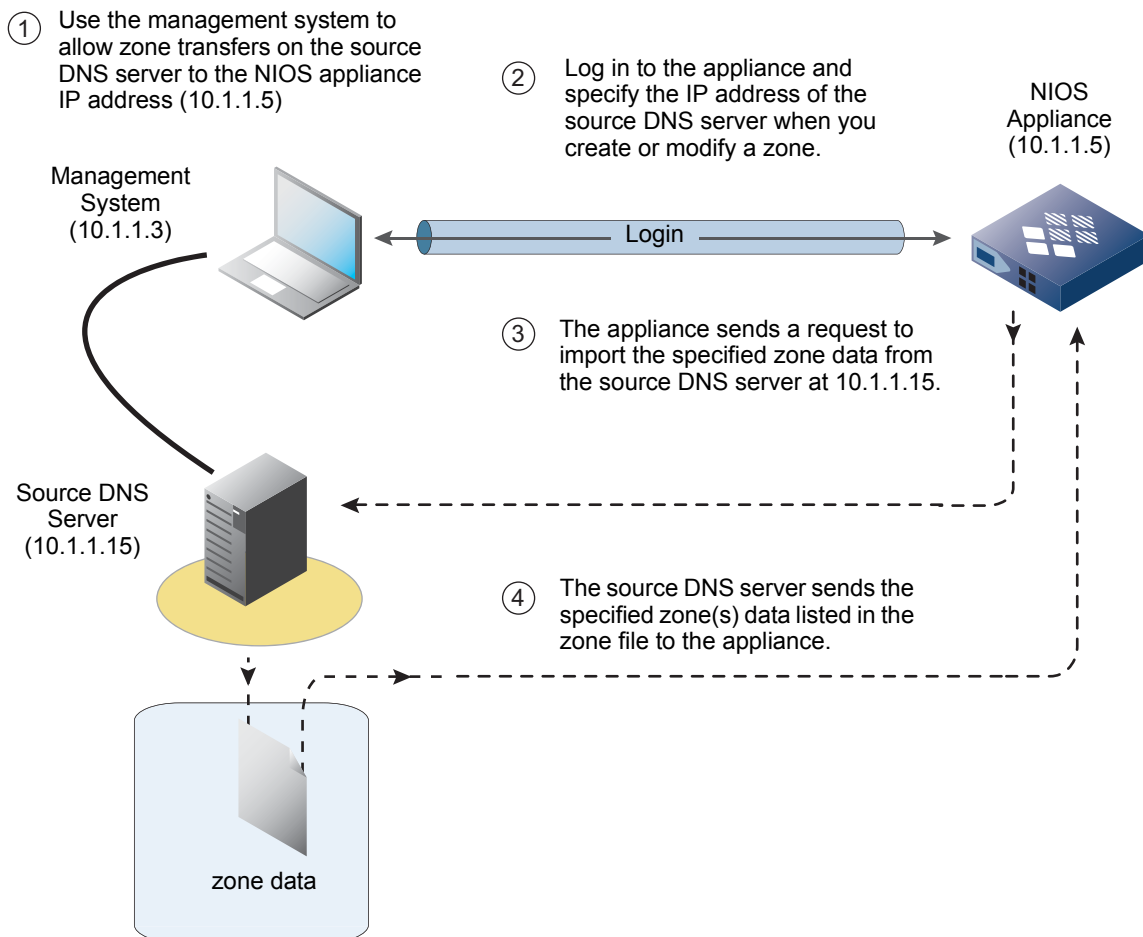
Note: If you apply a name server group to at least one zone or specify it as the default group, you cannot rename or remove it. To rename or remove a group, you must first disassociate it from all zones and unassign it as the default group.

IMPORTING ZONE DATA

Importing zone information alleviates having to manually enter data through the Infoblox GUI. You can import data from existing name servers, as well as from NIOS appliances running version 3.1r4 or later. You can import existing zone data when you create a new zone and when you edit an existing zone. You can import one zone (and its subzones) at a time.

For the remainder of this section, the name server that stores the existing zone data (which is imported) is referred to as the **source** name server (regardless of whether it is a third-party server or another NIOS appliance). The appliance that receives the zone data is referred to as the **destination** appliance. The following illustration shows the import zone data process.

Figure 18.3 Importing Zone Data Process



The appliance imports zone data through a zone transfer. Therefore, the source name server must be authoritative for the zone data being imported. You must also configure the source name server to allow zone transfers to the destination appliance. On the source name server, you might need to modify the *allow-transfer* substatement to include the IP address of the destination appliance prior to importing the data. If you are importing zone data to an HA pair, use the VIP (virtual IP) address shared by the HA pair. For a single independent appliance, use the LAN IP address. If you are importing zone data to a Grid, always use the IP address of the Grid Master.

If the source name server is an Infoblox appliance, you can configure it to allow zone transfers as described in [Enabling Zone Transfers](#) on page 583. Note that a NIOS appliance, acting as the primary name server for a zone, by default allows zone transfers to its secondary name servers.

Note: The appliance does not encode punycode when you import zone data containing punycode. For example, a zone data containing IDNs in punycode is stored in punycode for the data being imported. The data is managed in punycode only.

About Importing Data into a New Zone

When the appliance imports data to a newly created zone, it imports the existing A, CNAME, DNAME, SRV, TXT, MX, PTR, host, and bulk host records, but creates NS (and A records matching that NS record) and SOA records appropriate for the destination server. The NS and SOA records are auto-created when a destination appliance is specified as the primary or secondary name server for the new zone. If the imported zone has extra NS records, they are rewritten to specify the source server as an external secondary. Delegation is also added for any subzones. The subzone records are not imported.

About Importing Data into an Existing Zone

When you import zone data into an existing zone, the zone retains the NS and SOA records automatically created when the zone was originally created and replaces all other records—A, PTR, MX, TXT, SRV, CNAME, DNAME, host, and bulk host. The local appliance also retains subzones and records in the subzones that exist locally. If there are no duplicates, the destination appliance records are retained. If the imported zone has extra NS records, those records change to designate the source server as an external secondary.

Importing Data into Zones

In Grid Manager, you can import zone data when you create the zone using the *Add Authoritative Zone* wizard and when you edit an existing zone. For information on how to add a new zone through the wizard, see [Configuring Authoritative Zones](#) on page 616. The last step of the wizard provides the option to import zone data. The following procedure describes how to import data into an existing zone.

To import data into an existing zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box, and then click **Import Zone** in the Toolbar.
2. In the *Import Zone* dialog box, specify the following:
 - The IP address of the name server from which you want to import data.
 - Optionally, click the **Automatically create Infoblox host records from A records** check box.
3. Click **Import**.

When the local server successfully imports the zone data, a *Confirmation* message appears. If the local server cannot import the zone data, an *Error* message appears, recommending that you verify the correctness of the IP address of the remote server and zone information.

Note: If NIOS resolves the IP address of the imported zone data, an external secondary member is added to the list of name servers with the exact IP address. If NIOS cannot resolve the IP address of the imported zone data, it adds an external secondary member with the IP address 255.255.255.255 to the list of name servers.

CONFIGURING AUTHORITATIVE ZONE PROPERTIES

A zone inherits some of its properties from the Grid or from the member that serves it as a primary or secondary server. When you edit a zone, you can override properties set at the Grid or member level and modify the original zone settings, as well.

To configure authoritative zone properties:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box, and then click the Edit icon.
2. In the *Authoritative Zone* editor, you can do the following in each tab:
 - **General:** Modify the original zone settings, except the zone name.
 - **Name Servers:** Specify primary and secondary servers as described in [Assigning Zone Authority to Name Servers](#) on page 623.
 - **Settings:** Set certain properties if the primary server is a Grid member. If the zone's primary server is an external server, then all these fields, except **Don't use forwarders to resolve queries in subzones**, are read-only with the information derived from the SOA record of the zone.
 - The **Serial Number** field displays the zone's current serial number. You can change the serial number in an SOA record only if the primary server of the zone is a Grid member.
The serial number in the SOA record increments every time the record is modified. This serial number plays a key role when and how zone data is updated via zone transfers. The NIOS appliance allows you to change the serial number (in the SOA record) for the primary server so it is higher than the secondary server, thereby ensuring zone transfers come from the primary server (as they should).
 - Override the Grid or member TTL settings as described in [About Time To Live Settings](#) on page 557.
 - Override the email settings, as described in [Adding an Email Address to the SOA Record](#) on page 560.
 - Change the primary name server that is specified in the SOA MNAME of a zone, as described in [Changing the SOA Name for a Zone](#) on page 624.
 - **Don't use forwarders to resolve queries in subzones:** If the DNS members are configured to use forwarders to resolve queries that they cannot resolve locally, you can select this check box to disable the use of forwarders to resolve queries for data in the subzones.
 - **Queries:** Set restrictions for queries as described in [Controlling DNS Queries](#) on page 570.
 - **Zone Transfers:** Specify to which servers zone transfers are allowed as described in [Enabling Zone Transfers](#) on page 583.
 - **Updates:** Set dynamic DNS update properties as described in [Configuring DNS Servers for DDNS](#) on page 705.
 - **Active Directory:** Set parameters to allow zones to receive GSS-TSIG authenticated DDNS updates from DHCP clients and servers in an AD domain. For information, see [Supporting Active Directory](#) on page 709.
 - **Extensible Attributes:** Define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** Define administrative permissions. For information, see [About Administrative Permissions](#) on page 160.
3. Click **Toggle Expert Mode** if the editor is in basic mode. When the additional tabs appear, you can do the following in each tab:
 - **General:** Click the **Advanced** subtab and view the networks associated with the zone. This tab is visible only if the primary server is a Grid member, a Microsoft server, or unassigned.

If a zone is associated with one or more networks, the IP addresses of its host, A and AAAA records must belong to the associated networks. You cannot change the network associations in this editor. Navigate to the *DHCP Network* editor of the network, to change the zone associations. For information, see [Associating Networks with Zones](#) on page 813.

- **Host Naming:** Set restrictions for host names. For information, see [Specifying Hostname Policies](#) on page 592.
- **Shared Record Groups:** Add shared record groups to a zone. For information, see [About Shared Record Groups](#) on page 680.
- **DNSSEC:** Configure DNSSEC properties. For information, see [Chapter 21, DNSSEC](#), on page 733.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

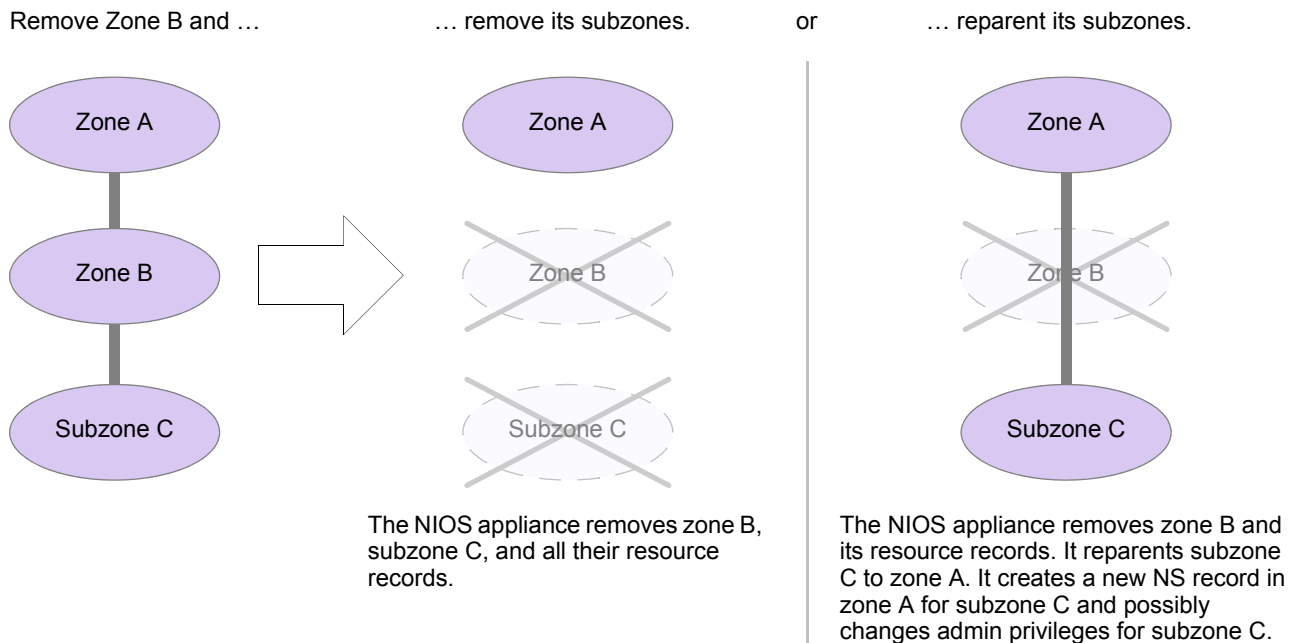
REMOVING ZONES

Depending on the configuration, you may or may not be able to delete or schedule the deletion of a zone and all its contents. Superusers can determine which group of users are allowed to delete or schedule the deletion of a zone and all its contents. For information about how to configure the recursive deletion of zones, see [Configuring Recursive Deletions of Networks and Zones](#) on page 269.

Note that you must have Read/Write permission to all the subzones and resource records in order to delete a zone. The possible effects of removing or re-parenting are illustrated in [Figure 18.4](#).

The appliance puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary. When you restore a parent object from the Recycle Bin, all its contents, if any, are re-parented to the restored parent object. For information about the Recycle Bin, see [Using the Recycle Bin](#) on page 64.


Figure 18.4 Removing or Reparenting Subzones



If you choose to reparent the subzones, be aware of the following caveats and possible effects of the reparenting:

- You cannot remove a zone and reparent its subzones if at least one of the subzones is a delegated zone. You must first remove any delegated subzones, and then you can remove the zone and reparent its subzones.
- If there are AD (Active Directory) subzones (`_msdcs`, `_sites`, `_tcp`, `_udp`, `domaindnszones`, `foresetdnszones`) and you opt to remove the parent zone only, the NIOS appliance reparents all subzones except the AD subzones, which it removes regardless of the removal option you specify.
- The subzone reparenting option is unavailable when you select multiple zones for removal.
- When you remove a zone and reparent its subzones, any subzone that inherited its admin access settings from its previous parent zone (as opposed to having specific access settings for the subzone) now receive their settings from its new parent zone, which might be different. See [Figure 18.5](#).

Figure 18.5 Changed Admin Access Settings after Reparenting Subzones



... the admin access settings for subzone C change because the privileges for its new parent zone (zone A) are different from those of its previous parent zone (zone B).

Before you remove zone B, subzone C inherits a “Deny” admin access setting from zone B. After the removal, subzone C inherits “Read/Write” access from its new parent zone, zone A.

Note that if you set a specific “Deny” admin access privilege for subzone C before removing its parent zone (zone B), subzone C retains its specified “Deny” setting.

Note: Instead of removing a zone, you can also disable it. For more information, see [Enabling and Disabling Zones](#) on page 621.

To remove a zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the check box of the zones you want to delete.
3. Click the Delete icon.
4. Select one of the following. Note that these options appear only if you are allowed to delete zones and all its contents. For information about how to configure this, see [Configuring Recursive Deletions of Networks and Zones](#) on page 269.
 - **Remove zone only:** Select this to remove the zone and all its content. The appliance reparents all subzones to the parent zone of the zone that you want to remove, except for the automatically created AD (Active Directory) subzones.
 - **Remove all subzones:** Select this to remove the selected zone, all its subzones, and all the resource records of the selected zone and its subzones.
5. Click **Yes**.

You can also schedule the deletion for a later time. Click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#) on page 76. For information about scheduling recursive deletions of zones, see [Scheduling Recursive Deletions of Network Containers and Zones](#) on page 76.

RESTORING ZONE DATA

After you import or delete a zone, if you want the original zone back, you can restore it using the Recycle Bin. When you import a zone for the first time, the appliance saves the zone and its resource records as a single object in the Recycle Bin. It keeps the subzones with the zone. See [Restoring Zone Data After a Zone Import Example](#) on page 636.

When you reimport data into a zone, the software saves the zones, its resource records, and the delegated subzones created by the previous import operation in the Recycle Bin. It keeps the subzones (not created during the zone import) with the zone. See [Restoring Zone Data After a Zone Reimport Example](#) on page 637

If the zone import succeeds, the system adds resource records from the source to the target zone. It also adds delegated subzones for the source subzones. If the zone import fails, the system does not create records and delegated subzones. In either case, you can retrieve the original zone and its subzones from the Recycle Bin as follows:

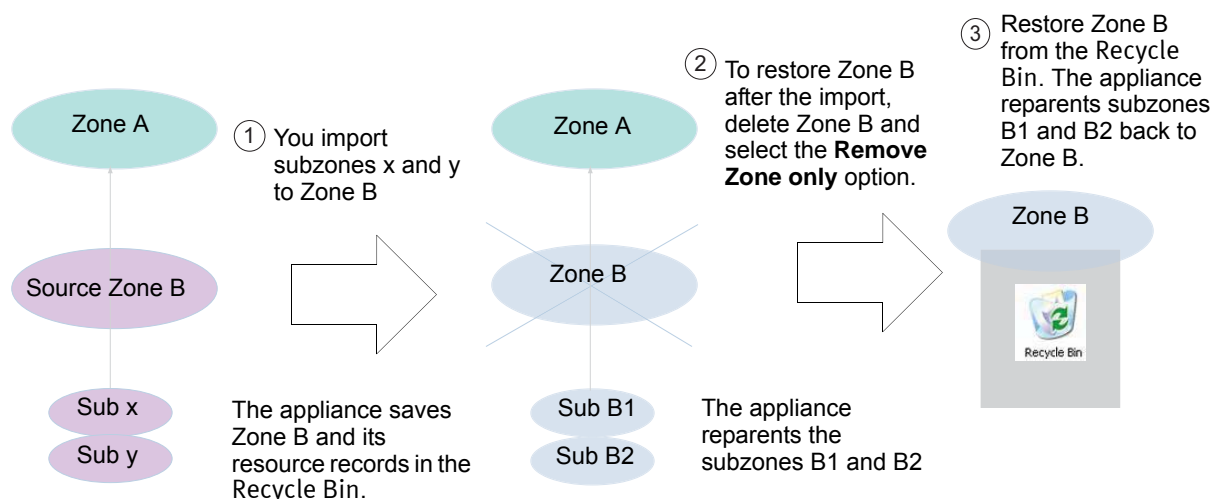
1. Delete the zone using the steps described in the section [Removing Zones](#) on page 634.
2. Select **Remove zone only** to remove the zone and its resource records. The NIOS appliance reparents all subzones to the parent zone of the zone that you remove. Do not select **Remove all subzones**.
Automatically created AD (Active Directory) subzones are an exception. Even if you select **Remove zone only**, the NIOS appliance still removes AD subzones.
3. In the Finder panel, click **Recycle Bin**.
4. Select the zone you want to restore and click the Restore icon.
The zone is restored back to its original state. The resource records are reparented back under it.

Restoring Zone Data After a Zone Import Example

In the example shown in [Figure 18.6](#):

1. import data from a source zone with subzones Sub x and Sub y into zone B with subzones Sub B1 and Sub B2. The appliance stores zone B and its resource records in the Recycle Bin.
To retrieve zone B after the import:
2. Delete subzone B using the **Remove zone only** option.
The appliance reparents subzones Sub B1 and Sub B2 to the Zone A, which is the zone above Zone B.
3. After the import, you can restore zone B from the Recycle Bin. The appliance reparents the subzones Sub B1 and Sub B2 back to zone B.

Figure 18.6 Restoring Zones After a Zone Import



Restoring Zone Data After a Zone Reimport Example

In the example shown in [Figure 18.7](#):

1. You reimport data from the source zone with subzones Sub x and Sub y into zone B with subzones Sub B1 and Sub B2.

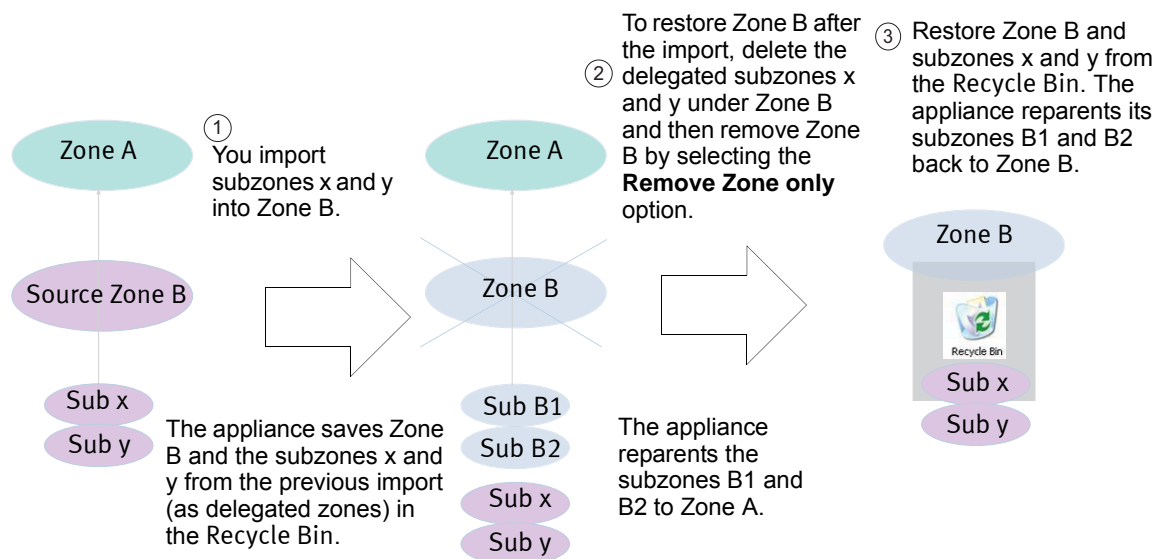
To retrieve zone B after the import:

2. Delete the delegated subzones x and y and then remove subzone B using the **Remove zone only** option.

The appliance stores zone B and its resource records and the previously-imported subzones Sub x and Sub y (as delegated subzones) in the Recycle Bin. It reparents subzones Sub B1 and Sub B2 to the zone above zone B (Zone A).

3. After the import, you can restore zone B and the subzones Sub x and Sub y from the Recycle Bin. The appliance reparents the subzones Sub B1 and Sub B2 back to zone B.

Figure 18.7 Restoring Zones After a Zone Reimport



CONFIGURING DELEGATED, FORWARD, AND STUB ZONES

In addition to authoritative zones, the NIOS appliance allows you to configure delegated, forward, and stub zones. A delegated zone is a zone managed by (delegated to) another name server who owns the authority for the zone. A forward zone is where queries are sent before being forwarded to other remote name servers. A stub zone contains records that identify the authoritative name servers in another zone. This section covers the following topics:

- [Configuring a Delegation](#)
- [Configuring a Forward Zone](#) on page 640
- [Configuring Stub Zones](#) on page 644

Configuring a Delegation

Instead of a local name server, remote name servers (which the local server knows) maintain delegated zone data. When the local name server receives a query for a delegated zone, it either responds with the NS record for the delegated zone server (if recursion is disabled on the local server) or it queries the delegated zone server on behalf of the resolver (if recursion is enabled).

For example, there is a remote office with its own name servers, and you want it to manage its own local data. On the name server at the main corporate office, define the remote office zone as delegated, and then specify the remote office name servers as authorities for the zone.

You can delegate a zone to one or more remote name servers, which are typically the authoritative primary and secondary servers for the zone. If recursion is enabled on the local name server, it queries multiple delegated name servers based on their round-trip times.

You can also configure TTL settings of auto-generated NS records and glue A and AAAA records for delegated zones in forward-mapping, IPv4 reverse-mapping, and IPv6 reverse-mapping zones. For information, see [About Time To Live Settings](#) on page 557.

The delegation must exist within an authoritative zone with a Grid primary server.

Configuring a Delegation for a Forward-Mapping Zone

To create a delegation for a forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the parent zone to open it.
Grid Manager displays the **Records** and **Subzones** tabs of the zone.
3. From the **Subzones** tab, click the Add icon -> **Zone** -> **Add Delegation**.
4. In the *Add Delegation* wizard, specify the following:
 - **Name:** This field displays a dot followed by the domain name of the current zone. Enter one or more labels before the dot to specify the domain name of the subzone.
 - **DNS View:** This field displays only when there is more than one DNS view in the network view. Displays the DNS view of the current zone.
 - **Comment:** Optionally, enter additional text about the zone.
 - **Disable:** Click this check box to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#) on page 621
 - **Lock:** Click this check box to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#) on page 621.
5. Click **Next** to define the name servers for the zone.
6. In the *Name Servers* panel, click the Add icon and specify the following information:
 - **Name:** Enter the name of a remote name server to which you want the local server to redirect queries for data for the zone. This is a name server that is authoritative for the delegated zone.
 - **Address:** Enter the IP address of the delegated server.

7. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes as described in [Using Extensible Attributes](#) on page 332.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Configuring a Delegation for a Reverse-Mapping Zone

To create a delegation for a reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the parent zone to open it.
Grid Manager displays the **Records** and **Subzones** tabs of the zone.
3. From the **Subzones** tab, click the Add icon -> **Zone** -> **Add Delegation**.
4. In the *Add Delegation* wizard, specify the following:
 - **IPv4 Network:** This field displays if you are creating a delegation zone for an IPv4 reverse-mapping zone. Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the Netmask drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.
To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the following fields:
RFC 2317 Prefix: Enter a prefix in this field. Prefixes can include alphanumeric characters.
Allow manual creation of PTR records in parent zone: Select this check box to allow users to create labels that correspond to IP addresses in the delegated address space in the parent zone.
For information about RFC 2317, see [Specifying an RFC 2317 Prefix](#) on page 618.
 - **IPv6 Network Prefix:** This field displays if you are creating a delegation zone for an IPv6 reverse-mapping zone. Enter the IPv6 prefix for the address space for which you want to define the reverse-mapping zone and select the prefix length from the drop-down list.
 - **Name:** This field displays a dot followed by the domain name of the current zone. Enter one or more labels before the dot to specify the domain name of the subzone.
 - **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
 - **Comment:** Optionally, enter additional text about the zone.
 - **Disable:** Select this option to temporarily disable this zone.
 - **Lock:** Select this option to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
5. Click **Next** to define the name servers for the zone.
6. In the *Name Servers* panel, click the Add icon and specify the following information:
 - **Name:** Enter the name of a remote name server to which you want the local server to redirect queries for data for the zone. This is a name server that is authoritative for the delegated zone.
 - **Address:** Enter the IP address of the delegated server.
7. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes as described in [Using Extensible Attributes](#) on page 332.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Configuring a Forward Zone

When you want to forward queries for data in a particular zone, define the zone as a forward zone and specify one or more name servers that can resolve queries for the zone. For example, define a forward zone so that the NIOS appliance forwards queries about a partner's internal site to a name server, which the partner hosts, configured just for other partners to access.

You can override the default forwarders for a forward-mapping zone at a Grid member level and configure custom forwarders. In other words, each Grid member can have its own forwarders for the forward zone. For example: a forward-mapping zone `foo.com` served by two Grid members M1 and M2 with M1 forwarding queries to 10.1.0.1 and 10.1.0.2 and M2 forwarding queries to 90.3.3.3 and 90.4.4.1. Note that the Grid member uses the default forwarders unless you override them at any level. For more information about domains and zones, see [About Domains and Zones](#) on page 622.

Note: The use of a forward zone is different from that of a forwarder. (A forwarder is a name server that performs recursive lookups on behalf of the name servers that forward queries to it. For more information, see [Using Forwarders](#) on page 569.) A NIOS appliance forwards queries to the name server of a forward zone because the name server can resolve queries for the zone. A NIOS appliance forwards queries to a forwarder regardless of zones.

Note that a name server can have only one definition for a zone in any given DNS view; a forward zone cannot be configured on a member that already has a zone with the same domain name configured on it in the same DNS view. To configure a forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Forward Zone**.
2. In the *Add Forward Zone* wizard, click **Add a forward forward-mapping zone** and click **Next**.
3. Enter the following information, and then click **Next**:
 - **Name:** Enter the domain name of the zone for which you want the NIOS appliance to forward queries.
 - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select the DNS view of the forward zone.
 - **Comment:** Enter a descriptive comment.
 - **Disable:** Click this check box to temporarily disable this zone.
 - **Lock:** Click this check box to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
4. In the Default Zone Forwarders section, click the Add icon and specify the default servers to which the NIOS appliance forwards queries for the zone:
 - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
 - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
 - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
5. In the Members section, you can select a Grid member or multiple members to serve the forward-mapping zone and use the default forwarders or you can override default forwarders and configure custom forwarders for the Grid members.
 - **Add:** Click the Add icon to select the NIOS appliance on which the forward zone is configured. For an independent deployment, select the local appliance (it is the only choice). If there are multiple Grid members, the *Member Selector* dialog box is displayed. Select the required member from the list and click **OK**.
 - **Name:** Displays the name of the Grid member.
 - **IPv4 Address:** Displays the IPv4 address of the Grid member.
 - **IPv6 Address:** Displays the IPv6 address of the Grid member.

- **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
- **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.

Note: Skip the following two steps if you want to use the default forwarders.

6. Select a member and click the Edit icon.
7. In the *Edit Per-Member Forwarders* editor, select the **Override Default Forwarders** check box to override the default forwarders. The Default Zone Forwarders table becomes available only after you select the **Override Default Forwarders** check box. Click the Add icon to specify the servers to which the NIOS appliance forwards queries for the zone:
 - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
 - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
 - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
 - Save the configuration. After successfully saving the configuration, the **Override Default Forwarders** column displays **Yes** and the **Custom Forwarders** column displays the IP address of the forwarders.
8. Save the configuration, or click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#) on page 332, and then optionally proceed to the next step where you define admin permissions as defined in [About Administrative Permissions](#) on page 160.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

9. Click **Restart** if it appears at the top of the screen.

To configure a forward IPv4 reverse-mapping zone:

1. From the **Data Management** tab, select the **Zones** tab, expand the Toolbar and click **Add -> Zone -> Add Forward Zone**.
2. In the *Add Forward Zone* wizard, click **Add a forward IPv4 reverse-mapping zone** and click **Next**.
3. Enter the following information, and then click **Next**:
 - **IPv4 Network:** Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the **Netmask** drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.
To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the **RFC 2317 Prefix** field. Enter a prefix in the text field. Prefixes can be alphanumeric characters. For information, see [Specifying an RFC 2317 Prefix](#) on page 618.
 - or
 - **Name:** Enter the domain name of the reverse-mapping zone.
 - **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
 - **Comment:** Optionally, enter additional information about the zone.
 - **Disable:** Click this check box to temporarily disable this zone.
 - **Lock:** Click this check box to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes.
4. In the Default Zone Forwarders section, click the Add icon and specify the servers to which the NIOS appliance forwards queries for the zone:
 - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.

- **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
 - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
5. In the **Members** section, you can select a Grid member or multiple members to serve the forward-mapping zone and use the default forwarders or you can override default forwarders and configure custom forwarders for the Grid members.
- **Add:** Click the Add icon to select the NIOS appliance on which the forward zone is configured. For an independent deployment, select the local appliance (it is the only choice). If there are multiple Grid members, the *Member Selector* dialog box is displayed. Select the required member from the list and click **OK**.
 - **Name:** Displays the name of the Grid member.
 - **IPv4 Address:** Displays the IPv4 address of the Grid member.
 - **IPv6 Address:** Displays the IPv6 address of the Grid member.
 - **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
 - **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.

Note: Skip the following two steps if you want to use the default forwarders.

6. Select a member and click the Edit icon.
7. In the *Edit Per-Member Forwarders* editor, select the **Override Default Forwarders** check box to override the default forwarders. The Default Zone Forwarders table becomes available only after you select the **Override Default Forwarders** check box. Click the Add icon to specify the servers to which the NIOS appliance forwards queries for the zone:
- **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
 - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
 - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
 - Save the configuration. After successfully saving the configuration, the **Override Default Forwarders** column displays **Yes** and the **Custom Forwarders** column displays the IP address of the forwarders.
- To configure forwarders for multiple members, repeat the steps for each Grid member.

8. Save the configuration, or click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#) on page 332.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

9. Click **Restart** if it appears at the top of the screen.

To configure a forward IPv6 reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Forward Zone**.
2. In the *Add Forward Zone* wizard, click **Add a forward IPv6 reverse-mapping zone** and click **Next**.
3. Enter the following zone information:
 - **IPv6 Network Address:** Enter the 128-bit IPv6 address for the address space for which you want to define the reverse-mapping zone. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to

2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Choose the network prefix that defines the IPv6 network address space.

or

Name: Enter the domain name of the reverse-mapping zone.

- **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
- **Comment:** Enter a descriptive comment about the zone.
- **Disable:** Click this check box to temporarily disable this zone.
- **Lock:** Click this check box to lock the zone so that you can make changes to it, and also prevent others making conflicting changes.

4. In the Default Zone Forwarders section, click the Add icon and specify the servers to which the NIOS appliance forwards queries for the zone:
 - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
 - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
 - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
5. In the Members section, you can select a Grid member or multiple members to serve the forward-mapping zone and use the default forwarders or you can override default forwarders and configure custom forwarders for the Grid members.
 - **Add:** Click the Add icon to select the NIOS appliance on which the forward zone is configured. For an independent deployment, select the local appliance (it is the only choice). If there are multiple Grid members, the *Member Selector* dialog box is displayed. Select the required member from the list and click **OK**.
 - **Name:** Displays the name of the Grid member.
 - **IPv4 Address:** Displays the IPv4 address of the Grid member.
 - **IPv6 Address:** Displays the IPv6 address of the Grid member.
 - **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
 - **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.

Note: Skip the following two steps if you want to use the default forwarders.

6. Select a member and click the Edit icon.
7. In the *Edit Per-Member Forwarders* editor, select the **Override Default Forwarders** check box to override the default forwarders. The Default Zone Forwarders table becomes available only after you select the **Override Default Forwarders** check box. Click the Add icon to specify the servers to which the NIOS appliance forwards queries for the zone:
 - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
 - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
 - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
 - Save the configuration. After successfully saving the configuration, the **Override Default Forwarders** column displays **Yes** and the **Custom Forwarders** column displays the IP address of the forwarders.

To configure forwarders for multiple members, repeat the steps for each Grid member.

8. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#) on page 332 or **save the configuration** and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Configuring Stub Zones

A stub zone contains records that identify the authoritative name servers in the zone. It does not contain resource records for resolving IP addresses to hosts in the zone. Instead, it contains the following records:

- SOA (Start of Authority) record of the zone
- NS (name server) records at the apex of the stub zone
- A (Address) records that map the name servers to their IP addresses

Stub zones, like secondary zones, obtain their records from other name servers. Their records are read only; therefore, administrators do not manually add, remove, or modify the records.

Stub zone records are also periodically refreshed, just like secondary zone records. However, secondary name servers contain a complete copy of the zone data on the primary server. Therefore, zone transfers from a primary server to a secondary server, or between secondary servers, can increase CPU usage and consume excessive bandwidth. A name server hosting a stub zone maintains a much smaller set of records; therefore, updates are less CPU intensive and consume less bandwidth.

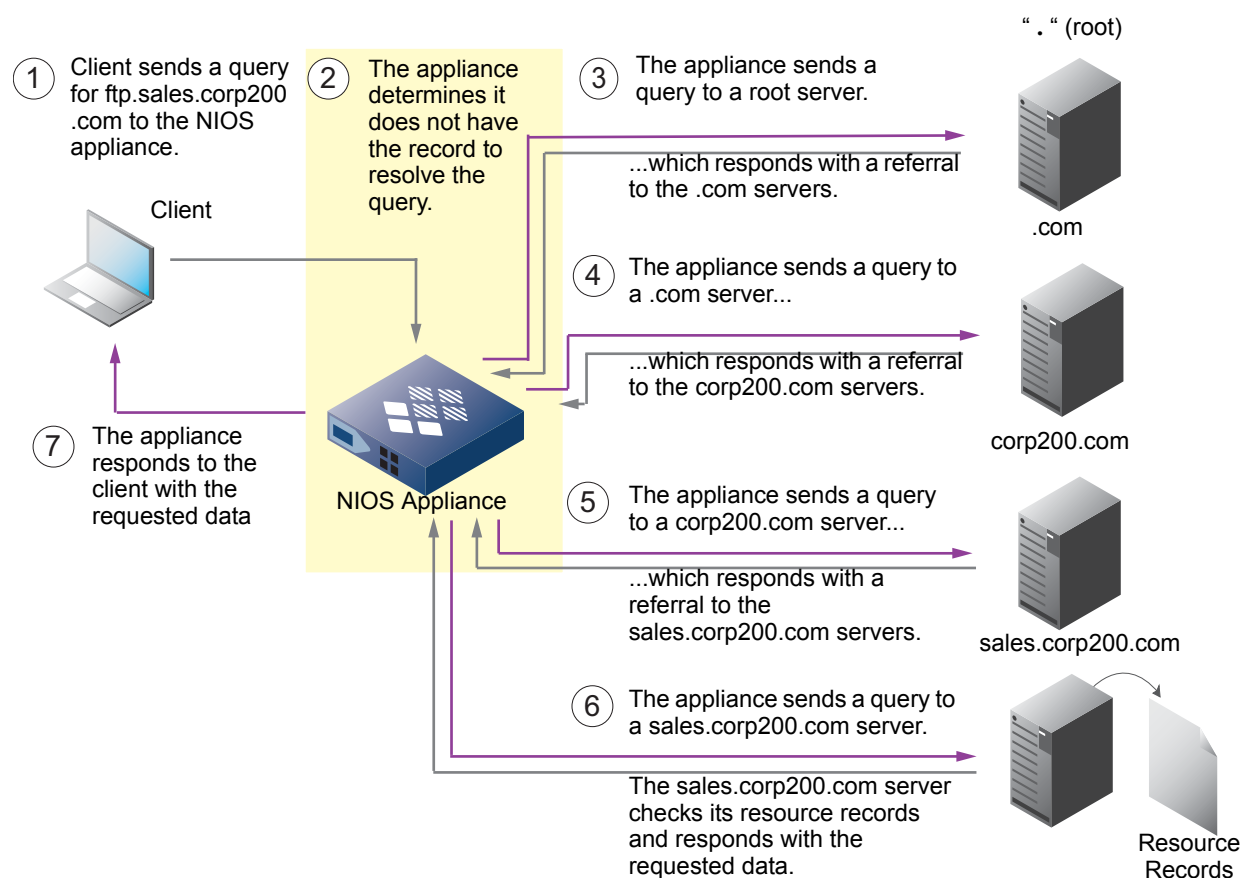
When a name server hosting a stub zone receives a query for a domain name that it determines is in the stub zone, the name server uses the records in the stub zone to locate the correct name server to query, eliminating the need to query the root server.

[Figure 18.8](#) and [Figure 18.9](#) illustrate how the NIOS appliance resolves a query for a domain name for which it is not authoritative. [Figure 18.8](#) illustrates how the appliance resolves a query when it does not have a stub zone.

[Figure 18.9](#) illustrates how the appliance resolves the query with a stub zone.

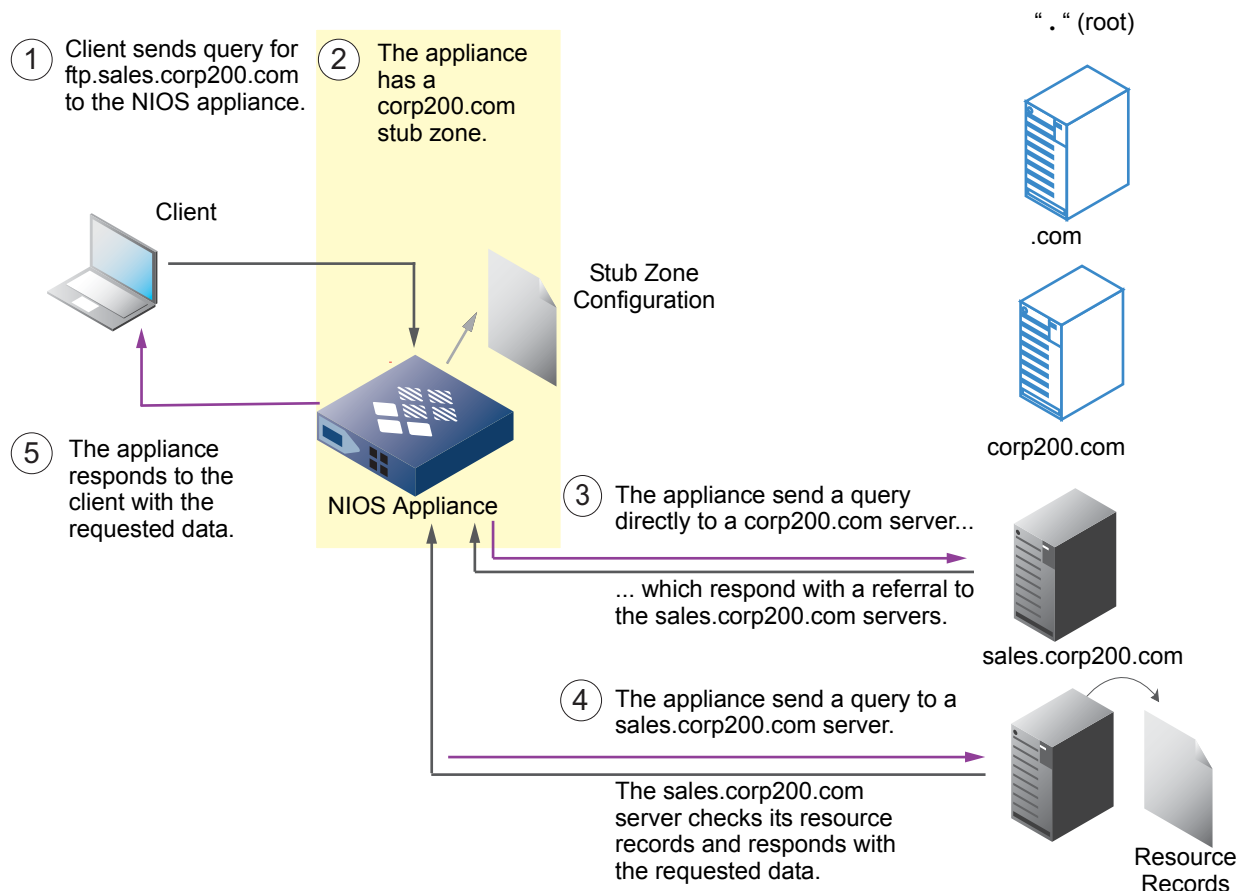
In [Figure 18.8](#), a client sends a query for ftp.sales.corp200.com to the NIOS appliance. When the appliance receives the request from the client, it checks if it has the data to resolve the query. If the appliance does not have the data, it tries to locate the authoritative name server for the requested domain name. It sends nonrecursive queries to a root name server and to the closest known name servers until it learns the correct authoritative name server to query.

Figure 18.8 Processing a Query without a Stub Zone



In [Figure 18.9](#), when the NIOS appliance receives the request for the domain name in corp200.com, it determines it does not have the resource records to resolve the query. It does, however, have a list of the authoritative name servers in the stub zone, corp200.com. The appliance then sends a query directly to the name server in corp200.com.


Figure 18.9 Processing a Query with a Stub Zone



Stub zones facilitate name resolution and alleviate name server traffic in your network. For example, the client in the previous examples is in corp100.com. The corp100.com and corp200.com zones are partners, and send all their communications through a VPN tunnel, as shown in [Figure 18.10](#) on page 647. The firewall protecting corp100.com is configured to send all messages for the 10.2.2.0/24 network through the VPN tunnel. Infoblox_A hosts the stub zone for corp200.com. Therefore, when the host in corp100.com sends a query for ftp.sales.corp200.com, Infoblox_A obtains the IP address of Infoblox_B (10.2.2.7) from its stub zone records and sends the query to the firewall protecting corp100.com.

Because the destination of the query is in the 10.2.2.0/24 network, the firewall (configured to encrypt all traffic to the network) sends the request through a VPN tunnel to Infoblox_B. Infoblox_B resolves the query and sends back the response through the VPN tunnel. All name server traffic went through the VPN tunnel to the internal servers, bypassing the root servers and external name servers.

Figure 18.10 Stub Zone Configuration



In parent-child zone configurations, using stub zones also eases the administration of name servers in both zones. For example, as shown in [Figure 18.10](#), sales.corp200.com is a child zone of corp200.com. On the corp200.com name servers, you can create either a delegated zone or a stub zone for sales.corp200.com.

When you create a delegated zone, you must first specify the name servers in the delegated zone and manually maintain information about these name servers. For example, if the administrator in sales.corp200.com changes the IP address of a name server or adds a new name server, the sales.corp100.com administrator must inform the corp200.com administrator to make the corresponding changes in the delegated zone records.

If, instead, you create a stub zone for sales.corp200.com, you set up the stub zone records once, and updates are then done automatically. The name servers in corp200.com that are hosting a stub zone for sales.corp200.com automatically obtain updates of the authoritative name servers in the child zone.

In addition, a name server that hosts a stub zone can cache the responses it receives. Therefore, when it receives a request for the same resource record, it can respond without querying another name server.

Creating Stub Zones

When you create a stub zone on the NIOS appliance, you specify the following:

- The Grid member that is hosting the stub zone
You can specify multiple appliances if you want the stub zones on multiple name servers. If you do, the appliances store identical records about the stub zone.
- The IP address of the primary server(s) that the NIOS appliance can query in the stub zone
The primary server can be a Grid member or an external primary server. If you specify multiple primary servers, the appliance queries the primary servers, starting with the first server on the list.

The primary server and the name server hosting the stub zone can belong to the same Grid, as long as the authoritative zone and the stub zone are in different DNS views. You cannot configure one zone as both authoritative and stub in the same view.

After you create a stub zone, the NIOS appliance does the following:

1. It sends a query to the primary server for the SOA (Start of Authority) record of the stub zone.
The primary server returns the SOA record.
2. Then, it sends a query for the NS (name server) records in the zone.
The primary server returns the NS records and the A (address) records of the name servers. (These A records are also called glue records.)
If the primary server is a NIOS appliance, you might have to manually create the A record and add it to the stub zone. A NIOS appliance that is the primary server for a zone always creates an NS record, but does not always create an A record.
 - The appliance automatically creates an A record when its host name belongs to the name space of the zone. For example, if the zone is corp100.com and the primary server host name is server1.corp100.com, the appliance automatically creates the NS and A records and sends these records when it is queried by the stub zone name server.
 - The appliance does not automatically create an A record when its host name is in a name space that is different from the zone. For example, if the zone is corp200.com and the primary server host name is server1.corp100.com, then the appliance creates the NS record only and sends it when it is queried by the stub zone name server. In this case, you must manually create the A record.

Maintaining Stub Zones

The NIOS appliance maintains the stub zone records and updates them based on the values in the SOA record as follows:

- The refresh interval indicates when the appliance sends a discrete query to the primary name server for the stub zone. The appliance learns about any changes in the stub zone and updates the NS and A records in the stub zone accordingly.
- If the update fails, the retry interval indicates when the appliance resends a discrete query.
- If the query continues to fail, the expiry value indicates when the appliance stops using the zone data.

Adding Stub Zones

To add a stub zone, you must identify the Infoblox appliance that hosts the stub zone, and provide the IP address of the primary server.

You can also add stub zones for Microsoft servers that are managed by Grid members. For information, see [Managing Microsoft Windows Servers](#) on page 953.

You can configure a stub zone for forward mapping or reverse mapping zones.

To add a forward-mapping stub zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Add Stub Zone* wizard, click **Add a stub forward-mapping zone** and click **Next**.
3. Specify the following, and then click **Next**:
 - **Name:** Enter the name for the stub zone.
 - **Comment:** Enter a useful comment, such as the admin to contact for the stub zone.
 - **Disable:** Click this check box to temporarily disable this zone.
 - **Lock:** Click this check box to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes.
4. In the *Master Name Servers* panel, click the Add icon and enter the **Name** and **IP Address** of the primary server in the stub zone, and then click **Next**.

If the primary server is a Grid member, you must enter the host name and IP address of the Grid member. The NIOS appliance does not validate these entries. Therefore, if you change the IP address of a Grid member listed here, you must update the Grid member information in this list as well.

You can specify multiple primary servers for redundancy. If the primary server is a NIOS appliance, the appliance must have the Minimal Response feature disabled so it can propagate the data to the stub server. For information about the Minimal Response feature, see [Specifying Minimal Responses](#) on page 565.

- Optionally, click the **Don't use forwarders to resolve queries in subzones** check box to indicate that the name servers hosting the stub zone must not use forwarders to resolve queries for domain names in the stub zone or in its subzones.

5. In the *Name Servers* panel, click the Add icon and select one of the following:

- **Add Infoblox Member:** Select this and select the Grid member that hosts the stub zone.
- **Add Microsoft Server:** Select this and select the Microsoft server that hosts the stub zone.

6. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#) on page 332.

7. Save the configuration and click **Restart** if it appears at the top of the screen

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

You can define two types of reverse-mapping stub zones, one for IPv4 addresses and one for IPv6 addresses.

To configure an IPv4 reverse-mapping stub zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Add Stub Zone* wizard, click **Add a stub IPv4 reverse-mapping zone** and click **Next**.

3. Specify the following:

- **IPv4 Network:** Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the **Netmask** drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.

To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the **RFC 2317 Prefix** field. Enter a prefix in the text field. Prefixes can be alphanumeric characters. For information, see [Specifying an RFC 2317 Prefix](#) on page 618.

or

Name: Enter the domain name of the reverse-mapping zone.

- **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
- **Comment:** Optionally, enter additional information about the zone.
- **Disable:** Click this check box to temporarily disable this zone.
- **Lock:** Click this check box to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes.

4. In the *Master Name Servers* panel, click the Add icon and enter the **Name** and **IP Address** of the primary server in the stub zone, and then click **Next**.

If the primary server is a Grid member, you must enter the host name and IP address of the Grid member. The NIOS appliance does not validate these entries. Therefore, if you change the IP address of a Grid member listed here, you must update the Grid member information in this list as well.

You can specify multiple primary servers for redundancy. If the primary server is a NIOS appliance, the appliance must have the Minimal Response feature disabled so it can propagate the data to the stub server. For information about the Minimal Response feature, see [Specifying Minimal Responses](#) on page 565.

- Optionally, click the **Don't use forwarders to resolve queries in subzones** check box to indicate that the name servers hosting the stub zone should not forward queries that end with the domain name of the stub zone to any configured forwarders.

5. In the *Name Servers* panel, click the Add icon and select one of the following:

- **Add Infoblox Member:** Select this and select the Grid member that hosts the stub zone.
- **Add Microsoft Server:** Select this and select the Microsoft server that hosts the stub zone.

6. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#) on page 332.
7. Save the configuration and click **Restart** if it appears at the top of the screen
or
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

To configure an IPv6 reverse-mapping stub zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Add Stub Zone* wizard, click **Add a stub IPv6 reverse-mapping zone** and click **Next**.
3. Specify the following:
 - **IPv6 Network Prefix** and **Prefix Length**: Enter the 128-bit IPv6 address for the address space for which you want to define the reverse-mapping zone. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. You can enter a slash and prefix length in the **IPv6 Network Prefix** field or you can choose a value from the **Prefix Length** drop-down list.
 - or
 - **Name**: Enter the domain name of the reverse-mapping zone.
 - **DNS View**: This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
 - **Comment**: Enter a descriptive comment about the zone.
 - **Disable**: Click this check box to temporarily disable this zone.
 - **Lock**: Click this check box to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
4. In the *Master Name Servers* panel, click the Add icon and enter the **Name** and **IP Address** of the primary server in the stub zone, and then click **Next**.
If the primary server is a Grid member, you must enter the host name and IP address of the Grid member. The NIOS appliance does not validate these entries. Therefore, if you change the IP address of a Grid member listed here, you must update the Grid member information in this list as well.
You can specify multiple primary servers for redundancy. If the primary server is a NIOS appliance, the appliance must have the Minimal Response feature disabled so it can propagate the data to the stub server. For information about the Minimal Response feature, see [Specifying Minimal Responses](#) on page 565.
 - Optionally, click the **Don't use forwarders to resolve queries in subzones** check box to indicate that the name servers hosting the stub zone should not forward queries that end with the domain name of the stub zone to any configured forwarders.
5. In the *Name Servers* panel, click the Add icon and select one of the following:
 - **Add Infoblox Member**: Select this and select the Grid member that hosts the stub zone.
 - **Add Microsoft Server**: Select this and select the Microsoft server that hosts the stub zone.
6. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#) on page 332.
7. Save the configuration and click **Restart** if it appears at the top of the screen
or
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [About Extensible Attributes](#) on page 322.

Viewing SOA Records

The timer values in the SOA record determine when the stub zone records are updated.

To view zone SOA record values:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box, and then click the Edit icon.
2. In the *Stub Zone* editor, click **Settings** to view the following:
 - **Serial number**—The number used by stub DNS servers to check if the zone has changed. If the serial number is higher than what the stub server currently has, a query is initiated. This number is automatically increased when changes are made to the zone or its records.
 - **Primary Name Server**—The domain name for the primary DNS server for the zone. The zone should contain a matching NS record.
 - **E-mail Address**—The e-mail address of the person responsible for maintaining the zone.
 - **Refresh**—The time lapse between checks the stub server makes for changes to the zone.
 - **Retry**—The time lapse after which the stub server checks for changes if the first refresh fails.
 - **Expire**—The time period the zone remains valid after repeated failures to refresh.
 - **Default TTL**: Specifies how long a name server can cache the record.
 - **Negative-caching TTL**: Specifies how long a name server caches negative responses from the name servers that are authoritative for the zone.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuration Example: Configuring a Stub Zone in a Grid

This example illustrates how to configure a stub zone and assign it to a Grid member. You configure a Grid, Corp100, with a single Grid Master and Grid member. The Grid member, member1.corp100.com, is the primary name server for the corp100.com zone in the internal view. The Grid Master, gm-corp100.com, hosts the stub zone for corp100.com in the external view. Thus, when the Grid Master receives a query for the corp100.com zone, it sends it directly to member1.corp100.com, the primary name server for the zone.

In this example, you configure the following:

1. Turn off minimal responses on member1.corp100.com, the primary name server for the corp100.com zone. See [Disable Minimal Responses](#).
2. Create the internal and external views. See [Create the Views](#).
3. Create the corp100.com authoritative zone and stub zone. See [Create the Zones](#).

Disable Minimal Responses

After you create the Grid, turn off minimal responses for member1.corp100.com. Disabling minimal responses ensures that member1.corp100.com propagates the required data to the server hosting the stub zone.

1. From the **Data Management** tab, select the **DNS** tab, click **Members** -> **member1.corp100.com** check box -> Edit icon.
2. In the *Member DNS Configuration* editor, click the **General** -> **Basic** tab.
3. Clear the **Return minimal responses** check box.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Create the Views

Create the internal and external views. To create each view:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add DNS View**.
2. In the *Add DNS View* wizard, enter the name of the view. In this example, enter either **External** or **Internal**.
3. Click **Save & New** and create the other DNS view.

Create the Zones

Create the corp100.com zone in the internal view and assign member1.corp100.com as the Grid primary server:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Auth Zone**.
2. In the *Forward Authoritative Zone* wizard, do the following:
 - Select **Add an authoritative forward-mapping zone** and click **Next**.
 - Enter the zone name, **corp100.com** and select the **Internal** DNS view. Click **Next**.
 - Select **Use this set of name servers** and select **member1.corp100.com** as the Grid primary server.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create the zone, you can view the NS and A records which were automatically created.

Create the stub zone, corp100.com, in the external view, assign gm-corp100.com as the stub member and member1.corp100.com as the stub primary server.

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Stub Zone* wizard, do the following:
 - Select **Add a stub forward-mapping zone** and click **Next**.
 - Enter the name of the stub zone, **corp100.com** and select the **External** DNS view. Click **Next**.
 - In the *Master Name Servers* panel, click the Add icon and enter the following for the primary name server, and then click **Next**:
 - Name:** member1.corp100.com
 - Address:** 10.35.0.222
 - In the *Name Servers* panel, click the Add icon and select **gm-corp100.com**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create the stub zone, the server hosting the stub zone, gm-corp100.com, sends queries to the primary server, member1.corp100.com, for the SOA and NS records. member1.corp100.com then returns its NS records and A (address) records.

VIEWING ZONES

To list zones, navigate to the **Data Management** tab -> **DNS** tab -> **Zones** panel. If there is more than one DNS view in the Grid, this panel lists the DNS views. Select a DNS view to list its zones. (For information, see [Listing DNS Views](#) on page 611.)

- Click **Toggle flat view** to display a flat list of all the zones in the view.
- Click **Toggle hierarchical view** to display only the apex zones.

In the hierarchical view, you can see one entry for the host that represents the entire host object. In a host record, there can be multiple DNS resource records (A, PTR, CNAME) and some DHCP data (fixed addresses) as well. In the flat view, each of the DNS resource records in the host are listed separately.

For example, the host called server1.infoblox.com contains 2 A records and an ALIAS (which is a host naming convention for CNAME records). If you view the infoblox.com zone using the hierarchical view option, you will see one entry host for server1.infoblox.com. In the flat view, you will see three records (one for each IP address/A record, and one host Alias for the CNAME). In the flat view, you cannot delete one piece of the host record. You can edit the host record and you can remove information. Deleting host records deletes the entire host record only.

This panel displays the following information for each zone, by default:

- **Name:** The domain name of the zone.
- **MS Sync Server:** When a zone is served by multiple Microsoft servers, this column shows which Microsoft server is actually performing the synchronization of that zone with the Grid.
- **Type:** The zone type. Possible values are **Authoritative**, **Forward**, **Stub** and **Delegation**.
- **Comment:** Comments that were entered for the zone.
- **Site:** Values that were entered for this pre-defined attribute.

You can also display the following columns:

- **Locked:** Displays **Yes** when a zone is locked by an admin, and displays **No** when the zone is unlocked.
- **Function:** Indicates whether the zone is a forward-mapping, or an IPv4 or IPv6 reverse-mapping zone.
- **Disabled:** This field displays **Yes** if the zone is disabled. Otherwise, this field displays **No**.
- **Signed:** This field displays **Yes** if the zone is a DNSSEC-signed zone. Otherwise, this field displays **No**.

You can do the following:

- List the resource records and subzones of a DNS zone.
 - Click a DNS zone name.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Edit the properties of a DNS zone.
 - Click the check box beside a DNS zone, and then click the Edit icon.
- Delete a DNS zone.
 - Click the check box beside a DNS zone, and then click the Delete icon.
- Export the list of DNS zones to a .csv file.
 - Click the Export icon.
- Print the list of DNS zones.
 - Click the Print icon.



Chapter 19 DNS Resource Records

This chapter provides general information about Infoblox host records and DNS resource records. The topics in this chapter include:

- [*About Bulk Hosts*](#) on page 656
 - [*Specifying Bulk Host Name Formats*](#) on page 656
 - [*Before Defining Bulk Host Name Formats*](#) on page 656
 - [*Adding Bulk Hosts*](#) on page 659
- [*Managing Resource Records*](#) on page 660
 - [*Managing A Records*](#) on page 660
 - [*Managing NS Records*](#) on page 662
 - [*Managing AAAA Records*](#) on page 662
 - [*Managing PTR Records*](#) on page 664
 - [*Managing MX Records*](#) on page 665
 - [*Managing SRV Records*](#) on page 666
 - [*Managing TXT Records*](#) on page 668
 - [*Managing CNAME Records*](#) on page 669
 - [*Managing DNAME Records*](#) on page 671
 - [*Managing NAPTR Records*](#) on page 676
 - [*Managing LBDN Records*](#) on page 678
 - [*Modifying, Disabling, and Deleting Host and Resource Records*](#) on page 679
- [*About Shared Record Groups*](#) on page 680
 - [*Shared Records Guidelines*](#) on page 681
 - [*Configuring Shared Record Groups*](#) on page 681
 - [*Managing Shared Resource Records*](#) on page 683
 - [*Managing Associated Zones*](#) on page 685
 - [*Configuration Example: Configuring Shared Records*](#) on page 686

ABOUT BULK HOSTS

If you need to add a large number of A and PTR records, you can have the NIOS appliance add them as a group and automatically assign host names based on a range of IP addresses and the host name format you specify. Such a group of records is called a *bulk host*, which the appliance manages and displays as a single bulk host record.

Specifying Bulk Host Name Formats

Bulk host name formats provide a flexible way to define bulk host names. You create multiple bulk host formats at the Grid level. Either select from the default bulk host formats or create your own. You can specify a different format for each bulk host. When you assign a bulk host name format to a bulk host in a zone, the system applies the zone's host name policy to it.

A bulk host name consists of a prefix, a suffix, and the name of the domain to which the host belongs. The prefix can contain any printable character that complies with the zone host name policy. It can also be blank. The suffix is derived from an IP address in the bulk host IP address range. The appliance also supports IDNs for bulk host names. You can use IDNs or their punycode representations while creating bulk hosts.

The following table summarizes how the appliance displays bulk host names that contain IDNs:

Input	NIOS Displays...	NIOS DNS Domain (Punycode in the GUI)	Conversion Guidelines
hello	hello	hello	No conversion
привітання	привітання	xn--80adk5aaihr3f9e	IDN to punycode
xn--80adk5aaihr3f9e	xn--80adk5aaihr3f9e	xn--80adk5aaihr3f9e	No conversion
\xyz format	\xyz format	\xyz format	No conversion

The suffix format is a string of ASCII characters that uses \$ (unpadded) or # (zero-padded) followed by 1,2,3,4 to refer to the first, second, third, or fourth IP address octet; it uses \$1,\$2,\$3,\$4 or #1,#2,#3,#4. **\$2** refers to the second unpadded octet and **#4** refers to the fourth zero-padded octet. For example:

The prefix of a bulk host = *info*

IP address = *10.19.32.133*

Domain name = *infoblox.com*.

If you specify the default four-octet format *-\$1-\$2-\$3-\$4*, the bulk host name is *info-10-19-23-133.infoblox.com*.

If you specify a custom name format such as **#1*#2*#3*#4*, the bulk host name is *info*010*019*023*133.infoblox.com*.

Before Defining Bulk Host Name Formats

Before you specify a bulk host name format, ensure that it complies with the following rules:

- The NIOS appliance uses *<prefix>xx-xx-xx-xx* for bulk hosts. Ensure that the bulk host name does not conflict with CNAMEs, DNAMEs, or host name aliases.
- When you add a bulk host, if you enable the **Automatically add reverse mapping** option and there is a CNAME record in the corresponding reverse zone that conflicts with a PTR record generated by the bulk host, the bulk host insertion fails and an error message appears. For example, if there is a CNAME with the alias **15** in a reverse zone **1.168.192.in-addr.arpa** and if you add a bulk host **foo/192.168.1.10/192.168.1.20** with the **Automatically add reverse mapping** option selected, the insertion fails and an error message appears because both the bulk host and the CNAME generate a record **15.1.168.192.in-addr.arpa** in the reverse zone.

- You cannot create or change a bulk host if a zone is locked by another user. If you select a different template for the Grid, it changes each record associated with the bulk host.
- You can define bulk host name formats only at the Grid level and override them at the bulk host level; not at the zone or bulk host object level.
- When you upgrade to NIOS 4.3r3 or earlier releases, the system migrates existing bulk hosts as follows:
 - If you did not customize the bulk host IP format, there is no action required. All migrated bulk hosts continue to use the Grid-level default four-octet format **-\$1-\$2-\$3-\$4**. See [Specifying Bulk Host Name Formats](#) on page 656.
 - If you customized the bulk host IP format, the system creates a new template called *Migrated Default* template. All migrated bulk hosts override the Grid default template and use the *Migrated Default* template.

Note: The NIOS appliance considers two bulk hosts that have the same prefix, start address, and end address as duplicate hosts; even if they use different bulk host formats.

Bulk Host Name Format Rules

[Table 19.1](#) describes the rules that you should follow when you create bulk host name formats. It also provides examples of valid and invalid formats for each rule.

Table 19.1 Bulk Host Name Format Rules and Examples

Rule	Example
The suffix format cannot have more than four octets.	-\$4-\$5 is invalid.
The octets must be in order.	-\$2-\$3-\$4 is valid but -\$3-\$2-\$4 is invalid.
Do not skip octets.	-\$2-\$3-\$4 is valid but -\$2-\$4 is invalid.
Do not use a combination of both the \$ and # symbols together as octet references; use only one of them.	-\$2-#3-\$4 is invalid.
The suffix format must contain at least the fourth octet. You must define at least one -\$4 or -#4 .	-\$4 is valid but -\$3 is invalid.
If the suffix format uses \$ references, it cannot be preceded by a digit. You must add a non-digit prefix to each \$ or # reference.	-\$2-\$3-\$4
The \ character is the designated escape character for the \$, # and \ characters. You cannot use the \$ or # symbols as separators unless you prefix them with an escape character \.	For the IP address 10.19.32.133 , the format \#-#1-#2-#3-#4 expands to #-010-019-032-133 .
The bulk host name format must comply with its zone host name policy.	You cannot insert a bulk host name format -\$4 in a zone that uses Allow Underscore as host name policy because the policy does not allow you to use the ? character in the host name.

Table 19.1 Bulk Host Name Format Rules and Examples

Rule	Example
The bulk host name must comply with the maximum label length.	The sum of the bulk host name prefix and suffix cannot be greater than 63 characters. When you enter a suffix format, the NIOS appliance determines the length of the longest bulk host defined, and checks that the sum of the bulk host prefix and suffix length does not exceed 63 characters; if it does, an error message appears.
The bulk host name cannot result in an FQDN with more than 255 characters.	
The NIOS appliance computes the maximum length of the bulk host suffix by expanding the bulk host IP format using 255.255.255.255.	For the format string -\$1-\$2-\$3-\$4 , the maximum length of the suffix is -255-255-255-255 ; that is, 16 characters. Therefore, the maximum length of the host prefix is 47 characters.
The bulk host name must not be the same as a CNAME/DNAME.	If there is a CNAME record with alias foo-003-015 , you cannot insert a bulk host foo/1.2.3.10/1.2.3.20 using template -#3-#4 because foo-003-015 is also one of the synthetic host names in the bulk host.
Each host name in the bulk host must be unique.	You cannot insert a bulk host foo/1.2.3.10/1.2.4.20 using the template -\$4 because the system resolves the host name foo-10 to both 1.2.3.10 and 1.2.4.10 . To ensure that the bulk host name is unique, use the template -\$3-\$4 .
You cannot insert a bulk host that violates the uniqueness of two bulk hosts that have the same prefix and use the same name format.	If there is a bulk host foo/1.2.3.10/1.2.4.20 using the template -\$3-\$4 , you cannot insert another bulk host foo/1.3.4.10/1.3.5.20 using the same template because the system resolves host name foo-4-15 to both 1.2.4.15 and 1.3.4.15 . Instead, use the template -\$2-\$3-\$4 to ensure that the two bulk hosts are unique.

The appliance provides four predefined formats. You can define additional formats or change the default format at the Grid level only. To define new bulk host name formats:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. Select the **Host Naming** tab of the *Grid DNS Properties* editor.
The Bulk Host Name Formats table displays four predefined name suffix formats. The following examples show the host name that each format generates for the zone test.com:
Four Octets: **\$1-\$2-\$3-\$4** (Default)—Generates foo-192-168-1-15.test.com.
Three Octets: **-\$2-\$3-\$4**—Generates foo-168-1-15.test.com
Two Octets: **-\$3-\$4**—Generates foo-1-15.test.com
One Octet: **-\$4**—Generates foo-15.test.com
For the IP address 10.100.0.10, the format **-\$1-\$2-\$3-\$4** generates the host name suffix **-10-100-0-10**. The format **#1-#2-#3-#4** generates the host name suffix **-010-100-000-010**.
3. Click **Add** to enter the name and format of a new bulk host name format.
4. Optionally, click the Default column of a format and select **Default** to make it the Grid default.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Adding Bulk Hosts

To add a bulk host:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Host -> Add Bulk Host**.
2. In the *Add Bulk Host* wizard, complete the following fields:
 - **Prefix:** If Grid Manager displays a zone name, enter a prefix (or series of characters) to insert at the beginning of each host name. The displayed zone name can either be the last selected zone or the zone from which you are adding the bulk host record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and enter a prefix for the bulk host record. You can enter any printable character that complies with the zone host name policy or you can also leave this blank.
 The sum of the bulk host prefix length and suffix length must not exceed 63 characters. When you enter a prefix, the NIOS appliance computes the maximum length of the bulk host suffix to verify that the total prefix and suffix length does not exceed 63 characters. If it does, the appliance displays an error message indicating the number of characters that you must remove to make a valid prefix.
 - **DNS View:** Displays the DNS view of the zone to which the bulk host records belong.
 - **Host Name Policy:** Displays the host name policy of the selected DNS zone.
 - **Name Format:** To override the default four-octet suffix format or the format set at the Grid level, and specify a different format, click **Override** and select a host name format from the **Name Formats** drop-down menu. The *Name Formats* drop-down menu lists the formats **Four Octets**, **Three Octets**, **Two Octets**, and **One Octet** along with any other bulk host name formats that you have defined.
 - **Starting IP Address:** Enter the first IP address in the range of addresses for the group.
 - **End IP Address:** Enter the last IP address in the range of addresses for the group.
 - **Comment:** Optionally, enter additional information for this record.
 - **Automatically Add Reverse Mapping:** Click to have the appliance automatically create a PTR record for each IP address within the bulk host range.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes for the bulk host record. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration and click **Restart** if it appears at the top of the screen

To modify or delete a bulk host, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Example 1 - Responding to DNS AXFR Queries

This example shows the responses the bulk host **foo/1.2.3.10/1.2.3.20** returns to DNS AXFR (Full Zone Transfers) queries.

If the bulk host uses the template **-\$3-\$4**, the query returns:

```
foo-3-10.test.com
foo-3-11.test.com
.....
foo-3-20.test.com
```

If the bulk host uses the template **-#2-#3-#4**, the query returns:

```
foo-002-003-010.test.com
foo-002-003-011.test.com
.....
foo-002-003-020.test.com
```


Example 2 - Importing Zones with Bulk Hosts

When you import zones with bulk hosts, the system selects the most specific match.

The following example can possibly match three octet, two octet, and one octet formats; however, the system selects the most specific four octet default format.

The query:

```
foo-1-2-3-4 IN A 1.2.3.4
```

```
foo-1-2-3-5 IN A 1.2.3.5
```

Results in the match:

```
foo/1.2.3.4/1.2.3.5(Four Octets)
```

Not in any of the following:

```
foo-1/1.2.3.4/1.2.3.5(Three Octets)
```

```
foo-1-2/1.2.3.4/1.2.3.5(Two Octets)
```

```
foo-1-2-3/1.2.3.4/1.2.3.5(One Octet)
```

MANAGING RESOURCE RECORDS

DNS resource records provide information about objects and hosts. DNS servers use these records to respond to queries for hosts and objects. The appliance supports IDNs for all DNS resource records. For information about IDNs, see [Support for Internationalized Domain Names](#) on page 93. Note that the appliance does not decode the IDN of a resource record to punycode. In other words, a record that contains a domain name in punycode is displayed in punycode and a record that contains an IDN is displayed in its native characters.

You can manage the following types of DNS resource records:

- A (IPv4 Address)—For information, see [Managing A Records](#) on page 660.
- NS (Name server)—For information, see [Managing NS Records](#) on page 662.
- AAAA (IPv6 Address)—For information, see [Managing AAAA Records](#) on page 662.
- PTR (Pointer)—For information, see [Managing PTR Records](#) on page 664.
- MX (Mail exchanger)—For information, see [Managing MX Records](#) on page 665.
- SRV (Service location)—For information, see [Managing SRV Records](#) on page 666.
- TXT (Text)—For information, see [Managing TXT Records](#) on page 668.
- CNAME (Canonical name)—For information, see [Managing CNAME Records](#) on page 669.
- DNAME—For information, see [Managing DNAME Records](#) on page 671.

Managing A Records

An A (address) record is a DNS resource record that maps a domain name to an IPv4 address. To define a specific name-to-address mapping, you can add an A record to a previously defined authoritative forward-mapping zone. If the zone is associated with one or more networks, the IP address must belong to one of the associated networks. For example, if the A record is in the corp100.com zone, which is associated with 10.1.0.0/16 network, then the IP addresses of the A record must belong to the 10.1.0.0/16 network. For information about associating zones and networks, see [Associating Networks with Zones](#) on page 813.

The appliance also supports wildcard A records. For example, you can use a wildcard A record in the corp100.com domain to map queries for names such as www1.corp100.com, ftp.corp100.com, main.corp100.com, and so on to the IP address of a public-facing web server. Note that wildcard names only apply when the domain name being queried does not match any resource record.

Note: If an A record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** field displays the record name in UTF-8 encoded format. For example, an A record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Name** field.

Adding A Records

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add A Record**.
2. In the *Add A Record* wizard, do the following:
 - **Name:** If Grid Manager displays a zone name, enter the hostname that you want to map to an IP address. The displayed zone name can either be the last selected zone or the zone from which you are adding the host record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box and then enter the hostname. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is corp100.com and you enter admin, then the FQDN becomes admin.corp100.com. Ensure that the domain name you enter complies with the hostname restriction policy defined for the zone. To create a wildcard A record, enter an asterisk (*) in this field.
 - **DNS View:** This field displays the DNS view to which the DNS zone belongs.
 - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
 - **Hostname Policy:** Displays the hostname policy of the zone.
 - In the **IP Addresses** section, click the Add icon and do one of the following:
 - Select **Add Address** to enter the IPv4 address to which you want the domain name to map.
 - or
 - Select **Next Available IPv4** to retrieve the next available IP address in a network. Note that the appliance displays an error message if the obtained next available IP address is already being used by other users. You can request for another unused IP address or enter a new one.

If the A record is in zone that has associated networks, the *Network Selector* dialog box lists the associated networks. If the zone has no network associations, the *Network Selector* dialog box lists the available networks. When you select a network, Grid Manager retrieves the next available IP address in that network.
 - **Comment:** Optionally, enter additional information about the A record.
 - **Create associated PTR record:** Select this option to automatically generate a PTR record that maps the specified IP address to the hostname. To create the PTR record, the reverse-mapping zone must be in the database.
 - **Disable:** Select this check box to disable the record. Clear the check box to enable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying A Records

When you modify an A record, you can do the following:

- In the **General** tab, you can change the information you previously entered through the wizard, as described in [Adding A Records](#) on page 661.
- The **Discovered Data** tab displays discovered data, if any, for the record. For information, see [Viewing Discovered Data](#) on page 510.

You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing NS Records

An NS record identifies an authoritative DNS server for a domain. Each authoritative DNS server must have an NS record. Grid Manager automatically creates an NS record when you assign a Grid member as the primary server for a zone. You can manually create NS records for other zones. NS records associated with one or more IP addresses are used for related A record and PTR record generation. You can configure an NS record for anycast IP addresses on the appliance. For more information about anycast, see [About Anycast Addressing for DNS](#) on page 761.

Adding NS Records

To add an NS record:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add NS Record**.
2. In the *Add NS Record* wizard, complete the following fields:
 - **Zone:** The displayed zone name can either be the last selected zone or the zone from which you are adding the NS record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box.
 - **DNS View:** Displays the DNS view to which the selected zone belongs.
 - **Hostname Policy:** Displays the hostname policy of the selected zone.
 - **Name Server:** Enter the host name that you want to configure as the name server for the zone. IDN is not supported in this field. You can use the punycode representation of an IDN in this field.
3. Click **Next** to enter IP addresses for the name server.
4. In the *Name Server Addresses* panel, click the Add icon and complete the following fields:
 - **Address:** Enter the IP address of the name server.
 - **Add PTR Record:** This field displays **Yes** by default, enabling the automatic generation of a PTR record for the IP address. You can select **No** to disable the generation of the PTR record.
5. Click **Next** to define extensible attributes, or save the configuration and click **Restart** if it appears at the top of the screen.

Modifying and Deleting NS Records

When you modify an NS record, you can change the following information:

- In the **General** tab, you can change the name server name.
- In the **Addresses** tab, you can do the following:
 - Delete an address by selecting it and clicking the Delete icon.
 - Add an address by clicking the Add icon, and then entering the IP address and completing the **Add PTR Record** field.

Managing AAAA Records

An AAAA (quad A address) record maps a domain name to an IPv6 address. To define a specific name-to-address mapping, add an AAAA record to a previously defined authoritative forward-mapping zone. If the zone is associated with one or more networks, the IP address must belong to one of the associated networks. For example, if the AAAA record is in the corp100.com zone, which is associated with the 1111:0001/32 network, then the IP addresses of the A record must belong to that network. For information about associating zones and networks, see [Associating Networks with Zones](#) on page 813.

Note: If an AAAA record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** field displays the record name in UTF-8 encoded format. For example, an AAAA record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Name** field.

Adding AAAA Records

To create an AAAA record:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add AAAA Record**.
2. In the *Add AAAA Record* wizard, complete the following:
 - **Name:** If Grid Manager displays a zone name, enter the hostname that you want to map to an IP address. The displayed zone name can either be the last selected zone or the zone from which you are adding the AAAA record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the hostname. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is corp100.com and you enter admin, then the FQDN becomes admin.corp100.com.
 - **DNS View:** Displays the DNS view to which the selected DNS zone belongs.
 - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
 - **Hostname Policy:** Displays the hostname policy of the zone.
 - **IP Address:** Enter the IPv6 address to which you want the domain name to map. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered.
 - **Comment:** Optionally, enter additional information about this record.
 - **Create associated PTR record:** Select this option to automatically generate a PTR record that maps the specified IP address to the hostname. To create the PTR record, the reverse-mapping zone must be in the database.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying AAAA Records

When you modify an AAAA record, you can do the following:

- In the **General** tab, you can change the information you previously entered through the wizard.
- In the **Discovered Data** tab, you can view discovered data, if any, for the record. For information, see [Viewing Discovered Data](#) on page 510.

You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing PTR Records

In a reverse-mapping zone, a PTR (pointer) record maps an IP address to a hostname. Before adding a PTR record to a reverse-mapping zone, you must first create the zone. You can also add PTR records to forward-mapping zones to support zeroconf (zero configuration networking), such as wide-area Bonjour. For information about the Bonjour protocol, refer to <http://www.apple.com/support/bonjour>. Though adding PTR records to forward-mapping zones supports some of the use cases in RFC 1101, it does not support the network name mapping use case described in the RFC. For more information, refer to <http://tools.ietf.org/html/rfc1101>.

Note: If a PTR record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** and **Domain Name** fields display the record name in UTF-8 encoded format. For example, a PTR record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Name** and **Domain Name** fields.

Adding PTR Records

To add a PTR record:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add PTR Record**.
2. In the *Add PTR Record* wizard, do the following:
 - **Name or IP Address:** From the drop-down list, select **Name** or **IP Address**. When you select **Name**, click **Select Zone** to select a zone, and then enter the host name for the PTR record. Note that when you launch this wizard from the **IPAM** tab, you can only select a reverse-mapping zone. When you launch this from a reverse-mapping zone, the IP address field is populated with the prefix that corresponds to the selected zone. When you launch this from a forward-mapping zone, you can only specify the host name, not an IP address.
When you select **IP Address**, enter the IPv4 or IPv6 address that you want to map to the domain name.
 - **DNS View:** If you entered an IP address, you must select the DNS view of the PTR record. If you entered a name, this field displays the DNS view of the selected zone.
 - **Domain Name:** Enter the domain name to which you want the PTR record to point. For example, you can enter corp100.com.
 - **Comment:** Optionally, enter information about the PTR record.
 - **Disable:** Select this check box to disable the record. Clear the check box to enable it.
3. Save the configuration, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Click **Restart** if it appears at the top of the screen.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

Note: When you add a PTR record to a forward-mapping zone, a message may appear on the top of the wizard if a Grid member is configured to ignore DNS queries for PTR records in forward-mapping zones. Contact Infoblox Technical Support for more information about this message.

Modifying PTR Records

Do the following to modify a PTR record:

- In the **General** tab, you can change the information you previously entered through the wizard. Note that you cannot change an IPv4 address to an IPv6 address or move a PTR record from a forward-mapping zone to a reverse-mapping zone and vice versa. When you modify a PTR record that belongs to a forward-mapping zone, you can only modify the name since there is no IP address for such record. For information, see [Adding PTR Records](#) on page 664.
- In the **Discovered Data** tab, you can view discovered data, if any, for the record. For information, see [Viewing Discovered Data](#) on page 510.

You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing MX Records

An MX (mail exchanger) record maps a domain name to a mail exchanger. A mail exchanger is a server that either delivers or forwards mail. You can specify one or more mail exchangers for a zone, as well as the preference for using each mail exchanger. A standard MX record applies to a particular domain or subdomain.

You can use a wildcard MX record to forward mail to one mail exchanger. For example, you can use a wildcard MX record in the corp100.com domain to forward mail for eng.corp100.com and sales.corp100.com to the same mail exchange, as long as the domain names do not have any matching resource record. Wildcards only apply when the domain name being queried does not match any record.

Note: If an MX record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Mail Destination** and **Mail Exchanger** fields display the record name in UTF-8 encoded format. For example, an MX record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Mail Destination** and **Mail Exchanger** fields.

See [Figure 19.1](#).

Figure 19.1 MX Records

The following MX records direct queries for one or more domains to the same mail exchanger:

An MX record for the mail exchanger that answers queries for just the corp100.com domain (and its corresponding A record):

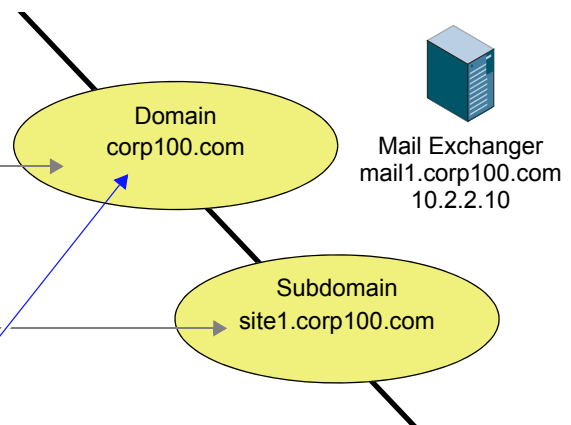
```
corp100.com IN MX 0 mail1.corp100.com
mail1.corp100.com IN A 10.2.2.10
```

An MX record for just site1.corp100.com, a subdomain of corp100.com:

```
site1.corp100.com IN MX 0
mail1.corp100.com
```

A wildcard MX record for the corp100.com domain:

```
*.corp100.com IN MX 0 mail1.corp100.com
```



Note: You must also create an A record for the host defined as a mail exchanger in an MX record.

Adding MX Records

To add an MX record from the Tasks Dashboard, see [Add MX Record](#) on page 106. You can also add MX records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add MX Record** from the Toolbar.

Modifying and Deleting MX Records

When you modify an MX record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing SRV Records

An SRV (service location) record directs queries to hosts that provide specific services. For example, if you have an FTP server, then you might create an SRV record that specifies the host which provides the service. You can specify more than one SRV record for a host. For more information about SRV records, see *RFC 2052, A DNS RR for specifying the location of services (DNS SRV)*.

Note: If an SRV record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** and **SRV Target** fields display the domain name in UTF-8 encoded format. For example, an SRV record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **Name** and **SRV Target** fields.

Adding SRV Records

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add** -> **Record** -> **Add SRV Record**.
2. In the *Add SRV Record* wizard, complete the following fields:
 - **Display input as:** Select the format in which you want the SRV record to be displayed. When you select **RFC 2782 format**, the appliance follows the `_service._protocol.name` format as defined in RFC 2782. When you select **Free format**, enter the entire name in the **Domain** field.
 - **Service:** Specify the service that the host provides. You can either select a service from the list or type in a service, if it is not on the list. For example, if you are creating a record for a host that provides FTP service, select **_ftp**. To distinguish the service name labels from the domain name, the service name is prefixed with an underscore. If the name of the service is defined at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>, use that name. Otherwise, you can use a locally-defined name.
 - **Protocol:** Specify the protocol that the host uses. You can either select a protocol from the list or type in a protocol, if it is not on the list. For example, if it uses TCP, select **_tcp**. To distinguish the protocol name labels from the domain name, the protocol name is prefixed with an underscore.
 - **Domain:** If Grid Manager displays a zone name, enter the name here to define an SRV record for a host or subdomain. The displayed zone name can either be the last selected zone or the zone from which you are adding the SRV record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name to define the SRV record. The NIOS appliance prefixes the name you enter to the domain name of the selected zone. For example, if you want to create an SRV record for a web server whose host name is www2.corp100.com and you define the SRV record in the corp100.com zone, enter www2 in this field. To define an SRV record for a domain whose name matches the selected zone, leave this field blank. The NIOS appliance automatically adds the domain name (the same as the zone name) to the SRV record. For example, if you want to create an SRV record for the corp100.com domain and you selected the corp100.com zone, leave this field blank.
 - **Preview:** After you have entered all the information, this field displays the FQDN, which is the concatenation of the **Service**, **Protocol**, and **Domain** fields.

- **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
- **Priority:** Select or enter an integer from 0 to 65535. The priority determines the order in which a client attempts to contact the target host; the domain name host with the lowest number has the highest priority and is queried first. Target hosts with the same priority are attempted in the order defined in the **Weight** field.
- **Weight:** Select or enter an integer from 0 to 65535. The weight allows you to distribute the load between target hosts. The higher the number, the more that host handles the load (compared to other target hosts). Larger weights give a target host a proportionately higher probability of being selected.
- **Port:** Specify the appropriate port number for the service running on the target host. You can use standard or nonstandard port numbers, depending on the requirements of your network. You can select a port number from the list or enter an integer from 0 to 65535.
- **Target:** Enter the canonical domain name of the host (not an alias); for example, www2.corp100.com

Note: In addition, you need to define an A record mapping the canonical name of the host to its IP address.

- **Comment:** Enter a descriptive comment for the record.
- **Disable:** Clear the check box to enable the record. Select the check box to disable it.

3. Save the configuration, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Click **Restart** if it appears at the top of the screen.

Modifying and Deleting SRV Records

Do the following to modify an SRV record:

- In the **General** tab, the **Display input as** field displays the format in which the SRV record was configured. For **RFC 2782 format**, the appliance matches the `_service._protocol.name` format and displays the corresponding information in the **Service** and **Protocol** fields. If the appliance cannot match the service and protocol, it displays the entire name in the **Domain** field. For **Free format**, the entire name is displayed in the **Domain** field. For more information about the other fields, see [Adding SRV Records](#) on page 666.

Note: The appliance does not match the service and protocol names to exactly how they appear in the drop-down lists. It only checks whether the first two parts of the names start with an underscore. If the first two parts do not start with an underscore, the appliance assumes it is a free format. For example, `_abc._xyz.name` is considered as RFC 2782 format even though `_abc` is not in the **Service** drop-down list, and `_xyz` is not in the **Protocol** drop-down list. Grid Manger displays `_abc` in the **Service** field and `_xyz` in the **Protocol** field. On the other hand, `"abc.xyz.name"` is considered as a free format because the first two parts do not start with underscores, and Grid Manager displays this in its entirety in the **Domain** field.

You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing TXT Records

A TXT (text record) record contains supplemental information for a host. For example, if you have a sales server that serves only North America, you can create a text record stating this fact. You can create more than one text record for a domain name.

Note: If a TXT record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** field displays the domain name in UTF-8 encoded format. For example, a TXT record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **Name** field.

Using TXT Records for SPF

SPF (Sender Policy Framework) is an anti-forgery mechanism designed to identify spam e-mail. SPF fights e-mail address forgery and makes it easier to identify spam, worms, and viruses. Domain owners identify sending mail servers in DNS. SMTP receivers verify the envelope sender address against this information, and can distinguish legitimate mail from spam before any message data is transmitted.

SPF makes it easy for a domain to say, “I only send mail from these machines. If any other machine claims that I'm sending mail from there, they're not valid.” For example, when an AOL user sends mail to you, an email server that belongs to AOL connects to an email server that belongs to you. AOL uses SPF to publish the addresses of its email servers. When the message comes in, your email servers can tell if the server that sent the email belongs to AOL or not.

You can use TXT records to store SPF data that identifies what machines send mail from a domain. You can think of these specialized TXT records as *reverse MX records* that e-mail servers can use to verify if a machine is a legitimate sender of an e-mail.

SPF Record Examples

```
corp100.com. IN TXT "v=spf1 mx -all"
corp100.net. IN TXT "v=spf1 a:mail.corp100.com -all"
corp100.net. IN TXT "v=spf1 include:corp100.com -all"
corp100.net. IN TXT "v=spf1 mx -all exp=getlost.corp100.com"
corp100.com. IN TXT "v=spf1 include:corp200.com -all"
```

Adding TXT Records

To add an TXT record from the Tasks Dashboard, see [Add TXT Record](#) on page 106. You can also add TXT records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add TXT Record** from the Toolbar.

Modifying and Deleting TXT Records

When you modify a TXT record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing CNAME Records

A CNAME record maps an alias to a canonical name. You can use CNAME records in both forward- and IPv4 reverse-mapping zones to serve two different purposes. (At this time, you cannot use CNAME records with IPv6 reverse-mapping zones.)

Note: If a CNAME record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Alias** and **Canonical Name** fields display the domain name in UTF-8 encoded format. For example, a CNAME record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **Canonical Name** and **Alias** fields.

CNAME Records in Forward-Mapping Zones

In a forward-mapping zone, a CNAME record maps an alias to a canonical (or official) name. CNAME records are often more convenient to use than canonical names because they can be shorter or more descriptive. For example, you can add a CNAME record that maps the alias *qa.engr* to the canonical name *qa.engr.corp100.com*.

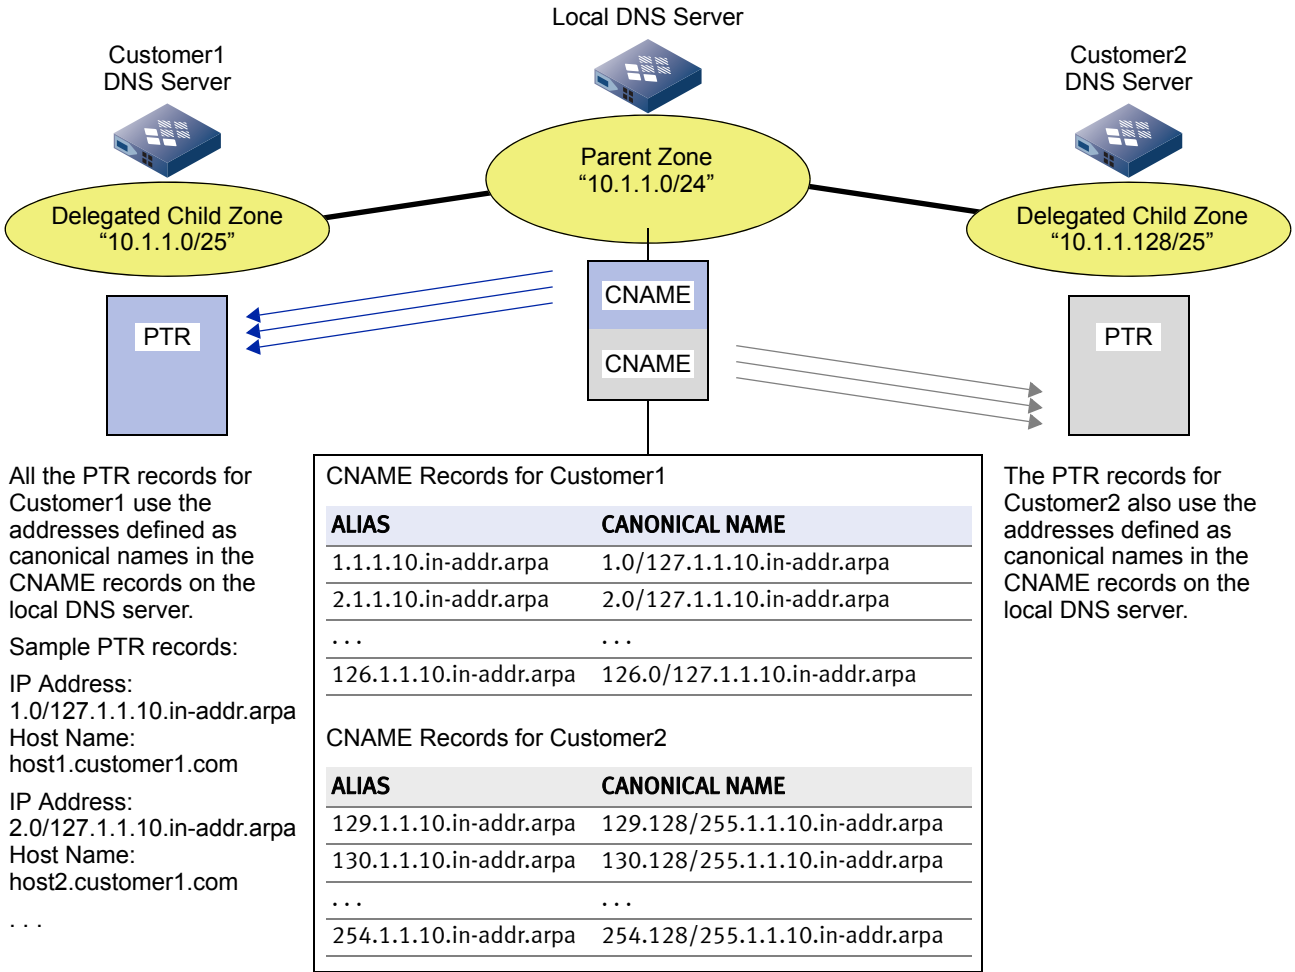
Note: A CNAME record does not have to be in the same zone as the canonical name to which it maps. In addition, a CNAME record cannot have the same name as any other record in that zone.

To add a CNAME record to a forward-mapping zone from the Tasks Dashboard, see [Add CNAME Record](#) on page 105. You can also add CNAME records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add CNAME Record** from the Toolbar.

CNAME Records in IPv4 Reverse-Mapping Zones

You can add CNAME records to an IPv4 reverse-mapping zone to create aliases to addresses maintained by a different name server when the reverse-mapping zone on the server is a delegated child zone with fewer than 256 addresses. This technique allows you to delegate responsibility for a reverse-mapping zone with an address space of fewer than 256 addresses to another authoritative name server. See [Figure 19.2](#) and *RFC 2317, Classless IN-ADDR.ARPA delegation*.

Figure 19.2 CNAME Records in a Reverse-Mapping Zone



You add CNAME records in the parent zone on your name server. The aliases defined in those CNAME records point to the addresses in PTR records in the child zone delegated to the other server.

When you define a reverse-mapping zone that has a netmask from /25 (255.255.255.128) to /31 (255.255.255.254), you must include an RFC 2317 prefix. This prefix can be anything, from the address range (examples: 0-127, 0/127) to descriptions (examples: first-network, customer1). On a NIOS appliance, creating such a reverse-mapping zone automatically generates all the necessary CNAME records. However, if you need to add them manually to a parent zone that has a child zone with fewer than 255 addresses.

Adding CNAME Records

To add a CNAME record to a forward-mapping or reverse-mapping zone from the Tasks Dashboard, see [Add CNAME Record](#) on page 105. You can also add CNAME records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Add CNAME Record** from the Toolbar.

Modifying and Deleting CNAME Records

When you modify a CNAME record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing DNAME Records

A DNAME record maps all the names in one domain to those in another domain, essentially substituting one domain name suffix with the other (see RFC 2672, *Non-Terminal DNS Name Redirection*). For example, adding a DNAME record to the corp100.com domain mapping “corp100.com” to “corp200.com” maps *name-x.corp100.com* to *name-x.corp200.com*:

Domain Name		Target Domain Name
server1.corp100.com	→	server1.corp200.com
server2.corp100.com	→	server2.corp200.com
server3.corp100.com	→	server3.corp200.com
...corp100.com	→	...corp200.com

Note: If a DNAME record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Alias** and **Target** fields display the domain name in UTF-8 encoded format. For example, a DNAME record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **Alias** and **Target** fields.

When a request arrives for a domain name to which a DNAME record applies, the NIOS appliance responds with a CNAME record that it dynamically creates based on the DNAME definition. For example, if there is a DNAME record

corp100.com. DNAME corp200.com.

and a request arrives for server1.corp100.com, the NIOS appliance responds with the following CNAME record:

server1.corp100.com. CNAME server1.corp200.com.

If responding to a name server running BIND 9.0.0 or later, the NIOS appliance also includes the DNAME record in its response, so that name server can also create its own CNAME records based on the cached DNAME definition.

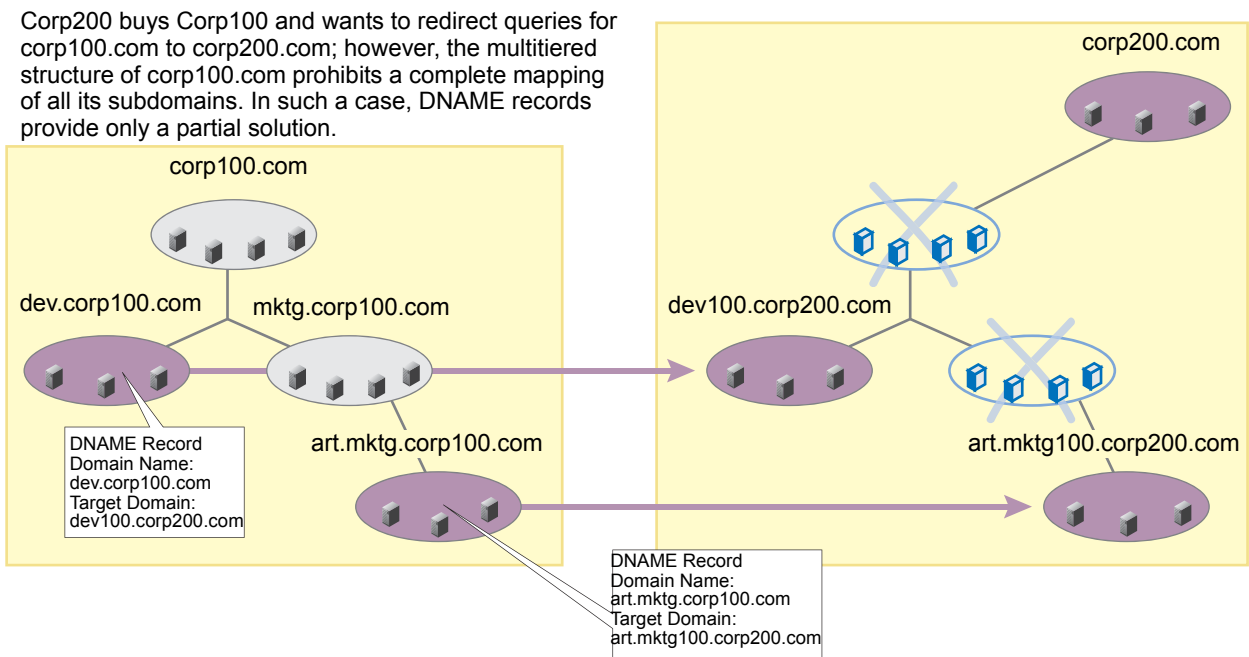
The following are two common scenarios for using DNAME records:

- One company buys another and wants people using both the old and new name spaces to reach the same hosts.
- A virtual Web hosting operation offers different “vanity” domain names that point to the same server or servers.

There are some restrictions that apply to the use of DNAME records:

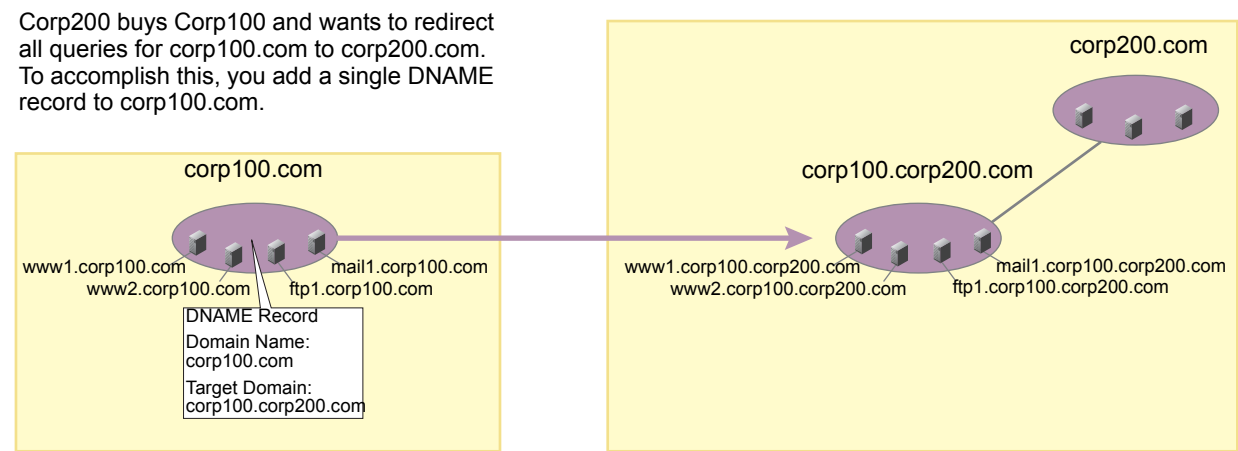
- You cannot have a CNAME record and a DNAME record for the same subdomain.
- You cannot use a DNAME record for a domain or subdomain that contains any subdomains. You can only map the lowest level subdomains (those that do not have any subdomains below them). For an example of using DNAME records in a multi-tiered domain structure, see [Figure 19.3](#) on page 672.

Figure 19.3 Adding DNAME Records for the Lowest Level Subdomains



In the case of a domain structure consisting of a single domain (no subdomains), adding a DNAME record redirects queries for every name in the domain to the target domain, as shown in [Figure 19.4](#).

Figure 19.4 Adding a DNAME Record for a Single Domain



When using a DNAME record, you must copy the resource records for the source domain to the zone containing the target domain, so that the DNS server providing service for the target domain can respond to the redirected queries. For the example in [Figure 19.4](#), copy these records:

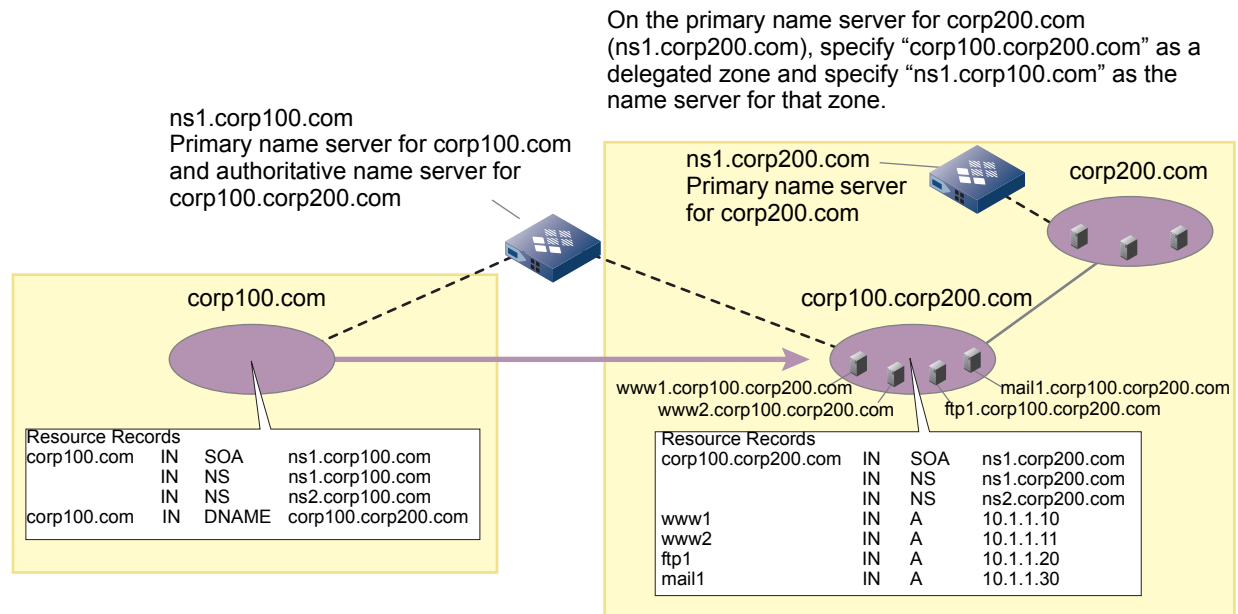
Copy from corp100.com	to corp100.corp200.com
www1 IN A 10.1.1.10	www1 IN A 10.1.1.10
www2 IN A 10.1.1.11	www2 IN A 10.1.1.11
ftp1 IN A 10.1.1.20	ftp1 IN A 10.1.1.20
mail1 IN A 10.1.1.30	mail1 IN A 10.1.1.30

After copying these records to the zone containing the corp100.corp200.com domain, delete them from the zone containing the corp100.com domain.

If DNS service for the source and target domain names is on different name servers, you can import the zone data from the NIOS appliance hosting the source domain to the appliance hosting the target domain. For information about this procedure, see [Importing Zone Data](#) on page 631.

If DNS service for the source and target domain names is on the same name server and the parent for the target domain is on a different server, you can delegate DNS services for the target domain name to the name server that provided—and continues to provide—DNS service for the source domain name (see [Figure 19.5](#) on page 673). By doing this, you can continue to maintain resource records on the same server, potentially simplifying the continuation of DNS administration.

Figure 19.5 Making the Target Zone a Delegated Zone



Note: This is a conceptual representation of domain name mapping and depicts the resulting hierarchical relationship of corp200.com as the parent zone for corp100.corp200.com. The hosts are not physically relocated.

The following tasks walk you through configuring the two appliances in [Figure 19.5](#) to redirect queries for corp100.com to corp100.corp200.com using a DNAME record:

On the ns1.corp100.com name server, do the following:

1. Create a new forward-mapping zone called corp100.corp200.com. See [Creating an Authoritative Forward-Mapping Zone](#) on page 617.
2. Copy all the resource records for the domain or subdomain to which the DNAME record is going to apply from corp100.com to corp100.corp200.com.

Note: Because you can only specify the records by type, not individually, you might have to copy some records that you do not want and then delete them from the corp100.corp200.com zone.

3. In the corp100.com zone, delete all the resource records for the domain or subdomain to which the DNAME record is going to apply.
4. Add a DNAME record to the corp100.com zone specifying “corp100.com” as the domain and “corp100.corp200.com” as the target domain. Adding a DNAME record is explained in the next section.
5. On the ns1.corp200.com name server, add corp100.corp200.com as a delegated zone and specify ns1.corp100.com as the name server for it. See [Configuring a Delegation](#) on page 638.

DNAME Records for Forward-Mapping Zones

To add a DNAME record to a forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add DNAME Record**.
2. In the *Add DNAME Record* wizard, complete the following fields:

Note: If you specify a subdomain in the Domain Name field when configuring a DNAME record and the subdomain is also a subzone, the DNAME record appears in the list view for the subzone, not in the list view for the parent zone selected in the process of adding the record.

- **Alias:** If Grid Manager displays a zone name, enter the name of a subdomain here. If you are adding a DNAME record for the entire zone, leave this field empty. This field is for adding a DNAME record for a subdomain within the selected zone. The displayed zone name can either be the last selected zone or the zone from which you are adding the CNAME record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name of a subdomain.
 - **Target:** Enter the domain name to which you want to map all the domain names specified in the **Alias** field.
 - **Comment:** Enter identifying text for this record, such as a meaningful note or reminder.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Save the configuration, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Click **Restart** if it appears at the top of the screen.

DNAME Records for Reverse-Mapping Zones

You can use DNAME records to redirect reverse lookups from one reverse-mapping zone to another. You can use DNAME records for reverse-mapping zones to simplify the management of subzones for classless address spaces larger than a class C subnet (a subnet with a 24-bit netmask).

RFC 2672, *Non-Terminal DNS Name Redirection*, includes an example showing the delegation of a subzone for an address space with a 22-bit netmask inside a zone for a larger space with a 16-bit netmask:

```
$ORIGIN 0.192.in-addr.arpa.
8/22      NS      ns.slash-22-holder.example.
8         DNAME    8.8/22
9         DNAME    9.8/22
10        DNAME    10.8/22
11        DNAME    11.8/22
```

The reverse-mapping zone 0.192.in-addr.arpa. applies to the address space 192.0.0.0/16. Within this zone is a subzone and subdomain with the abbreviated name 8/22. (Its full name is 8/22.0.192.in-addr.arpa.) This subdomain contains its own subdomains corresponding to the 1024 addresses in the 192.0.8.0/22 subnet:

- Subdomain 8/22 (8/22.0.192.in-addr.arpa)
 - Subdomain 8.8/22 for addresses 192.0.8.0 – 192.0.8.255 (or 192.0.8.0/24)
 - Subdomain 9.8/22 for addresses 192.0.9.0 – 192.0.9.255 (or 192.0.9.0/24)
 - Subdomain 10.8/22 for addresses 192.0.10.0 – 192.0.10.255 (or 192.0.10.0/24)
 - Subdomain 11.8/22 for addresses 192.0.11.0 – 192.0.11.255 (or 192.0.11.0/24)

The NS record delegates authority for the reverse-mapping subzone 8/22 to the DNS server ns.slash-22-holder.example.

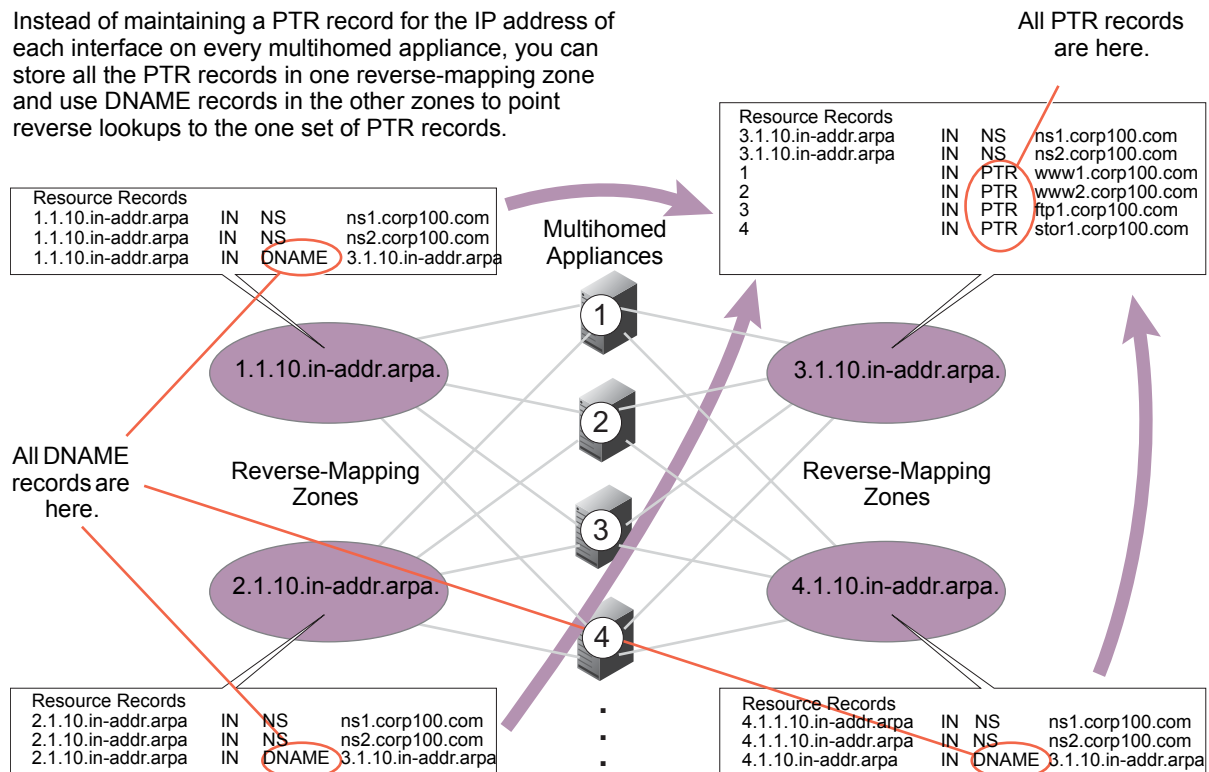
Finally, the DNAME records provide aliases mapping domain names that correspond to the 192.0.8.0/24, 192.0.9.0/24, 192.0.10.0/24, and 192.0.11.0/24 subnets to the respective subdomains 8.8/22, 9.8/22, 10.8/22, and 11.8/22 in the 8/22.0.192.in-addr.arpa subzone.

Note: NIOS appliances support DNAME records in reverse-mapping zones that map addresses to target zones with a classless address space larger than a class C subnet. However, NIOS appliances do not support such target zones.

You might also use DNAME records if you have a number of multihomed appliances whose IP addresses must be mapped to a single set of domain names. An example of this is shown in [Figure 19.6](#).

Figure 19.6 DNAME Records to Simplify DNS for Multihomed Appliances

Instead of maintaining a PTR record for the IP address of each interface on every multihomed appliance, you can store all the PTR records in one reverse-mapping zone and use DNAME records in the other zones to point reverse lookups to the one set of PTR records.



Note: If you specify a subdomain in the Domain Name field when configuring a DNAME record, and the subdomain is also a subzone, the DNAME record appears in the list view for the subzone, not in the list view for the parent zone that was selected when adding it.

To add a DNAME record to a reverse-mapping zone:

- From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add DNAME Record**.
- In the *Add DNAME Record* wizard, complete the following fields:

Note: If you specify a subdomain in the Domain Name field when configuring a DNAME record and the subdomain is also a subzone, the DNAME record appears in the list view for the subzone, not in the list view for the parent zone selected in the process of adding the record.

- Alias:** If Grid Manager displays a zone name, enter the name of a subdomain here. If you are adding a DNAME record for the entire zone, leave this field empty. This field is for adding a DNAME record for a subdomain within the selected zone. The displayed zone name can either be the last selected zone or the zone from which you are adding the CNAME record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name of a subdomain.

- **Target:** Type the name of the reverse-mapping zone to which you want to map all the addresses specified in the Domain Name field.
 - **Comments:** Enter identifying text for this record, such as a meaningful note or reminder.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Save the configuration, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Click **Restart** if it appears at the top of the screen.

Modifying and Deleting DNAME Records

When you modify a CNAME record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.

Managing NAPTR Records

A NAPTR (Name Authority Pointer) record specifies a rule that uses a substitution expression to rewrite a string into a domain name or URI (Uniform Resource Identifier). A URI is either a URL (Uniform Resource Locator) or URN (Uniform Resource Name) that identifies a resource on the Internet.

NAPTR records are usually used to map E.164 numbers to URIs or IP addresses. An E.164 number is a telephone number, 1-555-123- 4567 for example, in a format that begins with a country code, followed by a national destination code and a subscriber number. (E.164 is an international telephone numbering system recommended by the International Telecommunication Union.) Thus, NAPTR records allow us to use telephone numbers to reach devices, such as fax machines and VoIP phones, on the Internet.

To map an E.164 to a URI, the E.164 number must first be transformed into a domain name. ENUM (E.164 Number Mapping) specifies a method for converting E.164 numbers to domain names. For example, using the method specified by ENUM, the telephone number 1-555-123-4567 becomes the domain name 7.6.5.4.3.2.1.5.5.5.1.e164.arpa. For details about ENUM, refer to *RFC 3761, The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)*.

After the E.164 number is converted to a domain name, a DNS client can then perform a DNS lookup for the NAPTR records of the domain name. The following example illustrates how a DNS client processes NAPTR records.

In this example, the telephone number 1-555-123-4567 is converted to the domain name 7.6.5.4.3.2.1.5.5.5.1.e164.arpa. The DNS client then sends a query to the Infoblox DNS server for the NAPTR records associated with 7.6.5.4.3.2.1.5.5.5.1.e164.arpa. The Infoblox DNS server returns the following NAPTR record:

```
$ORIGIN 7.6.5.4.3.2.1.5.5.5.1.e164.arpa
IN NAPTR 10 100 "U" "sip + E2U" "!.^.*$!sip:jdoe@corp100.com!" .
```

Order
Preference

Flag

Service

Regular Expression

Replacement

The DNS client then examines the fields in the NAPTR record as follows:

- If a DNS client receives multiple NAPTR records for a domain name, the value in the Order field determines which record is processed first. It processes the record with the lowest value first.
- The DNS client uses the Preference value when the Order values are the same. Similar to the Preference field in MX records, this value indicates which NAPTR record the DNS client should process first when the records have the same Order value. It processes the record with the lowest value first.

In the example, the DNS client ignores the Order and Preference values because it received only one NAPTR record.

- The Flag field indicates whether the current lookup is terminal; that is, the current NAPTR record is the last NAPTR record for the lookup. It also provides information about the next step in the lookup process. The flags that are currently used are:

U: Indicates that the output maps to a URI (Uniform Record Identifier).

S: Indicates that the output is a domain name that has at least one SRV record. The DNS client must then send a query for the SRV record of the resulting domain name.

A: Indicates that the output is a domain name that has at least one A or AAAA record. The DNS client must then send a query for the A or AAAA record of the resulting domain name.

P: Indicates that the protocol specified in the Service field defines the next step or phase.

If the Flag field is blank, this indicates that the client must use the resulting domain name to look up other NAPTR records.

- The Service field specifies the service and protocol that are used to communicate with the host at the domain name. In the example, the service field specifies that SIP (Session Initiation Protocol) is used to contact the telephone service.
- The regular expression specifies the substitution expression that is applied to the original string of the client. In the example, the regular expression `!^.*$!sip:jdoe@corp100.com!` specifies that the domain name `7.6.5.4.3.2.1.5.5.1.e164.arpa` is replaced with `sip:jdoe@corp100.com`.
The regular expression in a NAPTR record is always applied to the original string of the client. It must not be applied to a domain name that resulted from a previous NAPTR rewrite.
- The Replacement field specifies the FQDN for the next lookup, if it was not specified in the regular expression.

Note: If a NAPTR record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Domain** and **Replacement** fields display the domain name in UTF-8 encoded format. For example, a NAPTR record with the domain name 电脑.test.com added through DDNS updates displays `\231\148\181\232\132\145.test.com` in the **Domain** and **Replacement** fields.

Adding a NAPTR Record

To add a NAPTR record:

- From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add NAPTR Record**.
- In the *Add NAPTR Record* wizard, complete the following fields:
 - Domain:** If Grid Manager displays a zone name, enter the domain name to which this resource record refers. The displayed zone name can either be the last selected zone or the zone from which you are adding the NAPTR record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter a domain name for the record. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the record. For example, if the zone name displayed is `corp100.com` and you enter `admin`, then the FQDN becomes `admin.corp100.com`.
 - DNS View:** Displays the DNS view of the selected zone.
 - Service:** Specifies the service and protocol used to reach the domain name that results from applying the regular expression or replacement. You can enter a service or select a service from the list.
 - Flags:** The flag indicates whether the resulting domain name is the endpoint URI or if it points to another record. Select one of the following:
 - U:** Indicates that the output maps to a URI.
 - S:** Indicates that the resulting domain name has at least one SRV record.
 - A:** Indicates that the resulting domain name has at least one A or AAAA record.
 - P:** Indicates that this record contains information specific to another application.

Leave this blank to indicate that the DNS client must use the resulting domain name to look up other NAPTR records. You can use the NAPTR records as a series of rules that are used to construct a URI or domain name.

- **Order:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. This value indicates the order in which the NAPTR records must be processed. The record with the lowest value is processed first.
 - **Preference:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. Similar to the Preference field in MX records, this value indicates which NAPTR record should be processed first when the records have the same Order value. The record with the lowest value is processed first.
 - **REGEX:** The regular expression that is used to rewrite the original string from the client into a domain name. RFC 2915 specifies the syntax of the regular expression. Note that the appliance validates the regular expression syntax between the first and second delimiter against the Python re module, which is not 100% compatible with POSIX Extended Regular Expression as specified in the RFC. For information about the Python re module, refer to <http://docs.python.org/release/2.5.1/lib/module-re.html>.
 - **Replacement:** This specifies the domain name for the next lookup. The default is a dot (.), which indicates that the regular expression in the **REGEX** field provides the replacement value. Alternatively, you can enter the replacement value in FQDN format.
 - **Comment:** Optionally, enter a descriptive comment for this record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Save the configuration, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Click **Restart** if it appears at the top of the screen.

Managing LBDN Records

When your Grid is integrated with GLBs (Global Load Balancers), you can add LBDNs (Load Balanced Domain Name) record to authoritative or delegated zones. The Infoblox name server does not actually serve these LBDNs, but they can be managed by the load balancer devices listed at the zone level. Note that you cannot add an LBDN if the zone is DNSSEC signed.

For information about how to add and modify LBDN records, see [Managing LBDN on GLBs](#) on page 1201.

Viewing Resource Records

You can view the configured resource records by navigating to the **Data Management** tab -> **DNS** tab -> **Zones** tab -> **zone** -> **Records** tab. Grid Manager displays the following information for each resource record in the zone:

- **Name:** The name of the record, if applicable. For host records, this field displays the canonical name of the host. For PTR record, this displays the PTR record name without the zone name.
- **Type:** The resource record type.
- **Data:** Data that the record contains. For host records, this field displays the IP address of the host. For PTR records, this displays the domain names.
- **Comment:** Comments that were entered for the resource record.
- **Site:** Values that were entered for this pre-defined attribute.

You can also display the following columns:

- **MS Delegation Addresses:** This column appears only if the primary server of the zone is a Microsoft server. It displays the IP addresses that are associated with an NS record.
- **TTL:** The TTL (time-to-live) value of the record.
- **Address:** The IPv4 or IPv6 address associated with the owner domain name in a reverse-mapping zone.
- **Shared:** Displays true for shared resource records. Otherwise, displays false.
- **Shared Record Group:** Displays the shared record group name of a shared record.
- **Disabled:** Indicates if the record is disabled.

You can do the following:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only.
- Edit the properties of a resource record.
 - Select the resource record, and then click the Edit icon.
- Delete a resource record.
 - Select the resource record, and then click the Delete icon.
- Export the list of resource records to a .csv file.
 - Click the Export icon.
- Print the list of resource records.
 - Click the Print icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
 1. In the filter section, click **Show Filter** and define filter criteria for the quick filter.
 2. Click **Save** and complete the configuration in the Save Quick Filter dialog box.

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].

Modifying, Disabling, and Deleting Host and Resource Records

You can modify, disable, or delete an existing host or DNS resource record. When physical repair or relocation of a network device occurs, you can disable a record instead of deleting it. When you disable a record, the NIOS appliance does not answer queries for it, nor does it include disabled records in zone transfers and zone imports. This avoids having to delete and then add the record again. When the changes to the physical device are complete, you can simply enable the host or resource record.

To modify or disable a host or resource record:

1. Use one of the following methods to retrieve the host or resource record:
 - Perform a global search.
 - Select it from a Smart Folder.
 - From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns_view* -> *zone* -> *host_record* or *resource_record*.
2. Select the record you want to modify and click the Modify icon.
3. In the host or resource record editor, you can do the following:
 - In the **General** tab, you can change most of the information, except for the read-only fields, such as the **DNS View** and **Host Name Policy**. You can select the **Disable** check box to disable the record.
 - In the **TTL** tab, you can modify the TTL setting. The NIOS appliance also allows you to specify TTL settings for each record. If you do not specify a TTL for a record, the appliance applies the default TTL value of the zone to each record. For information, see [About Time To Live Settings](#) on page 557.
 - In the **Extensible Attributes** tab, you can modify the attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - The **Permissions** tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#) on page 160.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

When you delete host and resource records, Grid Manager moves them to the Recycle Bin. You can use the Recycle Bin to store deleted DNS configuration objects and selectively restore objects to the active configuration at a later time. You can also permanently remove the objects from the Recycle Bin.

Note: You cannot delete automatically-generated records, such as NS records and SOA records.

To delete host and resource record:

1. Perform a global search to retrieve the record you want to delete.

or

From the **Data Management** tab, select the **DNS** tab, click the **Zones** tab-> *dns_view*-> *zone*-> *host_record* or *resource_record*.

2. Select the record and click the Delete icon.
3. When the confirmation dialog box displays, select **Yes**.

ABOUT SHARED RECORD GROUPS

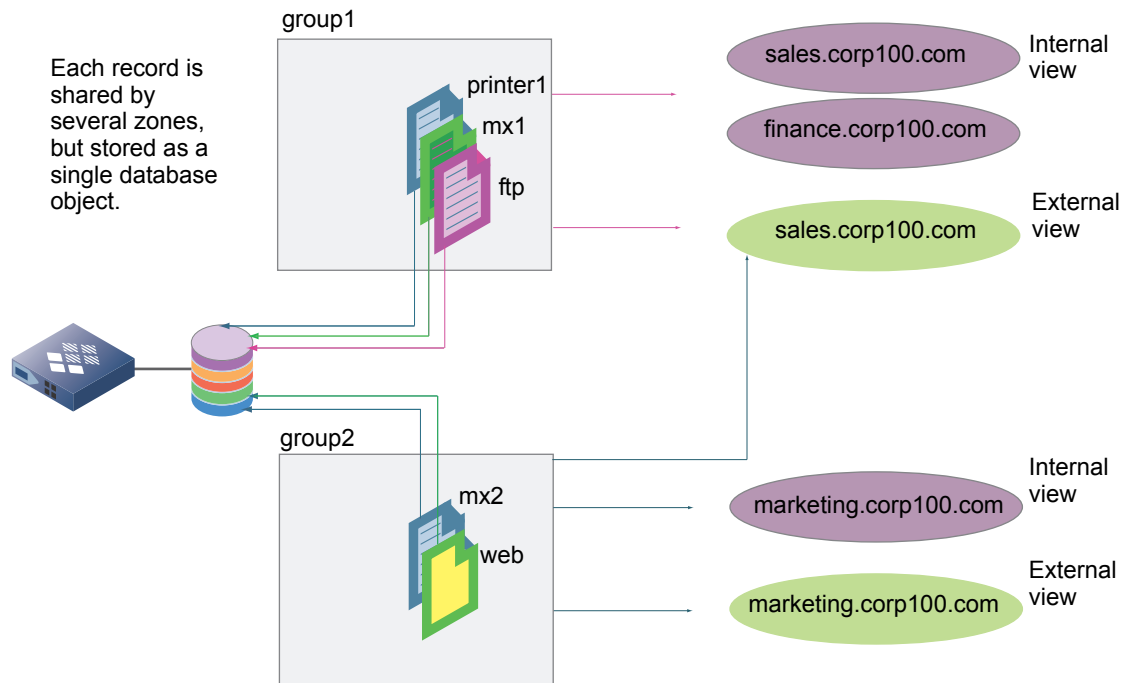
A shared record group is a set of resource records that you can add to multiple zones. You can create resource records in a group and share the group among multiple zones. The zones handle the shared resource records as any other resource record. You can include the following types of DNS resource records in a shared record group: A, SRV, MX, AAAA, and TXT.

Using shared record groups simplifies and expedites the administration of resource records. When you create or update a shared record, the appliance automatically updates it in all associated zones. In addition, shared resource records reduce the object count in the NIOS database; instead of creating the same record in multiple zones, you can use only one shared record. For example, for 10 zones and 500 records per zone, the object count decreases from 5278 objects to 781 objects.

[Figure 19.7](#) shows an example of how to create and use shared records.

In this example, there are two shared record groups. One group—group1— contains the A records ftp and printer1 and the MX record mx1, and the other group—group2—contains the A record web and the MX record mx2. The resource records in group1 are shared with the internal view zones sales.corp100.com and finance.corp100.com and the external view zone sales.corp100.com. The resource records in group2 are shared with the internal view zone marketing.corp100.com and the external view zones sales.corp100.com and marketing.corp100.com.

Figure 19.7 Creating Shared Records



Shared Records Guidelines

The following are guidelines for using shared records:

- You can include multiple shared A, AAAA, SRV, MX and TXT resource records in a group. You cannot include NS, CNAME, DNAME, PTR, host and bulk host records.
- You can add shared records to authoritative zones only. You cannot add shared records to forward zones, stub zones, or reverse mapping zones.
- Zones that contain shared records can also contain regular DNS records (not shared).
- When you change or delete a shared resource record, it changes the canonical source of the shared record and impacts all the zones that contain the record.
- You cannot copy shared records from a zone.
- You do not need to restart the appliance when you create, delete, or modify shared records.

Configuring Shared Record Groups

Before you can create shared resource records, you must first create the group to which they belong. The shared record group serves as a container for the shared resource records. The following are the tasks to configure a shared record group:

1. Create a shared record group and associate it with the appropriate zones. See [Creating a Shared Record Group](#) on page 682.
2. Create shared A, SRV, MX, AAAA, and TXT resource records, and add them into the shared record group. See [Managing Shared Resource Records](#) on page 683.

Creating a Shared Record Group

When you create a shared record group, the only requirement is that you give it a name. You can associate it with one or multiple zones when you first create the group or at a later time, by editing the shared record group. You can associate a shared record group with authoritative zones only. Associating the shared record group with a zone adds the shared records to the zone. The zone handles the shared records like any other resource records.

To create a shared record group:

1. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab, and then click the Add icon.
2. In the *Shared Record Group* wizard, specify the following:
 - **Name:** Enter the name of the shared record group. It can be up to 64 characters long and can contain any combination of printable characters. You can change the shared record group name even after you create the group. It does not impact the shared records in the group.
 - **Hostname Policy:** Click **Override** to supersede the hostname restriction policy set at the zone level or click **Inherit** to use the zone policy. This sets the hostname policy for the shared records in the group. See [Specifying Hostname Policies](#) on page 592.
 - **Comment:** Optionally, enter additional information about the shared record group.
3. Click **Next** to associate the shared record group with at least one zone.
4. Click the Add icon in the Associated Zones panel.
5. In the *Zone Selector* dialog box, select a zone by clicking the zone name. You can add multiple zones.
6. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
7. Save the configuration.

Viewing Shared Record Groups

You can view the configured shared record groups by navigating to the **Data Management** tab -> **DNS** tab -> **Shared Record Groups** tab. Grid Manager displays the following information about each shared record group:

- **Name:** The shared record group name.
- **Comment:** Comments that were entered for the shared record group.
- **Site:** Values that were entered for this pre-defined attribute.

You can do the following:

- List the shared resource records and associated zones in a shared record group.
 - Click a shared record group name.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Edit the properties of a shared record group.
 - Click the check box beside a shared record group, and then click the Edit icon.
- Delete a shared record group.
 - Click the check box beside a shared record group, and then click the Delete icon. Note that you must remove the zone associations in a shared record group before you delete it.
- Export the list of shared record groups to a .csv file.
 - Click the Export icon.
- Print the list of shared record groups.
 - Click the Print icon.

Modifying a Shared Record Group

When you edit a shared record group, you can do the following:

1. Perform a global search to retrieve the shared record group you want to modify.
or
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* check box, and then click the Edit icon.
2. The *Shared Record Group* editor contains the following tabs from which you can modify information:
 - **General:** You can change any of the information you entered when you created it, including its name. Changing the shared record group name does not impact the shared resource records in it.
 - **Extensible Attributes:** You can modify the attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab is displayed if you logged in as a superuser. For information, see [About Administrative Permissions](#) on page 160.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Deleting Shared Record Groups

Before you delete a shared record group, you must remove the zone associations in the group; otherwise, an error message appears when you delete. For information, see [Deleting Associated Zones](#) on page 686.

To delete a shared record group:

1. Perform a global search to retrieve the shared record group you want to modify.
or
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

Grid Manager moves the shared record group to the Recycle Bin, if enabled. Use the Recycle Bin feature to recover a deleted shared record group and retrieve the deleted zones. For information, see [Using the Recycle Bin](#) on page 64.

Managing Shared Resource Records

You can create shared A, AAAA, MX, SRV and TXT records. These resource records are similar to the non-shared resource records. The DNS server uses them to respond to queries in the same way as any other resource record. A shared resource record can belong to only one shared record group. This section describes how to add shared resource records to a group and how to modify and delete them. It includes the following sections:

- [Creating Shared Records](#)
- [Viewing Shared Records](#) on page 684
- [Modifying Shared Records](#) on page 685
- [Deleting Shared Records](#) on page 685

Creating Shared Records

After you create a shared record group, you can create its resource records.

To create a shared A, AAAA, MX, SRV or TXT record and add it to a group:

1. From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Add -> Shared Record**.
or
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group*. Expand the **Shared Records** tab and click the Add icon.

2. Select one of the following:
 - **Shared A Record**
 - **Shared AAAA Record**
 - **Shared MX Record**
 - **Shared SRV Record**
 - **Shared TXT Record**
3. Enter information in the *Shared Record* wizard. See the online Help or the following for information about each resource record:
 - For information about A records, see [Managing A Records](#) on page 660.
 - For information about AAAA records, see [Managing AAAA Records](#) on page 662.
 - For information about MX records, see [Managing MX Records](#) on page 665.
 - For information about SRV records, see [Managing SRV Records](#) on page 666.
 - For information about TXT records, see [Managing TXT Records](#) on page 668.
4. Save the configuration, or click **Next** to define extensible attributes for the shared record. For information, see [Using Extensible Attributes](#) on page 332.
5. Click **Restart** if it appears at the top of the screen.

Viewing Shared Records

You can view the shared records in a group and in a zone. To edit the shared record properties, click the shared record name and select the Edit icon.

To view the shared records in a group:

- From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* -> **Shared Records** tab.

To view the shared records in a zone:

- From the **Data Management** tab, select the **DNS** tab -> **Zones** tab and select a zone.

Grid Manager lists the following information about each shared record by default:

- **Name:** The shared record name.
- **Type:** Indicates the type of resource record, such as A, AAAA, MX, SRV or TXT records. Shared records are identified as **(Shared)**.
- **Data:** The data the shared resource record provides.
- **Comment:** Comments that were entered in the resource record.
- **Site:** Displays values that were entered for this pre-defined attribute.

You can display the following additional columns:

- **TTL:** The TTL value of the shared resource record.
- **Disabled:** Indicates whether the record is disabled.

You can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Edit the properties of a shared resource record.
 - Select the shared resource record, and then click the Edit icon.
- Delete a shared resource record.
 - Select the shared resource record, and then click the Delete icon.

- Export the list of shared resource records to a .csv file.
 - Click the Export icon.
- Print the list of shared resource records.
 - Click the Print icon.

Modifying Shared Records

You can modify, disable, or delete any shared record. When physical repair or relocation of a network device occurs, you can disable a record instead of deleting it. This alleviates having to delete, and then add the shared record again. When the changes to the physical device are complete, you can simply enable the shared record.

To modify or disable a shared record:

1. Perform a global search to retrieve the host or resource record you want to modify.
or
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* -> **Shared Records** tab.
2. Select the shared record you want to modify and click the Edit icon.
3. The *Shared Records* editor contains the following tabs from which you can modify information:
 - **General**: You can change most of the information, except for the read-only fields, such as the Host Name Policy. You can also select the **Disable** check box to disable the record.
 - **TTL**: You can modify the TTL setting. For information, see [About Time To Live Settings](#) on page 557.
 - **Extensible Attributes**: You can modify the attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions**: This tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#) on page 160.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Deleting Shared Records

To delete shared resource records:

1. Perform a global search to retrieve the record you want to delete.
or
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* -> **Shared Records** tab.
2. Select the shared record you want to delete and click the Delete icon.
3. When the confirmation dialog box displays, select **Yes**.
Grid Manager moves the shared records to the Recycle Bin, from which you can restore or permanently delete the records.

Managing Associated Zones

Typically, you associate a zone with a shared record group when you create the group. You can also add an associated zone to a shared record group after you create the group.

Creating Associated Zones

To associate a zone with a share record group:

1. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* -> **Associated Zones** tab, and then click the Add icon.
2. In the *Zone Selector* dialog box, select a zone by clicking the zone name.
The appliance adds the zone to the **Associated Zones** tab.

Viewing Associated Zones

To view the associated zones in a shared record group:

- From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* -> **Associated Zones** tab.

Grid Manager lists the following information about each associated zone by default:

- **Zone:** The zone associated with the shared record group.
- **DNS View:** The DNS view to which the zones belong.
- **Network View:** The network view associated with the DNS view.
- **Comment:** Comments that were entered for the shared record group.

You can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Associate another zone with the shared record group.
 - Click the Add icon and select a zone.
- Delete an associated zone.
 - Select the zone, and then click the Delete icon.
- Export the list of associated zones to a .csv file.
 - Click the Export icon.
- Print the list of shared associated zones.
 - Click the Print icon.

Deleting Associated Zones

To delete an associated zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared_record_group* -> **Associated Zones** tab.
2. Select the associated zone and click the Delete icon.
3. When the confirmation dialog box displays, select **Yes**.

Grid Manager removes the zone from the shared record group.

Configuration Example: Configuring Shared Records

The following example shows you how to configure shared records. In this example, you do the following:

- Create a shared record group: **group1**.
 - Associate it with three zones: **eng.com**, **sales.com**, and **marketing.com**.
 - Create an A record **www** and an MX record **mx1**.
1. Create a shared record group called **group1** and associate it with **eng.com**, **sales.com**, and **marketing.com**.
 - a. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab, and then click the Add icon.
 - b. In the first step of the *Shared Record Group* wizard, specify the following
Name: Enter **group1**.
 - c. Click **Next**.
 - d. Click the Add icon in the Associated Zones panel.
 - e. Select **eng.com** from the list of zones and click the select icon. Do the same for the **sales.com**, and **marketing.com** zones.
 - f. Save the configuration and click **Restart** if it appears at the top of the screen.

2. Add an A record **www** to **group1**.
 - a. Expand the Toolbar and click **Add -> Shared Record > Shared A Record**.
 - b. In the *Shared A Record* wizard, specify the following:
 - Name:** Enter **www**.
 - Shared Record Group:** Select **group1** from the drop-down list.
 - IP Address:** Enter the IP address **10.9.1.1**.
 - c. Save the configuration and click **Restart** if it appears at the top of the screen.
3. Add an MX record **mx1** into **group1**.
 - a. Expand the Toolbar and click **Add -> Shared Record > Shared MX Record**.
 - b. In the *Shared MX Record* wizard, specify the following:
 - Mail Destination:** Enter **mx1**.
 - Shared Record Group:** Select **group1** from the drop-down list.
 - Mail Exchanger:** Enter **www.infoblox.com**.
 - Preference:** Enter **10**.
 - Comment:** Enter **mail exchanger record for shared record group1**.
 - c. Save the configuration and click **Restart** if it appears at the top of the screen.



Chapter 20 Configuring DDNS Updates from DHCP

DDNS (Dynamic DNS) is a method to update DNS data (A, TXT, and PTR records) from sources such as DHCP servers and other systems that support DDNS updates, such as Microsoft Windows servers 2000, 2003, 2008, 2008 R2, 2012, and 2012 R2. This chapter provides conceptual information about DDNS and explains how to configure NIOS appliances running DHCP, DHCPv6 and DNS to support DDNS updates. It contains the following main sections:

- [*Understanding DDNS Updates from DHCP*](#) on page 691
- [*Configuring DHCP for DDNS*](#) on page 695
 - [*Enabling DDNS for IPv4 and IPv6 DHCP Clients*](#) on page 695
 - [*Sending Updates to DNS Servers*](#) on page 696
- [*Configuring DDNS Features*](#) on page 697
 - [*Resending DDNS Updates*](#) on page 697
 - [*Generating Host Names for DDNS Updates*](#) on page 698
 - [*Configuring DDNS Features*](#) on page 698
 - [*Replacing Host Names for DDNS Updates*](#) on page 699
- [*About the Client FQDN Option*](#) on page 701
 - [*Enabling FQDN Option Support*](#) on page 702
 - [*Sending Updates for DHCP Clients Using the FQDN Option*](#) on page 703
- [*Configuring DDNS Update Verification*](#) on page 703
- [*Configuring DNS Servers for DDNS*](#) on page 705
 - [*Enabling DNS Servers to Accept DDNS Updates*](#) on page 706
 - [*Forwarding Updates*](#) on page 707
- [*Supporting Active Directory*](#) on page 709
 - [*Sending DDNS Updates to a DNS Server*](#) on page 709
- [*About GSS-TSIG*](#) on page 710
 - [*Sending Secure DDNS Updates to a DNS Server in the Same Domain*](#) on page 711
 - [*Configuring DHCP to Send GSS-TSIG Updates in the Same Domain*](#) on page 712
 - [*Sending Secure DDNS Updates to a DNS Server in Another Domain*](#) on page 719
 - [*Configuring DHCP to Send GSS-TSIG Updates to Another Domain*](#) on page 720
 - [*Sending GSS-TSIG Updates to a DNS Server in Another Forest*](#) on page 722
- [*Accepting DDNS Updates from DHCP Clients*](#) on page 723
 - [*Supporting Active Directory and Unauthenticated DDNS Updates*](#) on page 723

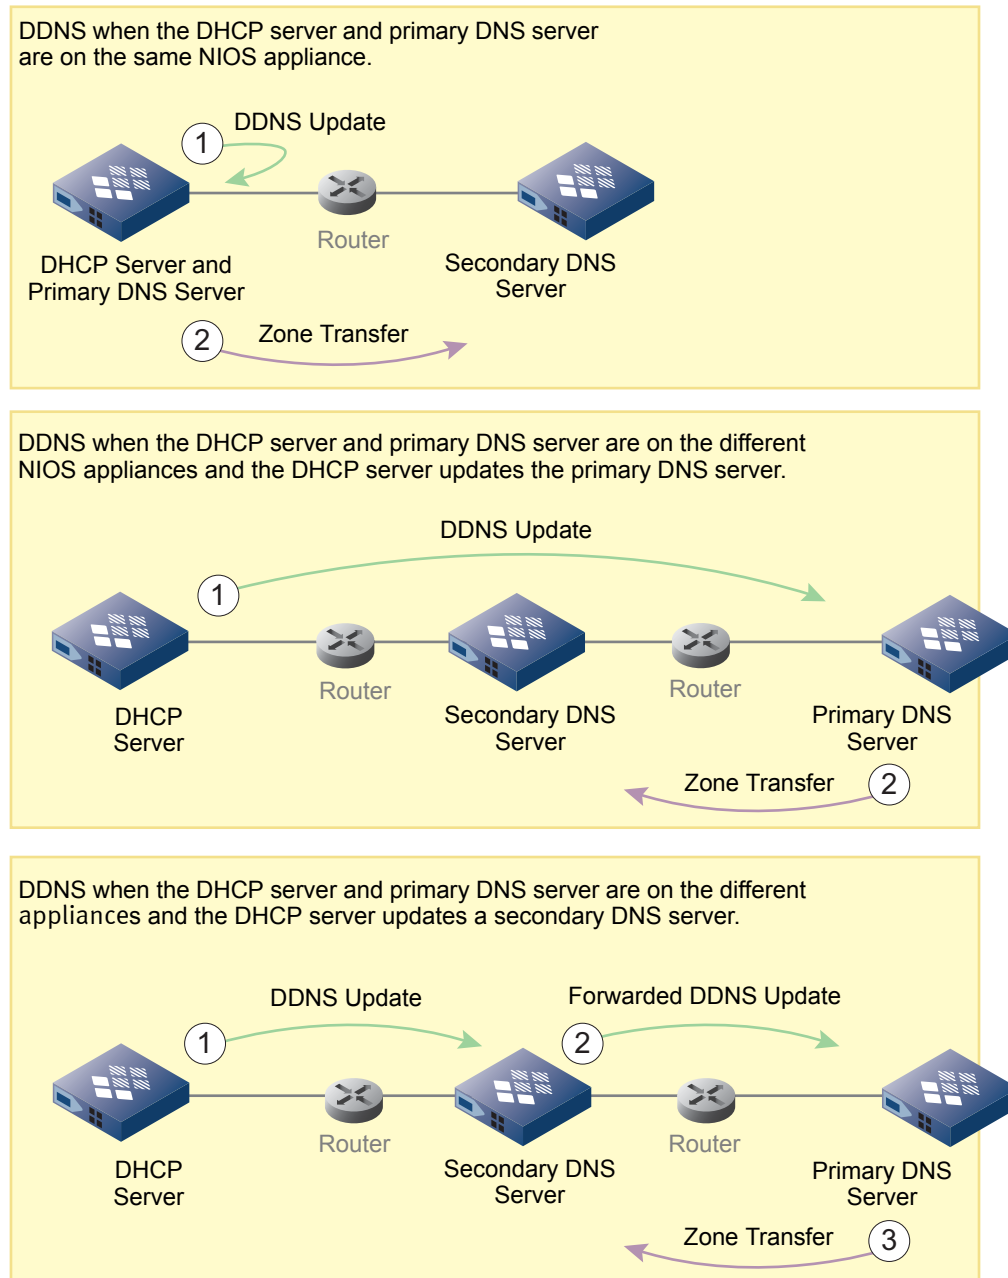
- [*Accepting GSS-TSIG-Authenticated Updates*](#) on page 725
 - [*Configuring DNS to Receive GSS-TSIG Updates*](#) on page 727

UNDERSTANDING DDNS UPDATES FROM DHCP

DHCP supports several DNS-related options (such as options 12, 15, and 81 for IPv4, and options 23, 24, and 39 for IPv6). With DDNS (Dynamic DNS) updates, a DHCP server or client can use the information in these options to inform a DNS server of dynamic domain name-to-IP address assignments.

To set up one or more NIOS appliances for DDNS updates originating from DHCP, you must configure at least one DHCP server and one DNS server. These servers might be on the same appliance or on separate appliances. Three possible arrangements for a DHCP server to update a DNS server are shown in [Figure 20.1](#).

Figure 20.1 Relationship of DHCP and DNS Servers for DDNS Updates

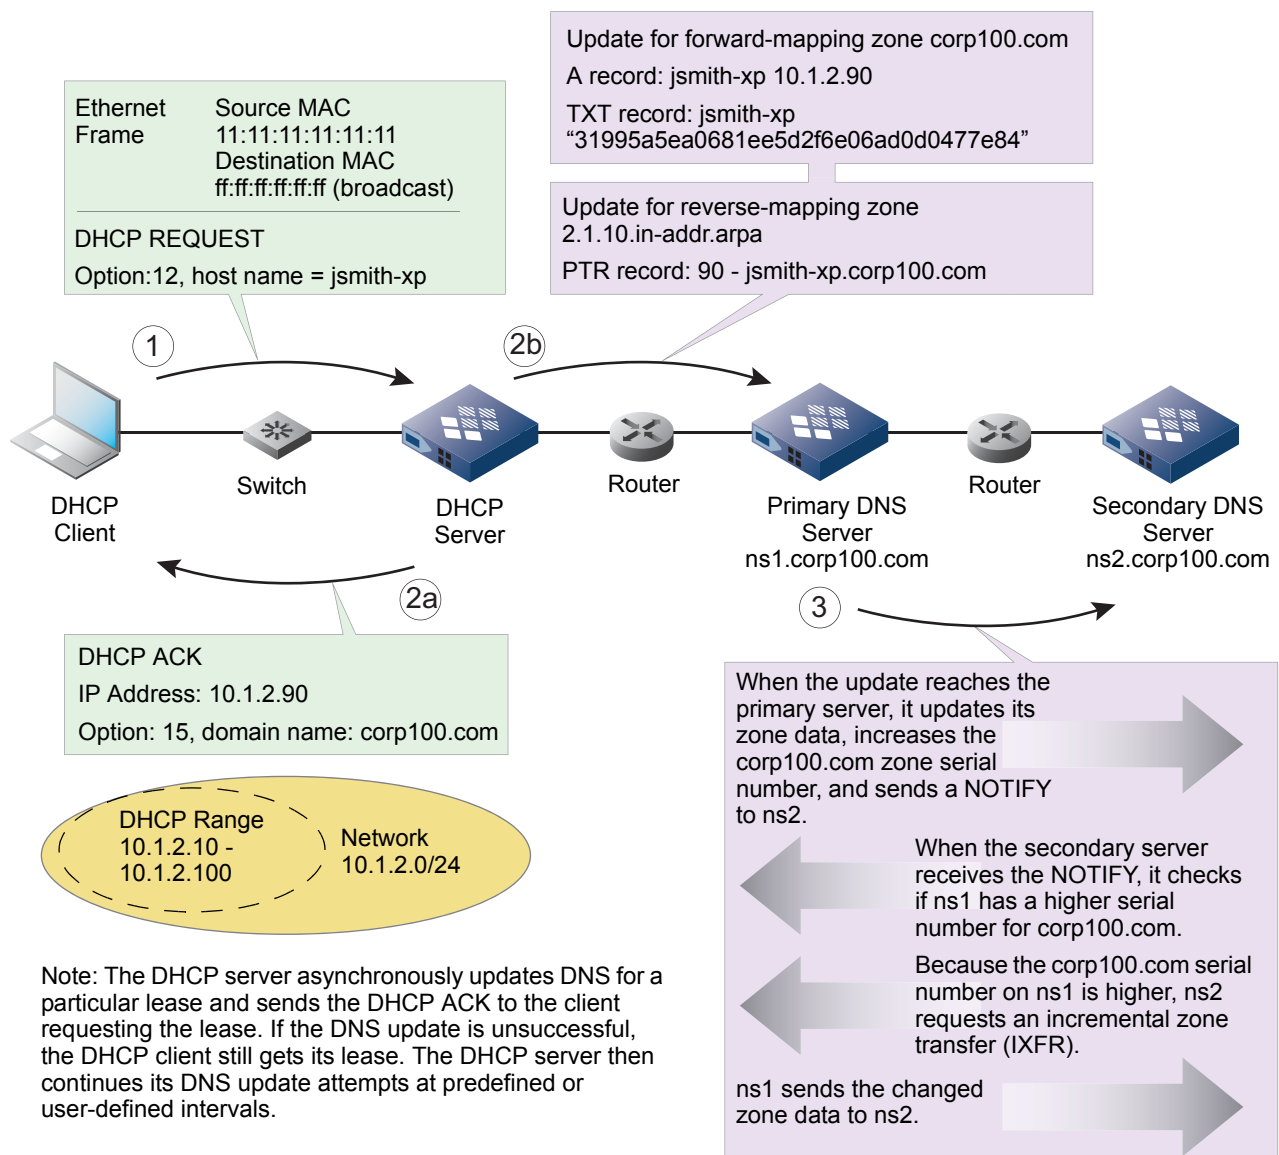


Here is a closer look at one setup for performing DDNS updates from a DHCP server (the steps relate to [Figure 20.2](#)).

1. When an IPv4 DHCP client requests an IP address, the client sends its host name (DHCP option 12). The client also includes its MAC address in the ethernet frame header.
2.
 - a. When the DHCP server responds with an IP address, it usually provides a domain name (DHCP option 15). The combined host name (from the client) and domain name (from the server) form an FQDN (fully qualified domain name), which the NIOS appliance associates with the IP address in the DHCP lease.
 - b. The DHCP server sends the A, TXT, and PTR records of the DHCP client to the primary DNS server to update its resource records with the dynamically associated FQDN + IP address.
3. The primary DNS server notifies its secondary servers of a change. The secondary servers confirm the need for a zone transfer, and the primary server sends the updated zone data to the secondary server, completing the update.

Note: For information about zone transfers, see [Enabling Zone Transfers](#) on page 583.

Figure 20.2 DDNS Update from a DHCP Server



To enable a DHCP server to send DDNS updates to a DNS server, you must configure both servers to support the updates. First, configure the DHCP server to do the following:

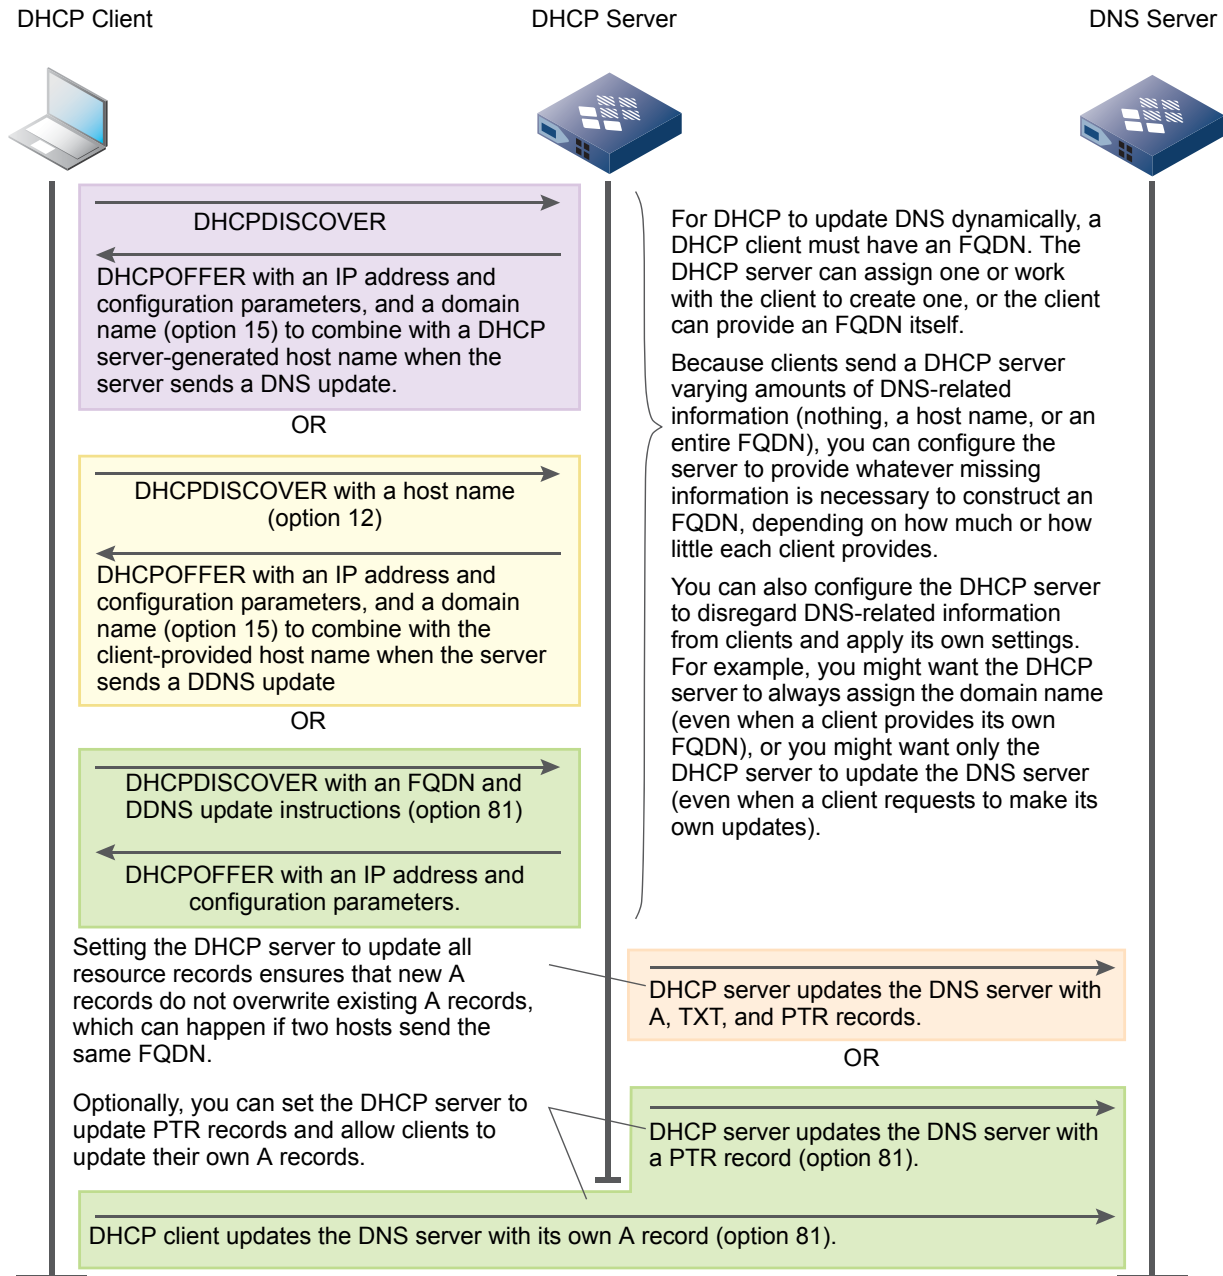
- Provide what is needed to create an FQDN: add a server-generated host name to a server-provided domain name, add a server-provided domain name to a client-supplied host name, or permit the client to provide its own FQDN
- Send updates to a DNS server

Then, configure the following on the DNS server:

- Accept updates from the DHCP server, a secondary DNS server, or a DHCP client
- If the DHCP server sends updates to a secondary DNS server, configure the secondary server to forward updates to the primary DNS server

When setting up DDNS, you can determine the amount of information that DHCP clients provide to a DHCP server—and vice versa—and where the DDNS updates originate. A summary of these options for IPv4 is shown in [Figure 20.3](#). It is similar for IPv6, except that the DHCP client and server exchange Request and Reply messages, AAAA records are updated instead of A records, and the FQDN option is option 39.

Figure 20.3 DHCP Clients and Server Providing DNS Information and Updates



You can configure the DHCP and DNS settings for DDNS at the Grid level, member level, and network and zone level. By applying the inheritance model in the NIOS appliance, settings made at the Grid level apply to all members in the Grid. Settings you make at the member level apply to all networks and zones configured on that member. Settings made at the network and zone level apply specifically to just that network and zone. When configuring independent appliances (that is, appliances that are not in a Grid), do not use the member-level settings. Instead, configure DDNS updates at the Grid level to apply to all zones and, if necessary, override the Grid-level settings on a per zone basis.

CONFIGURING DHCP FOR DDNS

Before a DHCP server can update DNS, the DHCP server needs to have an FQDN-to-IP address mapping. When a DHCP IPv4 client requests an IP address, it typically includes its host name in option 12 of the DHCPDISCOVER packet, and an IPv6 client includes its hostname in the Request packet. You can configure the NIOS appliance to include a domain name in option 15 of the IPv4 DHCPOFFER packet or in the IPv6 Reply packet. You specify this domain name in the **IPv4 DHCP Options** -> **Basic** and **IPv6 DHCP Options** -> **Basic** tabs of the *Grid DHCP Configuration* editor, *Member DHCP Configuration* editor, and the *Network editor*. For IPv4 clients you can also specify a domain name in the *DHCP Range* and *Fixed Address* editors.

Then, you can enable the DHCP server to send DDNS updates for IPv4 and IPv6 clients, as described in [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695. After you enable the DHCP server to send DDNS updates, you can do the following:

- Configure the DHCP server to send DDNS updates to DNS servers in the Grid. For information, see [Sending Updates to DNS Servers in the Grid](#) on page 696.
- Configure the DHCP server to send DDNS updates to external DNS servers. For information, see [Configuring DDNS Features](#) on page 697.
- Configure certain DDNS features. For information, see [Configuring DDNS Features](#) on page 698.
- Enable support for the FQDN option for IPv4 and IPv6 clients, and configure how the DHCP server updates DNS. For information, see [Enabling FQDN Option Support](#) on page 702.

Note: Whether you deploy NIOS appliance in a Grid or independently, they send updates to UDP port 53. Grid members do not send updates through a VPN tunnel; however, Grid members do authenticate updates between each other using TSIG (transaction signatures) based on an internal TSIG key.

Enabling DDNS for IPv4 and IPv6 DHCP Clients

You can enable the DHCP server to send DDNS updates for IPv4 clients at the Grid, member, shared network, network, address range, DHCP template, fixed address, and roaming host levels, and for IPv6 clients at the Grid, member, network, shared network, network template and roaming host levels.

You can specify a different domain name that the appliance uses specifically for DDNS updates. The appliance combines the hostname from the client and the domain name you specify to create the FQDN that it uses to update DNS. For IPv4 clients, you can specify the DDNS domain name at the network, network template, range, and range template levels. For IPv6 clients, you can specify the DDNS domain name at the Grid, member, network, shared network, and network template levels. You can also use the name of a roaming host record as the name of the client for DDNS updates, as described in [Setting Properties for Roaming Hosts](#) on page 867.

To enable DDNS and specify a DDNS domain name:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
Member: From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> **Members** -> *member* check box -> Edit icon.
Network: From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* check box -> Edit icon.
Network Template: From the **Data Management** tab, select the **DHCP** tab and click the **Templates** tab -> *DHCP_template* check box -> Edit icon.
Roaming Host: From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Roaming Hosts** -> *roaming_host* -> Edit icon.
 For IPv4 clients only:
IPv4 Address Range: From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* -> *addr_range* check box -> Edit icon.
IPv4 Fixed Address: From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* -> *ip_addr* check box -> Edit icon.

IPv4 Address Range/Fixed Address Template: From the **Data Management** tab, select the **DHCP** tab and click the **Templates** tab -> *DHCP_template* check box -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the **IPv4 DDNS** -> **Basic** tab or the **IPv6 DDNS** -> **Basic** tab, complete the following:
 - **Enable DDNS Updates:** Select this check box to enable DDNS updates. When setting properties for DHCP objects other than the Grid, you must click **Override** and select **Enable DDNS updates** for the DDNS settings to take effect.
 - **DDNS domain name:** Specify the domain name of the network that the appliance uses to update DNS. For IPv4 clients, you can specify this at the network, network template, range, and range template levels. For IPv6 clients, you can specify this at the Grid, member, network, shared network, and network template levels.
 - **DDNS Update TTL:** You can set the TTL used for A or AAAA and PTR records updated by the DHCP server. The default is shown as zero. If you do not enter a value here, the appliance by default sets the TTL to half of the DHCP lease time with a maximum of 3600 seconds. For example, a lease time of 1800 seconds results in a TTL of 900 seconds, and a lease time of 86400 seconds results in a TTL of 3600 seconds. For information about how to set the lease time, see [Defining Lease Times](#) on page 794.
 - **Update DNS on DHCP Lease Renewal:** Select this check box to enable the appliance to update DNS when a DHCP lease is renewed.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Sending Updates to DNS Servers

The DHCP server can send DDNS updates to DNS servers in the same Grid and to external DNS servers. When you enable the appliance to send updates to Grid members, you must specify the DNS view to be updated. If a network view has multiple DNS views, you can select only one DNS view for DDNS updates. For information about DNS views, see [Using Infoblox DNS Views](#) on page 602.

When you enable DDNS updates for a Grid, member, shared network, network, address range, DHCP template, fixed address, or roaming host, the DHCP server sends updates to authoritative zones using the domain name (as DHCP option 15) you define in the DHCP properties. You can also define forward-mapping zones that receive DDNS updates for DHCP clients that use option 81 to define the domain name. For information, see [About the Client FQDN Option](#) on page 701. To allow DDNS updates for clients using option 81, you must first enable the support for option 81. For information, see [Configuring DDNS Features](#) on page 697.

Sending Updates to DNS Servers in the Grid

You must specify the DNS view to be updated for each network view.

To configure the DHCP server to send updates to DNS servers in the same Grid:

1. If there are multiple network views in the Grid, select a network view.
 2. From the **Data Management** tab, select the **DHCP** tab, and then click **Configure DDNS** from the Toolbar.
 3. In the *DDNS Properties* editor, complete the following:
 - **DNS View:** If a network view has more than one DNS view, this field lists the associated DNS views. From the drop-down list, select the DNS view to which the DHCP server sends DDNS updates. Otherwise, the appliance uses the default DNS view.
 4. Save the configuration and click **Restart** if it appears at the top of the screen.
- The appliance sends DDNS updates to the appropriate zones in the selected DNS view. Note that you cannot delete a DNS view that has been selected for DDNS updates. By default, the DHCP server sends DDNS updates to zones using the domain name that you define for DHCP objects, such as networks and DHCP ranges.

Sending Updates for Zones on an External Name Server

The DHCP server can send dynamic updates to an external name server that you specify. For each network view, you can specify the zone to be updated and the IP address of the primary name server for that zone. You can add information for a forward and reverse zone. The DHCP server updates the A record in the forward zone and the PTR record in the reverse zone.

You can also use TSIG (transaction signatures) or GSS-TSIG to secure communications between the servers. TSIG uses the MD5 (Message Digest 5) algorithm and a shared secret key to create an HMAC (hashed message authentication code)—sometimes called a digital fingerprint—of each update. Both the DHCP server sending the update and the DNS server receiving it must share the same secret key. Also, it is important that the time stamps on the TSIG-authenticated updates and update responses be synchronized, or the participants reject them. Therefore, use an NTP server to set the time on all systems involved in TSIG authentication operations.

To send updates to a DNS server that is external to your Grid:

1. If there are multiple network views in the Grid, select a network view.
2. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Configure DDNS**.
3. In the **DDNS Updates to External Zones** section of the *DDNS Properties* editor, click the Add icon. Complete the following fields in the Add External DDNS Zone panel, and then click **Add**:
 - **Zone Name:** Enter the FQDN of a valid forward-mapping or reverse-mapping zone to which the DHCP server sends the updates. Do not enter the zone name in CIDR format. To specify a zone name in IDN, manually convert IDN to punycode and use the punycode representation.
 - **DNS Server Address:** Enter the IP address of the primary name server for that zone.
 - **Security:** Select one of the following security methods:
 - **None:** Select this to use unsecured DDNS updates. This is the default.
 - **TSIG:** Select this to use the standards-based TSIG key that uses the one-way hash function MD5 to secure transfers between name servers. You can either specify an existing key or generate a new key. To specify an existing key, complete the following:
Key Name: Enter the TSIG key name. The key name entered here must match the TSIG key name on the external name server.
Key Algorithm: Select either **HMAC-MD5** or **HMAC-SHA256**.
Key Data : To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
 - **GSS-TSIG:** For information about using GSS-TSIG, see [About GSS-TSIG](#) on page 710.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING DDNS FEATURES

You can enable the DHCP server to support certain DDNS features for IPv4 and IPv6 clients. These features affect the behavior of the DHCP server and how it handles DDNS updates. The following sections describe the different features you can set.

Resending DDNS Updates

You can enable the DHCP server to make repeated attempts to send DDNS updates to a DNS server. The DHCP server asynchronously updates DNS for a particular lease and sends the DHCP ACK to the client requesting the lease. If the update fails, the DHCP server still provides the lease and sends the DHCP ACK to the client. The DHCP server then continues to send the updates until it is successful or the lease of the client expires. You can change the default retry interval, which is five minutes.

You can enable this feature for the Grid and for individual Grid members.

Generating Host Names for DDNS Updates

Some IPv4 and IPv6 clients do not send a host name with their DHCP requests. When the DHCP server receives such a request, its default behavior is to provide a lease but not update DNS. You can configure the DHCP server to generate a host name and update DNS with this host name when it receives a DHCP request that does not include a host name. It generates a name in the following format: **dhcp-*ip_address***, where *ip_address* is the IP address of the lease. For example, if this feature is enabled and the DHCP server receives a DHCP REQUEST from an IPv4 DHCP client with IP address 10.1.1.1 and no host name, the DHCP server generates the name dhcp-10-1-1-1 and appends the domain name, if specified, for the DDNS update. Likewise, if an IPv6 client with IP address 2001:db8:a23:0:0:0:d sends a request, the DHCP server generates the name dhcp-2001-db8-a23-0-0-0-d and appends the domain name, if specified, for the DDNS update.

Updating DNS for IPv4 Clients with Fixed Addresses

By default, the DHCP server does not update DNS when it allocates an IPv4 or IPv6 fixed address to a client. You can configure the DHCP server to update the A and PTR record of IPv4 clients with a fixed address. When you enable this feature and the DHCP server adds A and PTR records for a fixed address, the DHCP server never discards the records. When the lease of the client terminates, you must delete the records manually. Note that the DHCP server does not send DDNS updates for IPv6 fixed addresses and hosts.

You can define fixed address settings for the Grid, Grid members, IPv4 networks, and IPv4 shared networks.

Configuring DDNS Features

You can configure DDNS features for a Grid, its member, IPv4 and IPv6 networks and shared networks, and IPv4 DHCP address ranges. You cannot set DDNS features for IPv6 DHCP ranges. To configure DDNS features:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Configuration**.
 - Member:** From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* check box -> Edit icon.
 - Network:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* check box -> Edit icon.
 - Shared Network:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Shared Networks** -> *shared_network* check box -> Edit icon.
 - DHCP Range:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* -> *addr_range* check box -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **DDNS** -> **Advanced** tab for the Grid and member, or the **DDNS** -> **Basic** tab for the network, do the following:
 - **Update Retry:** You can set this for a Grid and its members only.
 - **Retry Updates When Server Becomes Available:** Select this check box.
 - **Retry interval (Minutes):** You can optionally set the retry interval. The default is five minutes.
 - **Generate Hostname**
 - **Generate Hostname if not Sent by Client:** Select this check box to enable the DHCP server to generate a hostname and update DNS with this hostname, when the DHCP request of a client does not include a hostname.
 - **Fixed Address Updates:** You can set this for IPv4 fixed addresses only. This option is available in the **IPv4 DDNS Advanced** tab of the *Grid DHCP Properties* and *Member DHCP Properties* editors, and in the **IPv4 DDNS Basic** tab of the *IPv4 DHCP Network* and *Shared Network* editors.
 - **Update Fixed Addresses:** Select this check box to allow the DHCP server to send updates to DNS for IPv4 fixed addresses.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

When a lease expires, the DHCP server removes the A, AAAA, and PTR records that it updated. It does not remove any records that the client updated.

Replacing Host Names for DDNS Updates

In situations where you need to restrict the use of specific characters in a host name for DDNS updates, you can configure a hostname rewrite policy. Such policy accepts certain characters and replaces others in host names specified in IPv4 DHCP requests. When you create a hostname rewrite policy, you enter a list of valid characters that the appliance accepts in the host name. You also specify a character that the appliance uses to replace invalid characters. You can create multiple hostname rewrite policies on the appliance, but you can only enable one policy at any given time. The appliance provides a default policy that includes `a-z0-9_` as valid characters and dash (-) as the replacement character. You cannot modify or delete the default policy.

When you enable a hostname rewrite policy, the appliance replaces host names with the newly translated host name when it issues DHCP leases and sends DDNS updates for IPv4 DHCP clients. For information about how to add and enable a hostname rewrite policy, see [Adding and Enabling a Hostname Rewrite Policy](#) on page 699.

Before you enable a hostname rewrite policy, consider the following:

- You must enable DDNS updates before the hostname rewrite policy can take effect.
- You can use a hostname rewrite policy only if MS code pages are disabled.
- The policy supports only IPv4 DHCP clients.
- If DHCP option 81 support is enabled and updating DDNS is in the request, the appliance sends updates for A records directly to the DNS server and DHCP only updates the PTR record. When this happens, there can be a mismatch in the host name between the A and PTR records.
- Changes made to a hostname rewrite policy apply only to subsequent DDNS updates.

When an IPv4 DHCP client requests an IP address, it includes its host name in DHCP option 12. If you enable a hostname rewrite policy, the appliance uses the newly translated host name when it issues a lease to the client.

The client can also include a FQDN in option 81, in which it instructs the server whether to perform DDNS updates. If the client sends a FQDN in option 81, the appliance replaces the entire FQDN based on the policy. For example, if the FQDN in option 81 is `dev.bldg12.corp100.com`, the appliance replaces invalid characters in the entire FQDN even though the host name can be `dev` or `dev.bldg12`. For example, if your hostname rewrite policy specifies valid characters as `a-z` and the replacement character is `-`, the newly translated FQDN is `dev.bldg--corp---.com`. For information about client FQDN in option 81, see [About the Client FQDN Option](#) on page 701.

Note that when multiple IPv4 DHCP clients specify host names that map to the same translated host name, the appliance allocates leases to all clients, but it only sends DDNS updates to the first client request. When it tries to update DNS for subsequent clients, the updates fail.

You can add and enable a hostname rewrite policy for the entire Grid. You can also override the policy at a member level, as described in [Overriding a Grid Hostname Rewrite Policy](#) on page 700.

Adding and Enabling a Hostname Rewrite Policy

To add and enable a hostname rewrite policy, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Configuration**.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. In the **IPv4 DDNS -> Advanced** tab, click the Add icon in the **Hostname Rewrite Policy** section:
 - **Policy Name:** Enter the policy name. Each policy name must be unique.
 - **Valid Characters:** Enter a list of valid characters you want to keep in the host name. Ensure that you consider the following rules:
 - You can include only printable ASCII characters and space.
 - The appliance includes period (.) as a valid character by default. You do not need to specify it.

- You can also use shortcuts for a series or range of characters. For example, when you enter **a-d**, the appliance includes the following: A, B, C, D, a, b, c, and d. When you enter **0-5**, the appliance includes the following: 0, 1, 2, 3, 4, and 5. In a character range, ensure that the start character is less than the end character.
 - If you want to use dash (-) as a character, ensure that you put it in front of the valid character pattern. Otherwise, the appliance treats the string as a range of characters.
 - You can build a POSIX regular expression based on the string you enter here, but you cannot enter an empty string.
 - You cannot use the meta character (^) as a start or end character in a range. For example, **a-^** is invalid. You also cannot use duplicate characters as character sets. For example, **aa** is invalid.
- **Replace Invalid Characters with:** Enter a character the appliance uses to replace invalid characters. Only enter one printable ASCII character. You cannot enter multiple characters or use space as the replacement character.

To test the hostname policy before adding it to the system, enter a sample hostname in the **Sample Host Name** field, and then click **Test**. The appliance displays the translated hostname. You can change the policy and test it again until you get the desired result.

Click the Add icon to add the new hostname rewrite policy to the table. The appliance comes with a default policy that includes **a-z0-9_** as valid characters and dash (-) as the replacement character. Grid Manager displays the following for each policy:

- **Policy Name:** The name of the hostname rewrite policy.
- **Valid Characters:** Valid characters for the host name.
- **Replace Invalid Characters with:** The character used to replace invalid characters in the host name.

You can also select a hostname policy and click the Edit icon to modify it, or click the Delete icon to delete it. You cannot modify or delete the default policy. For information about how to modify a policy, see [Modifying a Hostname Rewrite Policy](#) on page 700.

4. Complete the following to enable the hostname rewrite policy:
 - **Enable hostname rewrite policy:** Select this check box to use a hostname rewrite policy for DHCP leases and DDNS updates for IPv4 DHCP clients. From the drop-down list, select the hostname policy you want to use.
5. Save the configuration.

Modifying a Hostname Rewrite Policy

To modify a hostname rewrite policy, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Configuration**.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. In the **IPv4 DDNS -> Advanced** tab, do the following in the **Hostname Rewrite Policy** section:
 - Select a policy from the table, and then click the Edit icon.
 - In the Edit Hostname Rewrite Policy section, modify and test the policy as described in [Adding and Enabling a Hostname Rewrite Policy](#) on page 699.

Note: If you enable the policy at the Grid level, you can modify all information, including the policy name. If you enable the policy at the member level, you can modify any information, except for the policy name.

4. Click **Save**. The appliance updates the policy in the table.
5. Save the configuration.

Overriding a Grid Hostname Rewrite Policy

You can override a Grid hostname rewrite policy at the member level. To override a Grid policy, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* check box -> *Edit* icon.

2. In the *Member DHCP Properties* editor, click **Toggle Advanced Mode**.
3. In the **IPv4 DDNS** -> **Advanced** tab, click **Override** in the **Hostname Rewrite Policy** section, and then complete the following:
 - **Enable hostname rewrite policy:** Select this check box to use a hostname rewrite policy, or deselect the check box to disable the policy.
 - When you enable this feature, select a policy that you want to use from the drop-down list. Grid Manager displays all hostname rewrite policies that you have configured on the appliance in the drop-down list. After you select a policy, Grid Manager displays the policy name, valid characters, and the replacement character.
4. Save the configuration.

ABOUT THE CLIENT FQDN OPTION

When an IPv4 DHCP client sends DHCP DISCOVER and DHCP REQUEST messages, it can include option 81, the Client FQDN option. An IPv6 DHCP client can include option 39, the Client FQDN option, when it sends Solicit and Request messages.

The Client FQDN option contains the FQDN (fully qualified domain name) of the client and instructions on whether the client or the server performs DDNS updates. You can configure the appliance to replace the FQDN in the option by defining a hostname rewrite policy. For information about adding and enabling a hostname rewrite policy, see [Replacing Host Names for DDNS Updates](#) on page 699.

The DHCP server can support option 81 for IPv4 and IPv6 clients, and use the host name or FQDN that the client provides for the update. It can also allow or deny the client's request to update DNS, according to the administrative policies of your organization. The DHCP server indicates its response in the DHCP OFFER message it sends back to an IPv4 client, and in the Reply message it sends back to an IPv6 client.

Sending Updates with the FQDN Option Enabled

When you enable the DHCP server to support the FQDN option, it uses the information provided by the IPv4 or IPv6 client to update DNS as follows:

- When an IPv4 or IPv6 DHCP client sends a DHCP request with the FQDN option, it can include either its FQDN or only its host name.
 - If the request includes the FQDN, the DHCP server uses this FQDN to update DNS. You can specify a list of forward-mapping zones to be updated for IPv4 and IPv6 clients using the FQDN option. For information, see [Sending Updates for DHCP Clients Using the FQDN Option](#) on page 703.
 - If the request includes the host name, the DHCP server provides the domain name. It combines the host name of the client and the domain name to create an FQDN for the client. It then updates DNS with the FQDN it created. (You can enter the domain name in the General page of the DHCP Properties window. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.)
- When a DHCP client sends a DHCP request with its hostname, the DHCP server adds the domain name you specified to create an FQDN for the client. It then updates DNS with the FQDN it created. For information about entering the domain name, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
- When a DHCP client does not send a host name, the DHCP server provides a lease but does not update DNS. You can configure the DHCP server to generate a host name and update DNS as described in [Generating Host Names for DDNS Updates](#) on page 698.
- If multiple DHCP clients specify the same FQDN or host name, the DHCP server allocates leases to the clients, but updates DNS only for the client that first sent the request. When it tries to update DNS for the succeeding clients, the update fails.

Sending Updates from DHCP Clients or a DHCP Server

When you enable the DHCP server to support the FQDN option, you must decide if you want the DHCP server to allow clients to update DNS. If you allow the client to update DNS, then the client updates its A or AAAA record only. The DHCP server always updates the PTR records. You can configure the DHCP server as follows:

- The DHCP server can allow clients to update DNS when they send the request in the FQDN option. This is useful for small sites where security is not an issue or in sites where clients move from one administrative domain to another and want to maintain the same FQDN regardless of administrative domain.

If you configure the DHCP server to allow clients to perform DDNS updates, you must also configure the DNS server to accept these updates from clients. Note that multiple clients can use the same name, resulting in multiple PTR records for one client name.

When a lease expires, the DHCP server does not delete the A or AAAA record, if it was added by the client.

- The DHCP server can refuse the DHCP client's request to update DNS and always perform the updates itself. When the DHCP server updates DNS, it uses the FQDN provided by the DHCP client. Select this option if your organization requires tighter control over your network and does not allow clients to update their own records.

If you do not enable support for the FQDN option and a client includes it in a DHCP request with its FQDN, the DHCP server does not use the FQDN of the client. Instead, it creates the FQDN by combining the host name from the client with the domain name specified in the Grid or Member DHCP Configuration editor.

Do the following to configure support for the FQDN option for both IPv4 and IPv6 clients:

- Enable support for the option and specify who performs the DDNS update. For more information, see [Enabling FQDN Option Support](#) on page 702.
- Specify the DNS zones and DNS view for the updates. For more information, see [Sending Updates for DHCP Clients Using the FQDN Option](#) on page 703.

Enabling FQDN Option Support

You can configure support for the FQDN option for IPv4 and IPv6 clients at the Grid, member, network and shared network levels.

To configure support for the FQDN Option (option 81) for IPv4 and (Option 39) for IPv6:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
Member: From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* check box -> Edit icon.
Network: From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* check box -> Edit icon.
Shared Network: From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Shared Networks** -> *shared_network* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **IPv4 DDNS** -> **Advanced** tab for the Grid and member, or the **IPv4 DDNS** -> **Basic** tab for the network, do the following:
 - **Option 81 Support**
 - **Enable Option 81 Support:** Select this to enable the support for option 81.
 - **DHCP server always updates DNS:** Select this to allow the DHCP server to update DNS, regardless of the requests from DHCP clients.
 - **DHCP server updates DNS if requested by client:** Select this to allow the DHCP server to update DNS only when requested by DHCP clients.
3. In the **IPv6 DDNS** -> **Advanced** tab for the Grid and member, or the **IPv6 DDNS** -> **Basic** tab for the network, do the following:
 - **FQDN Support:** Select **Enable FQDN Support** and select one of the following to indicate whether the DHCP server or the client performs the DDNS update.
 - DHCP always updates DNS

- DHCP updates DNS if requested by client

4. Save the configuration and click **Restart** if it appears at the top of the screen.

When a lease expires, the DHCP server removes the A or AAAA records and PTR records that it updated. It does not remove any records that the client updated.

Sending Updates for DHCP Clients Using the FQDN Option

You must specify the DNS view to be updated for each network view.

To send updates to zones for DHCP IPv4 and IPv6 clients using the FQDN option:

1. If there are multiple network views in the Grid, select a network view.
2. From the **Data Management** tab, select the **DHCP** tab, and then click **Configure DDNS** from the Toolbar.
3. In the *DDNS Properties* editor, complete the following:
 - **DNS View:** If a network view has more than one DNS view, this field lists the associated DNS views. From the drop-down list, select the DNS view to which the DHCP server sends DDNS updates. Otherwise, the appliance uses the default DNS view.
 - **Zones to Update for Hosts Using DHCP FQDN Option:** In this section, you can define forward-mapping zones to which the DHCP server sends DDNS updates for IPv4 and IPv6 DHCP clients that use the FQDN option. You must first enable support for the FQDN option before the DHCP server can send DDNS updates to these zones. By default, the DHCP server sends DDNS updates to zones using the domain name that you define for DHCP objects, such as networks and DHCP ranges. For clients using this option, the DHCP server uses the domain name defined in the option.
Click the Add icon to specify a forward-mapping zone. Note that the Forward-mapping Zone Selector dialog box displays only the DNS zones that are associated with the selected DNS view. The zones you select here are written to the `dhcpd.conf` file and the `dhcpdv6.conf` file as “zone” statements with the matching TSIG key of the DNS view, so the updates are sent to the correct DNS view.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING DDNS UPDATE VERIFICATION

The DHCP server can handle DDNS updates differently, depending on how stringently you configure record handling. You can configure the DHCP server to update records only after passing verification. You can adjust the way DHCP handles updates so the DHCP server updates records after passing less stringent verification requirements, or without any type of verification.

To provide a measure of protection against unintentional changes of DNS data, NIOS appliances support the generation and use of TXT records, as described in IETF draft, *draft-ietf-dhc-dhcp-dns-12.txt* and by the ISC (Internet Systems Consortium). When DHCP updates or deletes an A or AAAA record, the corresponding TXT record is checked first to verify the authenticity of the update. The TXT record is based on a hash of the DHCPID which is unique to each client, usually based in part on the MAC address or the DUID. If the client requests an update to DNS, the DHCP server first checks the TXT record to verify that it matches the client that originally inserted the record. This process provides assurance that the updates are from the same client. These security checks are based upon inserting a cryptographic hash of the DHCPID (DHCP Client Identifier) into a DNS TXT RR and then verifying that value before updating. For example, a sample client update adds the following records in DNS:

oxcart.lo0.net.	21600	IN A 172.31.1.20
oxcart.lo0.net.	21600	IN TXT "313ce164780d34b91486b7c489ed7467e6"
20.1.31.172.in-addr.arpa.	21600	IN PTR oxcart.lo0.net.

However, your DNS configuration might require that the NIOS appliance handle DNS record updates differently than described in draft-ietf-dhc-dhcp-dns-12.txt. Your specific requirements might benefit from less-stringent verification of the DHCID, or might require skipping verification entirely. Verification checks might cause complications in some specific cases described below:

- **Mobility:** The TXT record is based on the DHCID unique to each client and is usually based on the MAC address or DUID of the interface. Devices such as laptops that connect to both wired and wireless networks have different MAC addresses or DUIDs and different DHCID values for each interface. In this scenario, after either one of the network interfaces inserts a DNS record, updates are allowed from that interface only. This results in a disruption of service for DDNS updates when roaming between wired and wireless networks.
- **Migration:** The second problem occurs during a migration from non-ISC based systems to ISC systems. For example, if the user is migrating from a Microsoft-based system, the clients have A or AAAA and PTR records in the DDNS updates but no TXT records. As a result, new DDNS updates fail after the migration.
- **Mixed Environments:** The final problem occurs in mixed ISC and non-ISC environments. For example, assume that both Microsoft and ISC DHCP servers update DNS records on the appliance. Since the Microsoft DHCP server does not insert the TXT records, updates from ISC-based systems fail while updates from the Microsoft DHCP server are committed into the database.

The NIOS appliance offers four modes to handle DDNS updates as described in [Figure 20.4](#) on page 704:

Figure 20.4 DDNS Update Verification Mode

Mode	If a Record at Lease Grant	Then TXT Record at Lease Grant	Lease Grant Action	Lease Expire Action
Standard ISC	Exists	Must match	Delete A or AAAA, TXT if exists Add A or AAAA Add PTR	Delete PTR Delete A or AAAA, TXT if TXT matches and no other A or AAAA RRs
	No A or AAAA record	No check	Add A or AAAA, TXT Add PTR	
Check TXT only	Exists	Must exist	Delete A or AAAA, TXT Add A or AAAA, TXT Add PTR	Delete PTR Delete A or AAAA if TXT exists and no other A or AAAA RRs
	No A or AAAA record	No check	Add A or AAAA, TXT Add PTR	
ISC Transitional	Exists	No check	Delete A or AAAA, TXT if exists Add A or AAAA, TXT Add PTR	Delete PTR Delete A or AAAA, TXT if TXT matches and no other A or AAAA RRs
	No A or AAAA record	No check	Add A or AAAA, TXT Add PTR	
No TXT record	Exists	No check	Delete A or AAAA Add A or AAAA Add PTR	Delete PTR, A or AAAA
	No A or AAAA record	No check	Add A or AAAA Add PTR	

Depending on your expected usage, you must carefully consider the various options for update verification. The following section illustrates recommendations for each verification option:

- **Standard ISC:** This method is the most stringent option for verification of updates. This is the default.
- **ISC Transitional:** This method is useful during migrations from systems that do not support the TXT record to systems that are ISC-based.
- **Check TXT only:** This method is useful for the roaming laptop scenario. The NIOS appliance checks that a TXT record exists, but does not check the value of the TXT record.
- **No TXT record:** This method should be used with caution because anyone can send DDNS updates and overwrite records. This method is useful when both ISC and non-ISC-based DHCP servers and clients are updating the same zone. Infoblox recommends that you allocate a DNS zone for this authentication method, as a precaution.

Note: In certain situations, when a DHCP lease expires, the DHCP server might remove the TXT record even if there is no A or AAAA record.

You can enable this feature at the Grid level. To configure TXT record handling on the DHCP server:

1. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Configuration**.
2. In the **IPv4 DDNS -> Advanced** tab or the **IPv6 DDNS -> Advanced** tab, select one of the following from the **TXT (DHCID) Record Handling** drop-down list:
 - **Check Only:** Select this check box to enable minimal checking of DDNS updates. Specifically, A or AAAA records are modified only if a TXT record exists. The NIOS appliance checks that a TXT record exists, but does not check its value.
 - **ISC:** Select this check box to enable standard ISC (Internet Systems Consortium) handling for DDNS updates. Specifically, A or AAAA records are modified or deleted only if the TXT records match. This option is the default setting on the appliance.
 - **ISC Transitional:** Select this check box to enable less stringent handling of DDNS updates. Specifically, the NIOS appliance enables you to add or modify A or AAAA records whether or not TXT records exist. It checks whether a TXT record exists and then processes the update. If the appliance does not find a TXT record, it adds the record.
 - **No TXT Record:** Select this check box to disable TXT record checking. Specifically, A or AAAA records are added, modified, or deleted whether or not the TXT records match. No TXT records are added, and existing TXT records are ignored.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING DNS SERVERS FOR DDNS

For security reasons, an Infoblox DNS server does not accept DDNS updates by default. You must specify the sources from which you want to allow the DNS server to receive updates. You can configure the Infoblox DNS server to receive updates from specified DHCP clients, as described in [Enabling DNS Servers to Accept DDNS Updates](#), and to accept forwarded updates from another DNS server, as described in [Forwarding Updates](#) on page 707.

For protection against spoofed IP addresses, you can use TSIG (transaction signatures) to authenticate and verify updates.

TSIG uses the MD5 (Message Digest 5) algorithm and a shared secret key to create an HMAC (hashed message authentication code)—sometimes called a *digital fingerprint*—of each update. Both the DHCP server sending the update and the DNS server receiving it must share the same secret key. Also, it is important that the time stamps on the TSIG-authenticated updates and update responses be synchronized, or the participants reject them. Therefore, use an NTP server to set the time on all systems involved in TSIG authentication operations.

The TSIG key that you use can come from several places:

- You can use the key generation tool described in this section to create a new TSIG key to authenticate updates from the DHCP server.

- You can enter (copy and paste) a TSIG key that you previously generated for another purpose, such as for zone transfers.
- If the DHCP server is on a separate appliance and a TSIG key was previously generated on that appliance, you can enter (copy and paste) that TSIG key onto the local DNS server.

The TSIG key name and value that the DHCP and DNS servers use must be the same.

Note: Whether you deploy NIOS appliances in a Grid or independently, they send updates to UDP port 53. Grid members do not send updates through a VPN tunnel. Grid members do, however, authenticate updates between them using TSIG (transaction signatures) based on an internal TSIG key.

Enabling DNS Servers to Accept DDNS Updates

You can configure the Infoblox DNS server to receive updates from specified DHCP clients only. You can set this for the Grid so that the Grid members receive DDNS updates only from the specified sources. Note that you specify the IP addresses of the sources of the updates and not the actual IP addresses in the DNS records being updated.

To configure the DNS server to accept updates from the specified sources:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
Zones: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* -> *zone* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**, select the **Updates** tab.

Note: Ensure that you understand how the appliance handles match lists before you specify the list of IP sources for DDNS updates, as described in [Managing the Order of Match Lists](#) on page 380.

3. In the Allow updates from section, select one of the following:
 - **None:** Select this to deny DDNS updates from all DHCP clients. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance receives DDNS updates from the sources that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
 - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.

- **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of a remote name server. This name must match the name of the same TSIG key on other name servers.
 - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
 - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.

Note: You must enable GSS-TSIG signed updates to receive DDNS updates from TSIG key based ACEs. For information about how to enable this, see [Accepting GSS-TSIG Updates](#) on page 730.

- **Any Address/Network:** Select this to receive DDNS updates from any IP addresses.
After you have added access control entries, you can do the following:
 - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
 - Reorder the list of ACEs using the up and down arrows next to the table.
 - Select an ACE and click the Edit icon to modify the entry.
 - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

Allow GSS-TSIG signed updates: This check box is selected only if you have enabled GSS-TSIG signed updates.

4. Optionally, you can:
 - Modify an item on the list by selecting it and clicking the Edit icon.
 - Remove an item from the list by selecting it and clicking the Delete icon.
 - Move an item up or down the list. Select it and drag it to its new position, or click the up or down arrow. The appliance applies permissions to items in the order they are listed.
5. Save the configuration.

Forwarding Updates

When a secondary DNS server receives DDNS updates, it must forward the updates to the primary server because it cannot update zone data itself. In such situations, you must enable the secondary server to receive updates from the DHCP server, and then forward them to the primary DNS server.

To configure the secondary server to accept and forward updates for all zones:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
Zones: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns_view* -> *zone* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **Updates** tab, and then complete the following:
 - **Allow secondary name servers to forward updates:** Select this check box.

Forward updates from: This is available only for authoritative zones. Click **Add**. Depending on the item that you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:

- **None:** Select this to deny DDNS updates from any clients. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance receives DDNS updates from the sources that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
 - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
 - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
 - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
 - **Permission:** Select **Allow** or **Deny** from the drop-down list.
 - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
 - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of a remote name server. This name must match the name of the same TSIG key on other name servers.
 - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
 - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.

Note: You must enable GSS-TSIG signed updates to receive DDNS updates from TSIG key based ACEs. For information about how to enable this, see [Accepting GSS-TSIG Updates](#) on page 730.

- **Any Address/Network:** Select to allow or disallow the appliance to receive DDNS updates from any IP address.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

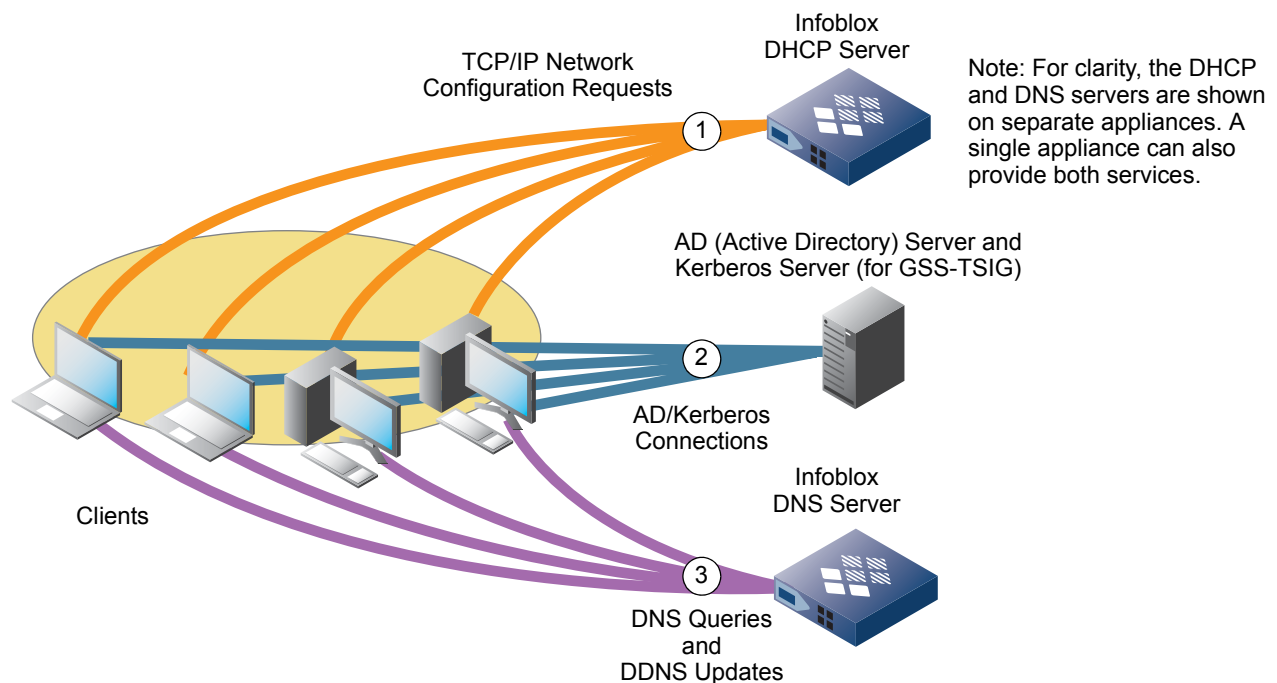
SUPPORTING ACTIVE DIRECTORY

Active Directory™ (AD) is a distributed directory service that authenticates network users and—by working with DHCP and DNS—provides the location of and authorizes access to services running on devices in a Windows® network.

You can integrate a NIOS appliance providing DHCP and DNS services with servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 with the Active Directory service installed. Assuming that you already have AD set up and it is currently in use, you can migrate DHCP and DNS services away from internal operations on the AD domain controller or from other third party DHCP and DNS systems to NIOS appliances that serve DHCP and DNS.

A NIOS appliance providing DHCP and DNS services to an AD environment can send and receive DDNS updates. In addition, a NIOS appliance can use GSS-TSIG (Generic Security Service-Transaction Signatures) authentication for DDNS updates. The basic DHCP, AD, and DNS services are shown in [Figure 20.5](#).

Figure 20.5 DHCP, Active Directory, and DNS



Sending DDNS Updates to a DNS Server

You can configure an Infoblox DHCP server to send unauthenticated or GSS-TSIG-authenticated DDNS updates to a DNS server in an AD domain. There are no special configurations to consider when configuring a NIOS appliance to send unauthenticated DDNS updates to the DNS server. (For information about configuring DHCP, see [Chapter 24, Configuring DHCP Properties](#), on page 791; and for information on configuring the DHCP server to send DDNS updates, see [Configuring DHCP for DDNS](#) on page 695.) For information about configuring a DHCP server to send GSS-TSIG authenticated updates, see [About GSS-TSIG](#) on page 710.

ABOUT GSS-TSIG

GSS-TSIG is used to authenticate DDNS updates. It is a modified form of TSIG authentication that uses the Kerberos v5 authentication system.

GSS-TSIG involves a set of client/server negotiations to establish a “security context”. It makes use of a Kerberos server (running on the AD domain controller) that functions as the Kerberos Key Distribution Center (KDC) and provides session tickets and temporary session keys to users and computers within an Active Directory domain. The client and server collaboratively create and mutually verify transaction signatures on messages that they exchange. Windows 2000 server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, and Windows Server 2012 R2 all support DDNS updates using GSS-TSIG.

Note: For information about GSS-TSIG, see *RFC 3645, Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)*.

A NIOS appliance can use GSS-TSIG authentication for DDNS updates for either one of the following:

- A NIOS appliance serving DHCP can send GSS-TSIG authenticated DDNS updates to a DNS server in an AD domain whose domain controller is running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2. The DNS server can be in the same AD domain as the DHCP server or in a different domain.
 - For information about sending secure DDNS updates to a DNS server in the same domain, see [Sending Secure DDNS Updates to a DNS Server in the Same Domain](#) on page 711.
 - For information about sending secure DDNS updates to a DNS server in a different domain, see [Sending Secure DDNS Updates to a DNS Server in Another Domain](#) on page 719
- A NIOS appliance serving DNS can accept GSS-TSIG authenticated DDNS updates from DHCP clients and servers in an AD domain whose domain controller is running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2.
 - For information, see [Accepting GSS-TSIG-Authenticated Updates](#) on page 725.

Note: A NIOS appliance cannot support both of these features at the same time.

Kerberos Authentication for GSS-TSIG

The keytab file contains only Kerberos keys for principals. It is possible to infer the KDC (Key Distribution Center) from the principal because Windows uses uppercase AD domain names for Kerberos realm names. You must provide the principal name. The principal name may contain Kerberos realm, and the DNS servers for the domain are available for DNS name resolution. Therefore, resolving `SRV _kerberos._tcp.REALM.` will return the appropriate KDC. New TGTs cannot be acquired when the KDC that issues the TGT fails. If the appliance has successfully authenticated before the KDC failure, the secure updates will continue until the session key and TGT expire. The default expiration on Windows is 10 hours. If the appliance restarts or reboots, secure updates are deferred until the KDC becomes available.

The following provides information about the traffic flow between the appliance and the KDC:

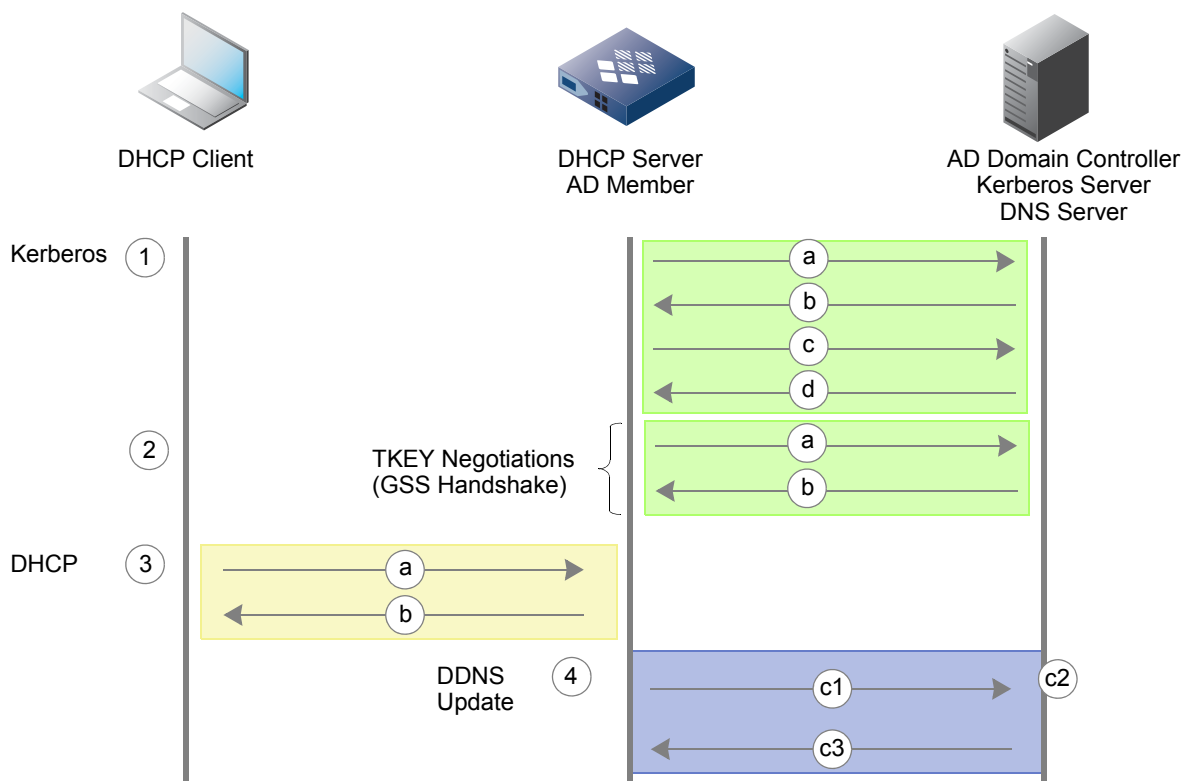
- Client uses keytab to get TGT for principal from KDC (AS-REQ/AS-REP).
- Client uses TGT to get session ticket from KDC (TGS-REQ/TGS-REP).
- Client uses session ticket to acquire TKEY from DNS server (TKEY/TKEY).
- Client uses TKEY to sign DNS updates (DNS-TSIG/DNS-TSIG).

The DNS server authenticates into the domain when the keytab file is generated on the KDC and its SPN is mapped to an account. The server’s private key is known to itself and to the KDC. The KDC generates the ticket and the DNS server allows the update.

Sending Secure DDNS Updates to a DNS Server in the Same Domain

An Infoblox DHCP server can send GSS-TSIG authenticated DDNS updates to a DNS server in an AD domain whose domain controller is running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2. The DHCP server, DNS server, and domain controller are all in the same AD domain. The process by which an Infoblox DHCP server dynamically updates resource records on a DNS server using GSS-TSIG authentication is shown in [Figure 20.6](#). In the illustration, the Kerberos Key Distribution Center (KDC) is running on an AD domain controller, which also provides DNS service.

Figure 20.6 An Infoblox DHCP Server Sends GSS-TSIG Updates to a DNS Server



After you enable the NIOS appliance to send GSS-TSIG authenticated updates to a DNS server, the following process occurs:

1. Kerberos – Login, and TGT and Service Ticket Assignments
 - a. The Infoblox appliance automatically logs in to the AD/Kerberos server.
 - b. The Kerberos server sends the appliance a TGT (ticket-granting ticket).
 - c. Using the TGT, the appliance requests a service ticket for the DNS server.
 - d. The Kerberos server replies with a service ticket for that server.
2. TKEY negotiations (GSS Handshake):
 - a. The appliance sends the DNS server a TKEY (transaction key) request. A Transaction Key record establishes shared secret keys for use with the TSIG resource record. For more information, see *RFC 2930, Secret Key Establishment for DNS (TKEY RR)*.
The request includes the service ticket. The service ticket includes the appliance's principal and proposed TSIG (transaction signature) key, along with other items such as a ticket lifetime and a timestamp.
 - b. The DNS server responds with a DNS server-signed TSIG, which is a "meta-record" that is never cached and never appears in zone data. A TSIG record is a signature of the update using an HMAC-MD5 hash that provides transaction-level authentication. For more information, see *RFC 2845, Secret Key Transaction Authentication for DNS (TSIG)*.

The two participants have established a security context.

When a DHCP client sends a request for an IP address to the DHCP server, the following occurs:

1. DHCP – IP Address and Network Parameters Assignment
 - a. The DHCP client requests an IP address.
 - b. The DHCP server assigns an IP address, subnet mask, gateway address, DNS server address, and a domain name.

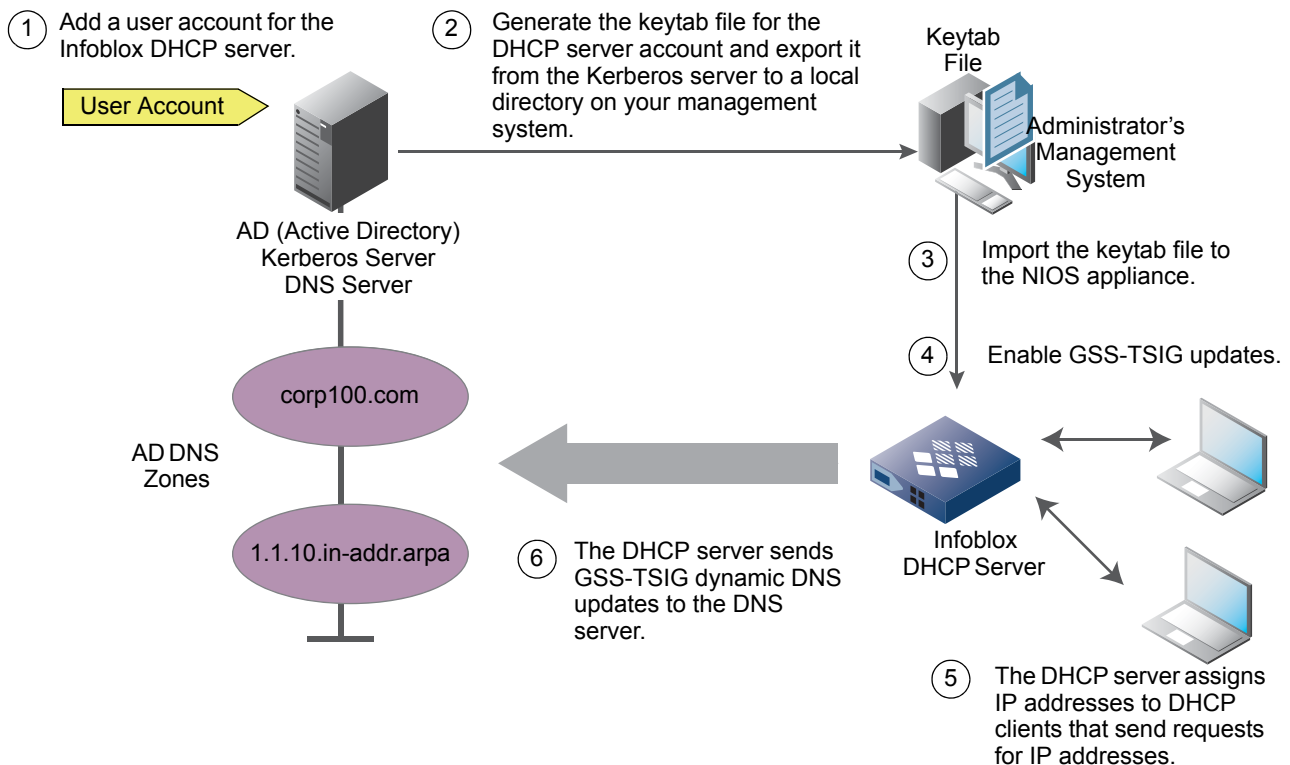
After the appliance assigns an IP address to the DHCP client, it sends the DDNS update to the DNS server as follows:

1. DDNS – Dynamic Update of the Client's Resource Records
 - c. GSS-TSIG-Authenticated DDNS Update
 1. The appliance sends an authenticated DDNS update, which may include the following resource records:
 - A or AAAA – Address record
 - or
 - PTR – Pointer record
 - TKEY – Transaction Key record
 - TSIG – TSIG record
 2. The DNS server verifies the DDNS update and allows it to complete.
 3. The DNS server sends a GSS-TSIG-authenticated response to the appliance, confirming the update.

Configuring DHCP to Send GSS-TSIG Updates in the Same Domain

Before configuring an Infoblox DHCP server to support GSS-TSIG, you must create a user account on the Kerberos server for the appliance. Then you must export the corresponding keytab file from the Kerberos server and import it onto the NIOS appliance. [Figure 20.7](#) illustrates the initial configuration tasks.

Figure 20.7 Adding an Infoblox DHCP Server to an AD Environment with GSS-TSIG Support



The Infoblox DHCP server can send GSS-TSIG-signed DDNS updates to a DNS server for one domain only, though multiple Infoblox DHCP servers can update that domain. If you want more than one Infoblox DHCP server to update a DNS domain, you can either import the same keytab file to the other Infoblox DHCP servers or generate and import a different keytab file. In a Grid, each member can update a different domain.

Note: For GSS-TSIG authentication to work properly, the system clock times of the Infoblox DHCP server, AD domain controller and DNS server must be synchronized. One approach is to use NTP and synchronize all three devices with the same NTP servers.

To use an AD domain controller as a Kerberos Key Distribution Center, complete the following tasks on an AD/Kerberos server:

1. Add a user account for the NIOS appliance to the AD domain controller. For information, see [Creating an AD User Account](#) on page 713.
2. Generate the keytab file for the NIOS appliance account and export it from the AD domain controller to a local directory on your management system. For information, see [Generating and Exporting the Keytab File](#) on page 713.

To configure a NIOS appliance to support AD and send GSS-TSIG secure DDNS updates to a DNS server, complete the following tasks on a NIOS appliance:

1. Import the keytab file from your management system to the appliance and enable GSS-TSIG dynamic updates at the Grid or member level. For information, see [Enabling GSS-TSIG Authentication for DHCP](#) on page 718.
2. Configure the appliance to send GSS-TSIG dynamic updates to forward-mapping and optionally, reverse-mapping zones on the DNS server. For information, see [Creating an External Zone for GSS-TSIG Updates](#) on page 718.

Creating an AD User Account

Connect to the AD domain controller and create a user account for the NIOS appliance.

Note: The name that you enter in the User logon name is the name that you later use when exporting the keytab file. This is also the principal name. The text in the First name, Initials, Last name, and Full name fields is irrelevant to this task.

The AD domain controller automatically creates a Kerberos account for this user. Note the following:

- If you define an expiration date for the user account and you later create a new account when the first one expires, the keytab for the corresponding Kerberos account changes. At that point, you must update the keytab file on the NIOS appliance (see [Generating and Exporting the Keytab File](#) and [Enabling GSS-TSIG Authentication for DHCP](#) on page 718). Optionally, if your security policy allows it, you can set the user account for the NIOS appliance so that it never expires.
- If the AD domain controller is running Windows Server 2003, the user account must have the DES encryption type enabled. You can enable this either in the Account tab of the AD domain controller when you create the user account or by specifying **+DesOnly** when you use the Ktpass tool to generate the keytab file. For instructions, see the next section, [Generating and Exporting the Keytab File](#).

Generating and Exporting the Keytab File

You can use the Ktpass tool to generate and export the keytab file for the Kerberos account. Note that the version of the Ktpass tool that you use must match the Windows version of the domain controller. For example, if you are using a domain controller running Windows Server 2008 or Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2, you must use the Ktpass tool for Windows Server 2008, Windows Server 2008 R2, Windows Server 2012 or Windows Server 2012 R2.

You enter different commands for generating and exporting the keytab file, depending on whether you are generating the keytab file from a server running Microsoft Windows 2000, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2010, Windows Server 2012, or Windows Server 2012 R2.

A Windows Server 2003 domain controller allows you to generate a keytab file with only one key for a principal. A Windows Server 2008, Windows Server 2008 R2 domain controller, Windows Server 2012 or Windows Server 2012 R2 allows you to generate a keytab file with multiple keys for one principal. This is useful when the KDC has principals with multiple encryption types. When the NIOS DHCP server uses a keytab with multiple keys, it negotiates a key based on those in the configured keytab file.

Note: The keytab file contains highly sensitive data for the NIOS appliance account. Ensure that you store and transport its contents securely.

Infoblox strongly recommends the following encryption types for compatibility purposes:

Microsoft Windows Server	Export keytab
Microsoft Windows 2000	Specify <code>/crypto DES-CBC-MD5</code> as the export keytab.
Microsoft Windows 2003	Specify <code>/crypto RC4-HMAC-NT</code> as the export keytab. <ul style="list-style-type: none"> You can also use <code>AES</code>, but <code>RC4</code> is set by default for Windows 2003 servers. Infoblox recommends that you do not use <code>DES</code>, though it is supported if you need it for compatibility with non-Windows systems.
Microsoft Windows 2008 and higher	Specify <code>/crypto all</code> as the export keytab. <ul style="list-style-type: none"> You can also use <code>AES</code> or <code>RC4</code>. Note that <code>DES</code> is not enabled by default on Windows Server 2008 and higher. You must enable <code>DES</code> in order to include it in the <code>crypto all</code> keytab. Infoblox recommends that you include <code>DES</code> for maximum compatibility, though you can still select non-<code>DES</code> crypto. Infoblox recommends that you do not use only <code>DES</code> as the export keytab.

Generating the Keytab on Windows 2000 Server

To export the keytab file using a Microsoft Windows 2000 Resource Kit:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:
`C:> ktpass -princ service_name/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab`

Note: The values are case-sensitive.

where:

- `service_name/instance`: The AD user name for the NIOS appliance and a character string. The AD user name must match the user logon name on the AD domain controller.
- `REALM`: The Kerberos realm in uppercase. It must match the realm (or domain name) specified in the `-mapuser` option.

For example:

```
C:> ktpass -princ DNS/ns1.corp100.com@CORP100.COM -mapuser ns1@corp100.com -pass 37Le37 -out ns1.keytab
```

Generating the Keytab on Windows Server 2003

The Ktpass tool is included in the Windows Server 2003 Support Tools. To export the keytab file using a Microsoft Windows 2003 Resource Kit:

1. Start a command prompt.
2. Enter the following command to generate the keytab file for the NIOS appliance user account:

```
ktpass -princ service_name/FQDN_instance@REALM -mapuser AD_username@REALM -pass password -out filename.ktb -ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

Note: The values are case-sensitive.

Example:

```
ktpass -princ DNS/ns1.corp100.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass 37Le37 -out ns1.ktb -ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

where

-princ = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: `DNS/ns1.corp100.com@GSS.LOCAL`

- DNS = Service name in uppercase format.
- ns1.corp100.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase format. This must be the same as the AD domain name.

-mapuser = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use `ksetup` without any parameters or arguments to see the current mapped settings and the default realm. Example: `ksetup /mapuser <Principal> <Account>`. To create an AD user account, see [Creating an AD User Account](#) on page 713.

- jsmith = The AD user name for the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the **-princ** option.

-pass = The AD user account password. The Ktpass command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- 37Le37 = The password of the user account for the NIOS appliance.

-out = The name of the keytab file that is generated.

- ns1.ktb = The name of the keytab file

-ptype = Sets the principal type. This must be **krb5_nt_principal**.

-crypto = Specifies the encryption type. This must be **des-cbc-md5**. DES-CBC-MD5 adheres to the MIT implementation and is used for compatibility.

+DesOnly = Specifies DES encryption for the account. This is set by default on the Windows Server 2008. Include this if you did not enable DES encryption for the account. Note that Windows 7 and Windows Server 2008 R2 do not support DES by default.

Note: Note that the Windows Server 2003 does not support AES encryption.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: ibtest-xu5nxd56.corp100.local
Using legacy password setting method
Successfully mapped dns/anywhere to dns.
Key created.
Output keytab to dns.ktb:
Keytab version: 0x502
keysize 56 dns/anywhere@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 5 etype
0x3 (DES-CBC-MD5) keylength 8 (0xbae610f11552c80b)
```

Generating the Keytab on Windows Server 2008 or Windows Server 2008 R2

On a Windows Server 2008 or Windows Server 2008 R2 domain controller, the Ktpass tool supports generating a keytab file with multiple keys for a single principal. To generate the keytab file using the Ktpass tool:

1. Start a command prompt.
2. Enter the following command to generate the keytab file for the NIOS appliance user account:
ktpass -princ *username@REALM* -mapuser *logon_name@REALM* -pass *password* -out *my.tab* -ptype krb5_nt_principal -crypto *encryption*

Example:

```
ktpass -princ DNS/ns1.corp100.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass 37Le37 -out ns1.keytab -ptype krb5_nt_principal -crypto RC4-HMAC-NT
```

where:

-princ = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: *DNS/ns1.corp100.com@GSS.LOCAL*

- *DNS* = Service name in uppercase format.
- *ns1.corp100.com* = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- *GSS.LOCAL* = The Kerberos realm in uppercase format. This must be the same as the AD domain name.

-mapuser = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use *ksetup* without any parameters or arguments to see the current mapped settings and the default realm. Example: *ksetup /mapuser <Principal> <Account>*. To create an AD user account, see [Creating an AD User Account](#) on page 713.

- *jsmith* = The AD user name for the NIOS appliance.
- *GSS.LOCAL* = The Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the **-princ** option.

-pass = The AD user account password. The Ktpass command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- *37Le37* = The password of the user account for the NIOS appliance.

-out = The name of the keytab file that is generated.

- *ns1.ktb* = The name of the keytab file

-ptype = Sets the principal type. This must be **krb5_nt_principal**.

-crypto = Specifies the encryption type. Note that the **RC4-HMAC-NT** encryption type is enabled by default. You can also use the following:

-crypto: all or **-crypto RC4-HMAC-NT**

You can optionally specify the following:

+DesOnly = Specifies DES encryption for the account. You must use this only when you use DES-CBC-MD5 for compatibility. Note that Windows 7 and Windows Server 2008 R2 do not support DES by default. However, you can enable DES on the Windows 2008 server. Include this option if you did not enable DES encryption for the account. For more information, refer to the information available in a third-party portal at:

<http://weblogic-wonders.com/weblogic/2010/11/30/windows-7-des-encryption-support-for-kerberos-authentication/>

Note: You must not use **+Desonly** with **/crypto all** or other non-DES encryption types.

+setpass = Sets a new AD user account password. This is required if the **+DesOnly** option is specified. When you use this encryption type, you must change the user's password. Otherwise, the ticket issued for the principal becomes unusable.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: qacert.test.local
Using legacy password setting method
Successfully mapped DNS/ns1.corp100.com to ns1.
```

```
Key created.
Output keytab to ns1.ktb:
Keytab version: 0x502
keysize 80 DNS/ns1.corp100.com@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype 0x12
(AES256-SHA1)
keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)
```

Generating the Keytab on Windows Server 2012 or Windows Server 2012 R2

On a Windows Server 2012 or Windows Server 2012 R2 domain controller, the Ktpass tool supports generating a keytab file with multiple keys for a single principal. Infoblox recommends that one of the keys include the encryption type DES-CBC-MD5 for compatibility purposes. Note that you must enable DES and then export using the `/crypto all` keytab. Most principals use DES-CBC-MD5 and it is the most compatible encryption type with other systems, such as MIT Kerberos.

To generate the keytab file using the Ktpass tool:

1. Start a command prompt.
2. Enter the following command to generate the keytab file for the NIOS appliance user account:

```
ktpass -princ username@REALM -mapuser login_name@REALM -pass password -out my.tab -ptype  
krb5_nt_principal -crypto encryption
```

Example:

```
ktpass -princ DNS/ns1.corp100.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass 37Le37 -out  
ns1.keytab -ptype krb5_nt_principal -crypto all
```

where:

-princ = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: `DNS/ns1.corp100.com@GSS.LOCAL`

- DNS = Service name in uppercase format.
- ns1.corp100.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase format. This must be the same as the AD domain name.

-mapuser = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use `ksetup` without any parameters or arguments to see the current mapped settings and the default realm. Example: `ksetup /mapuser <Principal> <Account>`. To create an AD user account, see [Creating an AD User Account](#) on page 713.

- jsmith = The AD user name for the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the **-princ** option.

-pass = The AD user account password. The Ktpass command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- 37Le37 = The password of the user account for the NIOS appliance.

-out = The name of the keytab file that is generated.

- ns1.ktb = The name of the keytab file

-ptype = Sets the principal type. This must be **krb5_nt_principal**.

-crypto = Specifies the encryption type. You can specify the following encryption types:

- **DES-CBC-CRC** = Specifies DES encryption for the account. This encryption type is used for compatibility.
- **DES-CBC-MD5** = Specifies DES encryption for the account. This encryption type adheres to the MIT implementation and is used for compatibility.
- **RC4-HMAC-NT** = Specifies 128-bit RC4-HMAC encryption for the account. This is enabled by default.
- **AES256-SHA1** = Specifies 256-bit AES encryption for the account.
- **AES128-SHA1** = Specifies 128-bit AES encryption for the account.
- **ALL** = Specifies all of the above encryption types. Infoblox recommends that you use the `/crypto all` encryption type.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: qacert.test.local
Using legacy password setting method
Successfully mapped DNS/ns1.corp100.com to ns1.
Key created.
Output keytab to ns1.keytab:
Keytab version: 0x502
keysize 80 DNS/ns1.corp100.com@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype 0x12
(AES256-SHA1)
keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)
```

Enabling GSS-TSIG Authentication for DHCP

You can enable GSS-TSIG authentication at the Grid or member level. When you enable GSS-TSIG authentication, make sure that you upload the keytab file from the Kerberos account for the Infoblox DHCP server.

The AD domain controller stores the keytab file in the directory in which you generated the keytab file. You can copy this file to a management system that connects to the NIOS appliance or launch the NIOS Grid Manager on the AD domain controller and import the keytab file to the NIOS appliance.

You can import keytab files to the Grid or to individual members.

To enable GSS-TSIG authentication and import Keytab files:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
Member: From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* check box -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the **IPv4 DDNS** -> **Basic** tab or the **IPv6 DDNS** -> **Basic** tab of the editor, complete the following:
 - **DDNS Updates:** Select the **Enable DDNS Updates** check box.
 - **GSS-TSIG:** Complete the following:
 - **Enable GSS-TSIG Updates:** Select this check box.
 - **Domain Controller:** Enter the resolvable host name or IP address of the AD domain controller that hosts the Key Distribution Center (KDC) for the domain.
 - **GSS-TSIG Key:** Select the name of the keytab file you are using for the Grid. This is only available if you have uploaded a keytab file.
To upload a keytab file, click **Manage Keytab Files**. In the *Keytab File Manager* dialog box, click the Add icon. Click **Browse**, navigate to the keytab file, select it, and then click **Upload**.
 - **Domain:** The appliance displays the name of the domain associated with the keytab file.
 - Click **Display** to list the external zones to which the Grid member can send secure DDNS updates.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Each time you export a keytab file from a Kerberos server running on Windows Server 2003, the version number of the keytab file increases incrementally. Because the version number of the keytab file that you import to the NIOS appliance must match the version that is in use on the Kerberos server, you should select the last keytab file that is exported from the Kerberos server if you have exported multiple keytab files.

Creating an External Zone for GSS-TSIG Updates

For each network view, you specify the zone to be updated, the IP address of the primary DNS server for that zone, and the security method, GSS-TSIG. The zone must be in the same AD domain as the member that is sending the updates.

You can add information for a forward and reverse zone. The DHCP server updates the A record in the forward zone and the PTR record in the reverse zone.

To enable the NIOS appliance to send dynamic updates to a DNS server using GSS-TSIG for authentication:

1. If there are multiple network views in the Grid, select a network view.

2. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Configure DDNS**.
3. In the DDNS Updates to External Zones table of the *DDNS Properties* editor, click the Add icon and complete the following fields in the Add External DDNS Zone panel:
 - **Zone Name:** Enter the name of the zone that receives the updates. You can specify both forward-mapping and reverse-mapping zones.
 - **DNS Server Address:** Enter the IP address of the primary name server for that zone.
 - **Security:** Select **GSS-TSIG**.
 - **AD Domain:** Select the AD domain associated with the keytab file.
 - **DNS Principal:** The name and domain of the DNS server receiving the DDNS updates. Note that this is not the same as the Kerberos principal you specified when you generated the keytab file.
Use the following format when you complete this field: **DNS/dns_server_fqdn@ad_domain**
dns_server_fqdn: This is the FQDN of the DNS server. You can use the “dig” command to perform a DNS lookup to obtain the FQDN of the DNS server as it appears on the SOA record.
ad_domain: This is the AD domain of the DNS server.
 - Click **Test GSS-TSIG** to list the Grid members that are allowed to send GSS-TSIG updates to the DNS server.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Verifying the Configuration

After you configure the AD domain controller and the Infoblox DHCP server, you can view the syslog of the Infoblox DHCP server to verify if it successfully established a security context with the AD domain controller. The DHCP server displays a series of messages similar to the following:

```
dhcpcd: Enabled GSS-TSIG for zone corp100. using principal jdoe/anywhere@CORP100.LOCAL.
dhcpcd: GSS-TSIG security thread has started.
dhcpcd: GSS-TSIG security update starting at 1222389338.
dhcpcd: Acquiring GSS-TSIG credential for jdoe/anywhere@CORP100.LOCAL.
dhcpcd: Acquired GSS-TSIG credential for jdoe/anywhere@CORP100.LOCAL(good for 3568s).
dhcpcd: Security context established with server 10.34.123.4 for principal
jdoe/anywhere@CORP100.LOCAL (good for 568s).
dhcpcd: GSS-TSIG security update complete at 1222389338. Next update in 360s.
```

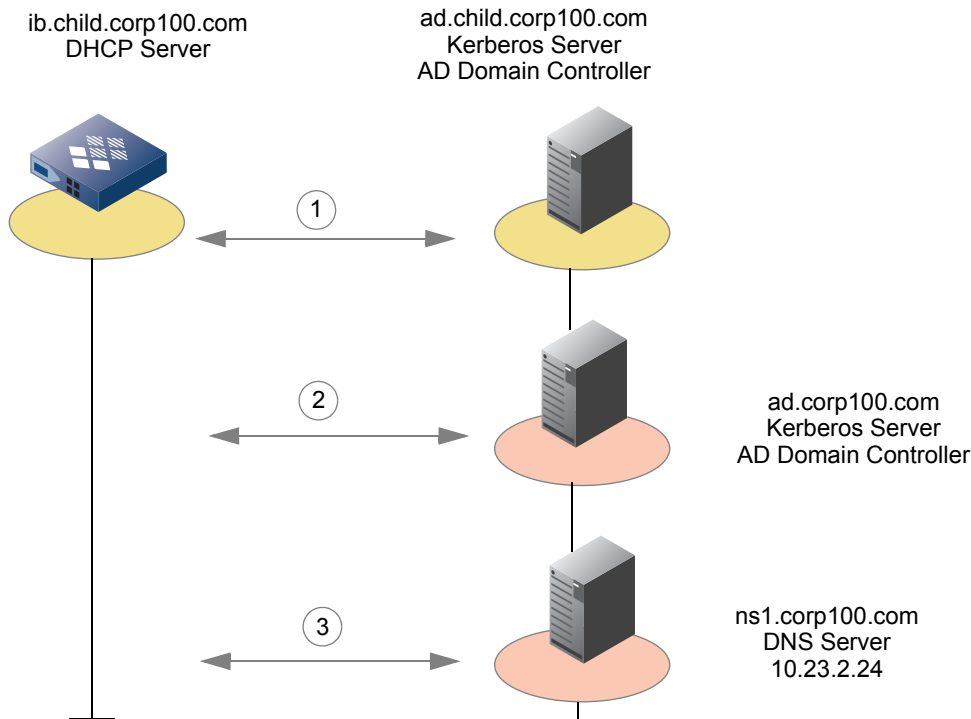
In addition, you can log in to the Infoblox CLI and use the `show dhcp_gss_tsig` CLI command to troubleshoot your configuration. For information about this command, refer to the *Infoblox CLI Guide*.

Sending Secure DDNS Updates to a DNS Server in Another Domain

Domain and forest trust relationships provide clients authenticated access to resources in other domains. Some trusts are automatically created, such as the two-way, direct trust between parent and child domains in a forest. Other trusts must be created manually. Refer to the Microsoft Active Directory documentation for information on establishing trusts between domains.

Once a direct trust exists between two AD domains, a KDC from one domain can grant a referral to the KDC of the other domain. The Infoblox DHCP server can then use the referral to request access to services in the other domain.

In [Figure 20.8](#), the Infoblox DHCP server in the child.corp100.com domain needs to send GSS-TSIG authenticated DDNS updates to the DNS server in its parent domain, corp100.com domain. There is an automatic two-way trust between the domains because corp100.com domain is the parent of child.corp100.com domain.

Figure 20.8 Sending Secure DDNS Updates to a DNS Server in Another Domain

After you configure the Infoblox DHCP server and AD domain controller, the following occurs:

1. Kerberos – In Same Domain

The Infoblox DHCP server uses the TGT (ticket-granting ticket) from the AD/Kerberos server, `ad.child.corp100.com`, to request a service ticket for `DNS/ns1.corp100.com@CORP100.COM`. The Kerberos server replies with a referral ticket for the Kerberos server in the `corp100.com` domain, `ad.corp100.com`.

2. Kerberos – In the Other Domain

The Infoblox DHCP server uses the referral ticket and requests a service ticket from `ad.corp100.com` for `DNS/ns1.corp100.com@CORP100.COM`. The Kerberos server replies with a service ticket for `DNS/ns1.corp100.com@CORP100.COM`.

3. TKEY Negotiations (GSS Handshake)

The Infoblox DHCP server sends the DNS server `ns1.corp100.com` a TKEY (transaction key) request, which includes the service ticket. The DNS server replies with a TKEY response that includes a TSIG (transaction signature). The Infoblox appliance and the DNS server have established a security context, enabling the DHCP server to send DDNS updates to the DNS server.

Configuring DHCP to Send GSS-TSIG Updates to Another Domain

Before the DHCP server can send secure DDNS updates to a DNS server in a different domain, you must ensure that a direct trust relationship exists between the domain of the DHCP server and that of the DNS server. (For information, refer to the Active Directory documentation.)

Following are the tasks to configure the AD domain controller and the Infoblox DHCP server for secure updates to another domain. All the configuration is done on the AD domain controller for the domain of the DHCP server and on the Infoblox DHCP server.:

1. Complete the following tasks on the AD domain controller for the domain of the DHCP server:

- a. Add a user account for the Infoblox DHCP server. In the configuration example, the user account is `ibdhcp`. For information, see [Creating an AD User Account](#) on page 713.
 - b. Generate the keytab file for the Infoblox DHCP server and export it from the AD domain controller to a local directory on your management system. For the DHCP server in [Figure 20.8](#), the principal is `ibdhcp/ib.child.corp100.com@CHILD.CORP100.COM`. For information, see [Generating and Exporting the Keytab File](#) on page 713.
2. Complete the following tasks on the Infoblox DHCP server:
 - a. Import the keytab file from your management system to the appliance and enable GSS-TSIG dynamic updates at the Grid or member level. For information, see [Enabling GSS-TSIG Authentication for DHCP](#) on page 718
 - b. Configure the external forward-mapping zone for the DDNS updates. Note that the DNS principal uses the domain of the DNS server, regardless of the domain of the DHCP server. For the DNS server in [Figure 20.8](#), the DNS principal is `DNS/ns1.corp100.com@CORP100.COM`. For information, see [Creating an External Zone for GSS-TSIG Updates](#) on page 718.

Configuration Example

Following are the steps to configure the example shown in [Figure 20.8](#):

On the AD domain controller:

1. Create a user account for the Infoblox DHCP server. The user account is `ibdhcp`.
2. Generate the keytab file and export it to your management system. If the domain controller is running Windows Server 2003:

```
ktpass -princ ibdhcp/ib.child.corp100.com@CHILD.CORP100.COM -mapuser
ibdhcp@CHILD.CORP100.COM -pass infoblox -out ibdhcp.ktb -ptype krb5_nt_principal -crypto
des-cbc-md5 +desonly
```

On the Infoblox DHCP server:

1. Enable GSS-TSIG at the member level.
2. From the **DHCP** tab, click the **Members** tab -> *member* check box -> Edit icon.
3. In the **DDNS** -> **Basic** tab of the editor, complete the following:
 - **Override**: Select this check box.
 - **DDNS Updates**: Select the **Enable DDNS Updates** check box.
 - **GSS-TSIG**: Select **Override** and complete the following:
 - **Enable GSS-TSIG Updates**: Select this check box.
 - **Domain Controller**: Enter `ad.child.corp100.com`. This is the KDC in the domain of the DHCP server.
 - **GSS-TSIG Key**: Click **Manage Keytab Files**. In the *Keytab File Manager* dialog box, click the Add icon. Click **Browse**, navigate to the keytab file, select it, and then click **Upload**.
Select the keytab file that you just uploaded, `ibdhcp/ib.child.corp100.com@CHILD.CORP100.COM`.
 - **Domain**: The appliance displays the name of the domain associated with the key, which is `child.corp100.com`.
 - Click **Test GSS-TSIG** to list the external zones to which the Grid member can send secure DDNS updates.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
5. Configure the external forward mapping zone, `corp100.com`.
 - a. From the **DHCP** tab, expand the Toolbar and click **Configure DDNS**.
 - b. In the DNS Updates to External Zones table of the *DDNS Properties* editor, click the Add icon and complete the following fields in the Add External DDNS Zone panel:
 - **Zone Name**: Enter `corp100.com`.
 - **DNS Server Address**: Enter the IP address of the primary DNS server to which the Infoblox DHCP server sends DDNS updates. In the example, the DNS server is `ns.corp100.com`. Therefore, enter its IP address, which is `10.23.2.24`.

- **Security:** Select **GSS-TSIG**.
- **AD Domain:** Select **child.corp100.com**.
- **DNS Principal:** Enter **DNS/ns1.corp100.com@CORP100.COM**.
- Click **Test GSS-TSIG** to list the Grid members that are allowed to send GSS-TSIG updates to the DNS server.

6. Save the configuration and click **Restart** if it appears at the top of the screen.

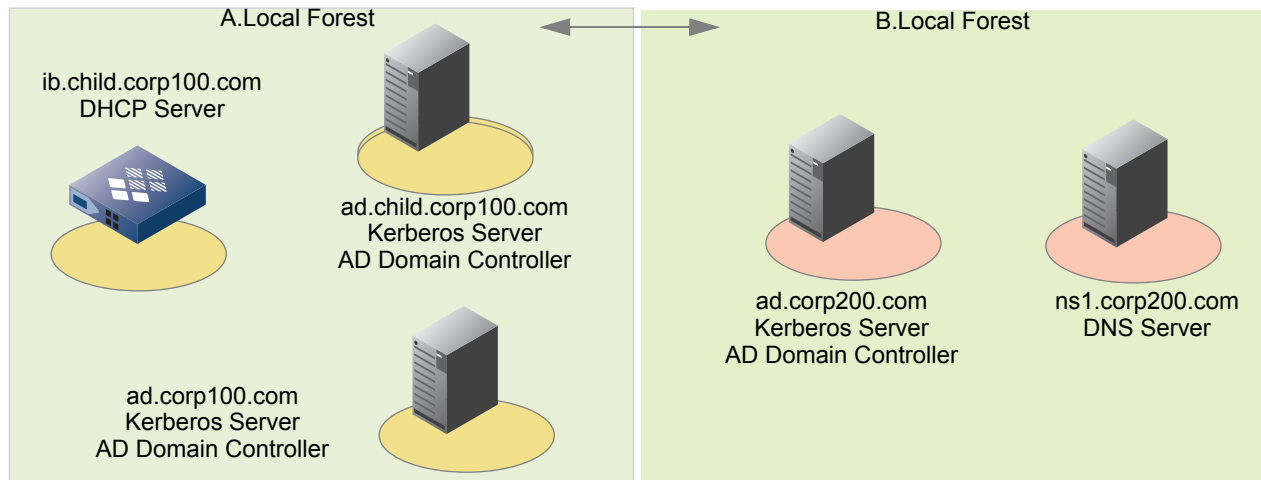
Sending GSS-TSIG Updates to a DNS Server in Another Forest

The Infoblox DHCP server can also send secure DDNS updates to a DNS server that belongs to a domain in another forest, as long as a forest trust exists. Refer to the Microsoft Active Directory documentation for information on establishing forest trusts.

Similar to the authentication process between domains, the authentication process between forests also uses referrals. The appliance follows the referral chain until it reaches the domain controller of the domain in which the service is located. Note that forest trusts are not transitive. For example, if the DHCP server is in forest A and the DNS server is in forest C, a direct trust must exist between forest A and forest C for the DDNS updates to succeed. Having a trust between forest A and B, and between forest B and C is not sufficient.

In [Figure 20.9](#), a trust exists between the A.Local forest and the B.Local forest. The Infoblox DHCP server in the A.Local forest needs to dynamically update the DNS server in the B.Local forest.

Figure 20.9 Sending Secure DDNS Updates to a DNS Server in Another Forest



The following authentication process occurs:

1. Kerberos – In Same Domain

The Infoblox appliance uses the TGT (ticket-granting ticket) from the AD/Kerberos server, ad.child.corp100.com, to request a service ticket for DNS/ns1.corp200.com@CORP200.COM. The Kerberos server does not find the principal name in its domain database and after consulting the global catalog, it replies with a referral ticket for its parent domain.

2. Kerberos – Referral Chain

The appliance contacts a domain controller in corp100.com and requests a referral to a domain controller in the corp200.com domain in B.Local Forest.

When it receives the referral, the DHCP server contacts the domain controller and requests a service ticket for the DNS server, ns1.corp200.com. The domain controller replies with a service ticket for DNS/ns1.corp200.com@CORP200.COM.

3. TKEY Negotiations (GSS Handshake)

The Infoblox appliance sends the DNS server ns1.corp200.com a TKEY (transaction key) request, which includes the service ticket. The DNS server replies with a TKEY response that includes a TSIG (transaction signature). The Infoblox appliance and the DNS server have established a security context.

Configuring DHCP to Send GSS-TSIG Updates to a Different Forest

Configuring the Infoblox DHCP server for dynamic updates to a DNS server in another forest is similar to the configuration used to send dynamic updates to another domain in the same forest. For information, see [Configuring DHCP to Send GSS-TSIG Updates to Another Domain](#) on page 720.

ACCEPTING DDNS UPDATES FROM DHCP CLIENTS

A NIOS appliance serving DNS can support Active Directory and accept both unauthenticated and GSS-TSIG authenticated updates from DHCP clients, DHCP servers, and AD domain controllers. The appliance supports servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows 2012 R2 with the Active Directory service installed.

When adding a NIOS appliance that serves DNS to an AD environment, you must configure the AD/Kerberos server and NIOS appliance as follows—based on whether or not you want the DNS server to support DDNS updates using GSS-TSIG authentication:

- AD/Kerberos Server
 1. Enable zone transfers to the NIOS appliance.
 2. (For GSS-TSIG) Create a user account for the NIOS appliance that it can use for authentication.
 3. (For GSS-TSIG) Generate the keytab file of the DNS server and save it to your management system.
- NIOS Appliance
 4. (GSS-TSIG) Enable GSS-TSIG support.
 5. (GSS-TSIG) Import the keytab file of the DNS server from your management system to the NIOS appliance.
 6. (GSS-TSIG) Enable GSS-TSIG authentication.
 7. Add a forward-mapping zone and give it a name matching the AD DNS zone whose resource records you want to import.
 8. Specify the domain controller from which the appliance can receive DDNS updates. An AD domain controller replicates its data among other domain controllers within its AD domain and among domain controllers in other domains.
 9. Import zone data from the specified domain controller.
 10. Enable the acceptance of DDNS updates from the AD domain controller and from the DHCP clients and servers whose addresses the DHCP server assigns. You can set this at the Grid, member, and zone levels.
 11. (For GSS-TSIG) Enable acceptance of GSS-TSIG DDNS updates from the AD domain controller and from the addresses that the DHCP server assigns. You can set this at the Grid, member, and zone levels.

As you can see from the above task list, adding a NIOS appliance that serves DNS to an AD environment without GSS-TSIG support involves four simple steps. To include GSS-TSIG support, there are several additional steps.

Supporting Active Directory and Unauthenticated DDNS Updates

Before configuring the NIOS appliance, configure the AD domain controller to permit zone transfers to the IP address of the appliance. Then on the appliance, you can do the following to configure a forward-mapping zone to support AD (Active Directory) and receive unauthenticated DDNS updates from DHCP clients, DHCP servers, and AD domain controllers.

- Create a forward-mapping zone, as described in [Creating an Authoritative Forward-Mapping Zone](#) on page 617. Give it a name that matches the AD DNS zone whose resource records you want to import.

- Specify the domain controllers from which the appliance can receive updates, as described in [Configuring AD Support](#) on page 724
- Import the zone data from the domain controller. For information, see [Importing Data into Zones](#) on page 632.
- Enable the appliance to accept DDNS updates from the DHCP clients and servers whose addresses the DHCP server assigns. You can set this at the Grid, member, and zone levels. For information, see [Enabling DNS Servers to Accept DDNS Updates](#) on page 706.

Configuring AD Support

You can configure a forward-mapping zone to support AD from the *Active Directory* wizard or from the **Active Directory** tab of the *Authoritative Zone* editor. This section describes both methods.

To configure AD support using the *Active Directory* wizard:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Configure Active Directory**. Note that from the **Zones** tab, you must select a zone before you click **Configure Active Directory**.

2. In the *Active Directory* wizard, complete the following, and then click **Next**:

- **Select Zone:** Click this and select a zone. The name of the zone must match the name in the AD domain controller so the zone transfer from the AD domain controller to the NIOS appliance can succeed.
- **Allow unsigned updates from Domain Controllers:** Select this option.

If you configured DNS resolvers in the Grid, the appliance sends DNS queries for the names and addresses of the AD domain's domain controllers. Since the name of the zone that you selected is the same as the AD domain name on the domain controller, the appliance can then send a DNS query for the SRV records attached to the domain name. It also sends a DNS query for the A record of each domain controller to determine its IP address. The query results are listed in the next panel.

3. You can edit the list of domain controllers, if necessary. Click **Next** to proceed to the next step.

- To add a domain controller, click the Add icon and specify the IP address.
- To delete a domain controller from the list, select it and click the Delete icon.

4. Complete the following:

- **Do you want to create underscore zones to hold the records added by the Domain Controllers?**

This option allows the appliance to create the following subzones that the DNS server must have to answer AD-related DNS queries:

```
_msdcs.zone
_sites.zone
_tcp.zone
_udp.zone
domaindnszones.zone
forestdnszones.zone
```

Note that these zones are automatically generated. You cannot edit these zones or import data into them. They cannot be modified, thus providing protection against forged updates.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

To configure AD support using the *Authoritative Zone* editor:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box -> Edit icon.

2. In the *Authoritative Zone* editor, select the **Active Directory** tab and do the following:

- **Allow unsigned updates from these Domain Controllers:** Select this check box and specify the AD domain controllers from which the appliance can receive DDNS updates.
- **Automatically create underscore zones:** Select this check box to automatically create the subzones.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

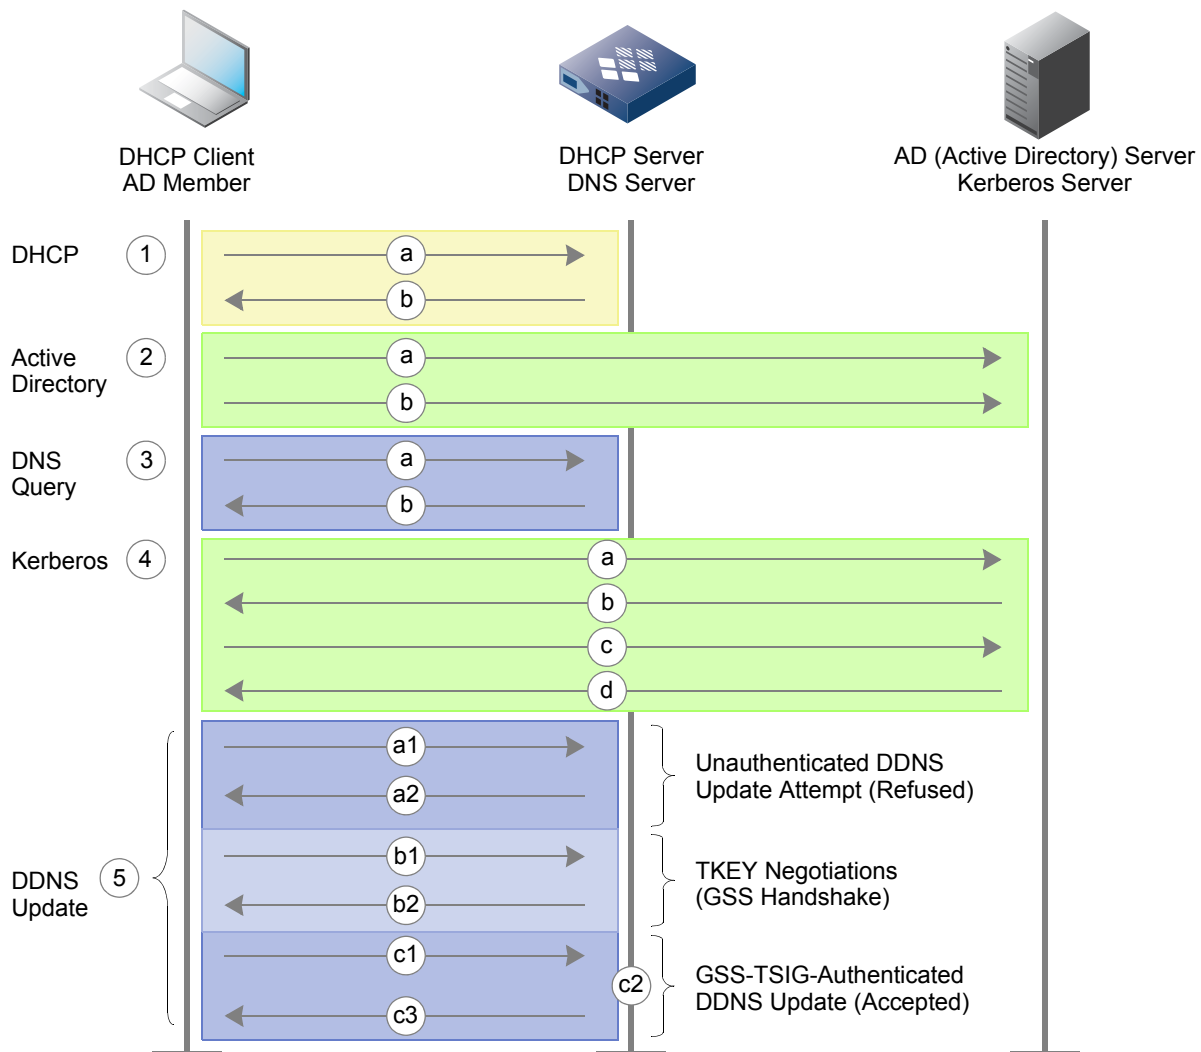
You can then import zone data, as described in [Importing Data into Zones](#) on page 632.

ACCEPTING GSS-TSIG-AUTHENTICATED UPDATES

A NIOS appliance can support Active Directory and process secure GSS-TSIG-authenticated DDNS updates from DHCP clients, DHCP servers, and AD domain controllers. The appliance supports servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 with the Active Directory service installed. The process in which a DHCP client dynamically updates its resource records on a DNS server using GSS-TSIG authentication is shown in [Figure 20.10](#). The illustration also shows the relationship of the clients, the DHCP server, the DNS server, and the Kerberos server (running on the AD domain controller).

Note: For explanations of the alphanumerically notated steps in [Figure 20.10](#), see the section following the illustration.

Figure 20.10 Authenticating DDNS Updates with GSS-TSIG



1. DHCP – IP Address and Network Parameters Assignment
 - a. The DHCP client requests an IP address.
 - b. The DHCP server assigns an IP address, subnet mask, gateway address, and a DNS server address.

2. Active Directory – Computer and User Logins

- a. The computer sends a DNS request to locate the AD domain controller, and then logs in to the domain controller.

Note: Computer accounts have passwords that the AD domain controller and computer maintain automatically. There are two passwords for each computer: a computer account password and a private key password. By default, both passwords are automatically changed every 30 days.

- b. The user manually logs in to a domain.

3. DNS – Query for the Kerberos Server

- a. The computer (or *client*) automatically sends a query for `_kerberos._udp.dc._msdcs.dom_name` to the DNS server whose IP address it received through DHCP.
- b. The NIOS appliance replies with the name of the Kerberos server.

4. Kerberos – Login, and TGT and Service Ticket Assignments

- a. The client automatically logs in to the Kerberos server.
- b. The Kerberos server sends the client a TGT (ticket-granting ticket).
- c. Using the TGT, the AD member requests a service ticket for the DNS server.
- d. The Kerberos server replies with a service ticket for that server.

5. DDNS – Dynamic Update of the Client’s Resource Records

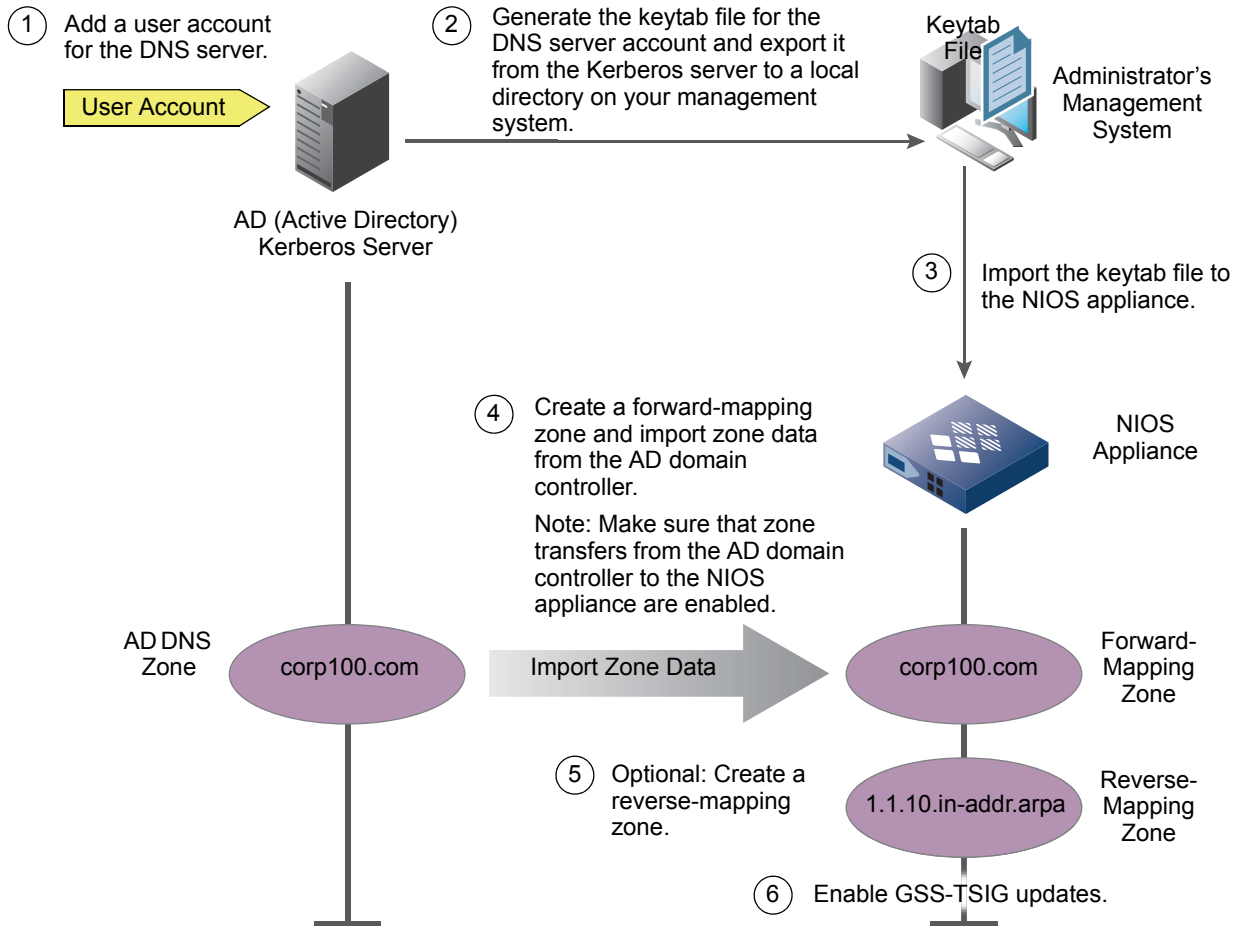
- a. Unauthenticated DDNS Update Attempt (Refused)
 1. The client sends an unauthenticated DDNS update.
 2. The DNS server refuses the update.
- b. TKEY negotiations (GSS Handshake):
 1. The client sends the DNS server a TKEY (transaction key) request. A Transaction Key record establishes shared secret keys for use with the TSIG resource record. For more information, see *RFC 2930, Secret Key Establishment for DNS (TKEY RR)*.
The request includes the service ticket. The service ticket includes the appliance’s principal and proposed TSIG (transaction signature) key, along with other items such as a ticket lifetime and a timestamp.
 2. The DNS server responds with a DNS server-signed TSIG, which is a “meta-record” that is never cached and never appears in zone data. A TSIG record is a signature of the update using an HMAC-MD5 hash that provides transaction-level authentication. For more information, see *RFC 2845, Secret Key Transaction Authentication for DNS (TSIG)*.
The two participants have established a security context.
- c. GSS-TSIG-Authenticated DDNS Update (Accepted)
 1. The client sends an authenticated DDNS update, which includes the following resource records:
 - A – Address record
 - or
 - PTR – Pointer record
 - TKEY – Transaction Key record
 - TSIG – TSIG record
 2. The DNS server authenticates the DDNS update and processes it.
 3. The DNS server sends a GSS-TSIG-authenticated response to the AD member, confirming the update.

Note: For GSS-TSIG authentication to work properly, the system clock times of the Infoblox DHCP server, AD domain controller and DNS server must be synchronized. One approach is to use NTP and synchronize all three devices with the same NTP servers.

Configuring DNS to Receive GSS-TSIG Updates

You can configure an appliance to support Active Directory and accept secure DDNS updates from clients using GSS-TSIG. The initial configuration tasks are shown in [Figure 20.11](#). The appliance supports servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, or Windows Server 2012 R2 with the Active Directory service installed.

Figure 20.11 Adding a NIOS Appliance to an AD Environment with GSS-TSIG Support



On an already functioning AD domain controller:

1. Enable zone transfers to the NIOS appliance.
2. Add a user account for the NIOS appliance serving DNS. A corresponding account on the Kerberos server is automatically created. For information, see [Creating an AD User Account](#) on page 728.
3. Export the keytab file for the NIOS appliance account from the Kerberos server to a local directory on your management system. For information, see [Generating and Exporting the Keytab File](#) on page 728.

On an Infoblox appliance:

1. Import the keytab file from your management system to the Infoblox appliance and enable GSS-TSIG authentication on the appliance. For information, see [Importing the Keytab File and Enabling GSS-TSIG Authentication](#) on page 730.
2. Configure a forward-mapping zone with the same name as the AD zone. For information, see [Creating an Authoritative Forward-Mapping Zone](#) on page 617.

3. (Optional) Create a reverse-mapping zone for the network address space that corresponds to the domain name space in the forward-mapping zone. For information, see [Creating an Authoritative Reverse-Mapping Zone](#) on page 618.
4. Import the zone data from the AD domain controller. For information, see [Importing Zone Data](#) on page 631.
5. Enable the acceptance of GSS-TSIG-signed updates from the AD controller and from the DHCP clients and servers whose addresses the DHCP server assigns. For information, see [Accepting GSS-TSIG Updates](#) on page 730.

Creating an AD User Account

Connect to the AD domain controller and create a user account for the NIOS appliance.

Note: The name you enter in the User logon name is the name that you later use when exporting the keytab file. This is also the principal name. The text in the First name, Initials, Last name, and Full name fields is irrelevant to this task.

The AD domain controller automatically creates a Kerberos account for this user with an accompanying keytab. Note the following:

- If you define an expiration date for the user account and you later create a new account when the first one expires, the keytab for the corresponding Kerberos account changes. At that point, you must update the keytab file on the NIOS appliance (see [Generating and Exporting the Keytab File](#) and [Importing the Keytab File and Enabling GSS-TSIG Authentication](#) on page 730). Optionally, if your security policy allows it, you can set the user account for the NIOS appliance so that it never expires.
- If the AD domain controller is running Windows Server 2003, the user account must have the DES encryption type enabled. You can enable this either in the Account tab when you create the user account or by specifying **+DesOnly** when you use the Ktpass tool to generate the keytab file.

Generating and Exporting the Keytab File

You can generate and export the keytab file for the Kerberos account by using the Ktpass tool. Note that the version of the Ktpass tool that you use must match the Windows version of the domain controller. For example, if you are using a domain controller running Windows Server 2008 or Windows Server 2008 R2, you must use the Ktpass tool for Windows Server 2008 or Windows Server 2008 R2.

You enter different commands for generating and exporting the keytab file, depending on whether you are generating the keytab file from a server running Microsoft Windows 2000, Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2.

Generating the Keytab on Windows 2000

To export the keytab file using a Microsoft Windows 2000 Resource Kit:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:
`C:> ktpass -princ service_name/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab`

For example:

```
C:> ktpass -princ DNS/ns1.corp100.com@CORP100.COM -mapuser ns1@corp100.com -pass 37Le37 -out ns1.keytab
```

Generating the Keytab on Windows Server 2003

The Ktpass tool is included in the Windows Server 2003 Support Tools. To export the keytab file using a Microsoft Windows 2003 Resource Kit:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:


```
ktpass -princ DNS/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab
-ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

For example:

```
ktpass -princ DNS/ns1.corp100.com@CORP100.COM -mapuser ns1@corp100.com -pass 37Le37 -out
ns1.keytab -ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

where:

-princ = Kerberos principal

- DNS = Service name in uppercase format
- ns1.corp100.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance
- CORP100.COM = The Kerberos realm in uppercase format; this must be the same as the AD domain name

-mapuser = Maps the Kerberos principal name to the AD user account

- ns1@corp100.com = The AD user name for the NIOS appliance

-pass = The AD user account password

- 37Le37 = The password of the user account for the NIOS appliance

-out = Exports the keytab file

- ns1.keytab = The name of the keytab file

-ptype = Sets the principal type. This must be **krb5_nt_principal**.

-crypto = Specifies the encryption type. This must be **des-cbc-md5**.

+DesOnly = Specifies DES encryption for the account. Include this if you did not enable DES encryption for the account.

Generating the Keytab on Windows Server 2008/Windows Server 2008 R2

A Windows Server 2008 or Windows Server 2008 R2 domain controller allows you to generate a keytab file with multiple keys for one principal. The Infoblox DNS server accepts GSS-TSIG updates from DHCP clients that provide a Kerberos ticket for any of the keys in its configured keytab. To generate the keytab file using the Ktpass tool:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:

```
ktpass -princ DNS/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab
-ptype krb5_nt_principal -crypto encryption
```

For example:

```
ktpass -princ DNS/ns1.corp100.com@CORP100.COM -mapuser ns1@corp100.com -pass 37Le37 -out
ns1.keytab -ptype krb5_nt_principal -crypto AES256-SHA1
```

where:

-princ = Kerberos principal

- DNS = Service name in uppercase format
- ns1.corp100.com = Instance in FQDN format; this is the same as the DNS name of the NIOS appliance
- CORP100.COM = The Kerberos realm in uppercase; this must be the same as the AD domain name

-mapuser = Maps the Kerberos principal name to the AD user account

- ns1@corp100.com = The AD user name for the NIOS appliance

-pass = The AD user account password

- 37Le37 = The password of the user account for the NIOS appliance

-out = Exports the keytab file

- ns1.keytab = The name of the keytab file

-ptype = Sets the principal type. This must be **krb5_nt_principal**.

-crypto = Specifies the encryption type.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```

Targeting domain controller: qacert.test.local
Using legacy password setting method
Successfully mapped DNS/ns1.corp100.com to ns1.
Key created.
Output keytab to ns1.keytab:
Keytab version: 0x502
keysize 80 DNS/ns1.corp100.com@CORP100.COM ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype 0x12
(AES256-SHA1)
keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)

```

Note: The keytab file contains highly sensitive data for the NIOS appliance account. Ensure that you store and transport its contents securely.

Modifying an AD User Account

To change any AD user account information (login, password, etc):

1. Remove the previous user account from AD.
2. Create a new user for GSS-TSIG mapping.
3. Generate a new keytab file.
4. Import the keytab file to the DNS server.

Importing the Keytab File and Enabling GSS-TSIG Authentication

Before you can enable GSS-TSIG authentication, you must import the keytab file from the Kerberos account for the NIOS appliance. To import the keytab file:

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* check box -> Edit icon.
2. In the *Member DNS Properties* editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **GSS-TSIG** and do the following:

If a principal name and version number are listed, there is a keytab file loaded on the appliance. Compare this information with that for the NIOS appliance account on the Kerberos server to make sure that they match. If there is no keytab file on the NIOS appliance or if the loaded keytab file does not match that on the Kerberos server, you must load the correct keytab file

 - Click **Upload**, click **Browse** to navigate to the keytab file, and then click **Upload**.
 - **Enable GSS-TSIG authentication of clients:** Select this check box.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Each time you export a keytab file from a Kerberos server running on Windows Server 2003, the version number of the keytab file increases incrementally. Because the version number on the keytab file that you import to the NIOS appliance must match the version that is in use on the Kerberos server, you should select the last keytab file that is exported from the Kerberos server if you have exported multiple keytab files. (A Kerberos server running on Windows 2000 does not increase the version number of keytab files with each export.)

Accepting GSS-TSIG Updates

You can allow a Grid or specific members or zones to accept GSS-TSIG signed updates from domain controllers and DHCP clients and servers, as follows:

1. Grid: From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
 Member: From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* check box -> Edit icon.
 Zone: From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box -> Edit icon.
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Select the **Updates** tab and do the following in the **Basic** subtab:
 - **Allow GSS-TSIG signed updates:** Select this option.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can then use the *Active Directory* wizard or navigate to the **Active Directory** tab of the *Authoritative Zone* editor to enable the appliance to create underscore zones for the records hosted by domain controllers and to allow GSS-TSIG signed updates to the underscore zones.

To use the *Active Directory* wizard:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Configure Active Directory**.
2. In the *Configure Active Directory* wizard, complete the following, and then click **Next**:
 - **Select Zone:** Click this and select a zone. The name of the zone must match the name in the AD domain controller so the zone transfer from the AD domain controller to the NIOS appliance can succeed.
 - **Allow GSS-TSIG-signed (secure) updates from Domain Controllers:** Select this option.
3. Complete the following:
 - **Do you want to create underscore zones to hold the records added by the Domain Controllers?**
This option allows the appliance to create the following subzones that the DNS server must have to answer AD-related DNS queries:
 - `_msdcs.zone`
 - `_sites.zone`
 - `_tcp.zone`
 - `_udp.zone`
 - `domaindnszones.zone`
 - `forestdnszones.zone`
 Note that these zones are automatically generated. You cannot edit these zones or import data into them.
 - **Allow GSS-TSIG-signed updates to underscore zones:** Select this check box to allow underscore zones to accept GSS-TSIG signed updates.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

To use the *Authoritative Zone* editor:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box -> Edit icon.
2. In the *Authoritative Zone* editor, select the **Active Directory** tab and do the following:
 - **Allow unsigned updates from these Domain Controllers:** Clear this check box.
 - **Automatically create underscore zones:** (select)
This option automatically creates the following subzones that the DNS server must have to answer AD-related DNS queries:
 - `_msdcs.zone`
 - `_sites.zone`
 - `_tcp.zone`
 - `_udp.zone`
 - `domaindnszones.zone`
 - `forestdnszones.zone`
 Note that these zones are automatically generated and cannot be manually edited.
 - **Allow GSS-TSIG-signed updates to underscore zones:** Select this check box to allow underscore zones to accept GSS-TSIG signed updates.
3. Save the configuration and click **Restart** if it appears at the top of the screen.



Chapter 21 DNSSEC

This chapter provides general information about DNSSEC. The topics in this chapter include:

- [*About DNSSEC*](#) on page 734
 - [*DNSSEC Resource Records*](#) on page 735
 - [*DNSKEY Resource Records*](#) on page 735
 - [*RRSIG Resource Records*](#) on page 737
 - [*NSEC/NSEC3 Resource Records*](#) on page 738
 - [*NSEC3PARAM Resource Records*](#) on page 739
 - [*DS Resource Records*](#) on page 740
- [*Configuring DNSSEC on a Grid*](#) on page 741
- [*Enabling DNSSEC*](#) on page 743
- [*Setting DNSSEC Parameters*](#) on page 743
 - [*About the DNSKEY Algorithm*](#) on page 743
 - [*About Key Rollovers*](#) on page 744
 - [*RRSIG Signatures*](#) on page 745
 - [*Configuring DNSSEC Parameters*](#) on page 745
- [*Signing a Zone*](#) on page 746
- [*Managing Signed Zones*](#) on page 747
 - [*Importing a Keyset*](#) on page 748
 - [*Exporting Trust Anchors*](#) on page 748
 - [*Checking Key-Signing Keys*](#) on page 749
 - [*Rolling Key-Signing Keys*](#) on page 749
 - [*Unsigning a Zone*](#) on page 749
 - [*Deleting and Restoring Signed Zones*](#) on page 749
- [*About HSM Signing*](#) on page 750
 - [*Configuring a SafeNet HSM Device*](#) on page 750
 - [*Adding and Managing a Thales HSM Group*](#) on page 752
 - [*Synchronizing the HSM Group*](#) on page 754
 - [*Monitoring the HSM Group*](#) on page 753
 - [*Enabling HSM Signing*](#) on page 753
- [*Configuring Grid Members to Support DNSSEC as Secondary Servers*](#) on page 754
- [*Enabling Recursion and Validation for Signed Zones*](#) on page 755
 - [*Enabling Recursion and Validation for Signed Zones*](#) on page 755
 - [*Enabling DNSSEC Validation*](#) on page 755

ABOUT DNSSEC

DNSSEC (DNS Security Extensions) provides mechanisms for authenticating the source of DNS data and ensuring its integrity. It protects DNS data from certain attacks, such as man-in the middle attacks and cache poisoning. A man-in-the middle attack occurs when an attacker intercepts responses to queries and inserts false records. Cache poisoning can occur when a client accepts maliciously created data. DNSSEC helps you avoid such attacks on your networks.

DNSSEC provides changes to the DNS protocol and additional resource records (RRs) as described in the following RFCs:

- *RFC 4033, DNS Security Introduction and Requirements*
- *RFC 4034, Resource Records for the DNS Security Extensions*
- *RFC 4035, DNSSEC Protocol Modifications*
- *RFC 4641, DNSSEC Operational Practices*
- *RFC 4956, DNS Security (DNSSEC) Opt-In*
- *RFC 4986, Requirements Related to DNS Security (DNSSEC) Trust Anchor Rollover*
- *RFC 5155, DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*
- *RFC 5702, Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC*

DNSSEC uses public key cryptography to authenticate the source of DNS responses and to ensure that DNS responses were not modified during transit. Public key cryptography uses an asymmetric key algorithm. With asymmetric keys, one key is used to decrypt data that was encrypted using the other key.

In DNSSEC, the primary name server of a zone generates at least one public/private key pair. It “signs” each data set in the zone by running it through a one-way hash, and then encrypting the hash value with the private key. The public key is stored in an RR type introduced by DNSSEC, the DNSKEY RR. Resolvers use the DNSKEY record to decrypt the hash value. If the hash values match, then the resolver is assured of the authenticity of the message.

In addition to the DNSKEY record, DNSSEC also introduces new RRs which DNS servers can use to authenticate the non-existence of servers, zones, or resource records. For information about the DNSSEC RRs, see [DNSSEC Resource Records](#) on page 735.

DNSSEC uses the EDNSO message extension. Resolvers include the EDNS OPT pseudo-RR with the DO (DNSSEC OK) bit set to indicate that they are requesting DNSSEC data. A DNS client or resolver sets the EDNS DO bit when it sends a query for data in a signed zone. When the DNS server receives such a query, it includes the additional DNSSEC records in its response, according to the DNSSEC standard rules. In addition, because DNSSEC messages are often large, the EDNSO message extension also provides mechanisms for handling larger DNS UDP messages. For information about EDNSO, refer to *RFC 2671, Extension Mechanisms for DNS (EDNS0)*. For information about the DO bit, refer to *RFC 3225, Indicating Resolver Support of DNSSEC*.

WARNING: NOTE THAT WHEN YOU DISABLE EDNSO ON THE APPLIANCE, ALL OUTGOING DNSSEC QUERIES TO ZONES WITHIN TRUSTED ANCHORS WILL FAIL EVEN IF DNSSEC VALIDATION IS ENABLED. TO ENSURE THAT DNSSEC FUNCTIONS PROPERLY, DO NOT DISABLE EDNSO ON THE APPLIANCE. FOR MORE INFORMATION, SEE [USING EXTENSION MECHANISMS FOR DNS \(EDNSO\)](#) ON PAGE 564.

DNSSEC also supports new data in the packet header, the CD (Checking Disabled) bit and the AD (Authenticated Data) bit. The CD bit is used by resolvers in their DNS queries and the AD bit is used by recursive name servers in their responses to queries.

A resolver can set the CD bit in its query to indicate that the name server should not validate the DNS response and that the resolver takes responsibility for validating the DNS data it receives.

A name server that has successfully validated the data in a DNS response sets the AD (Authenticated Data) bit in the message header to indicate that all resource records in its response have been validated and are authentic. Note that unless the connection between the DNS server and client has been secured, such as through TSIG, the client cannot rely on the AD bit to indicate valid data. The data could have been changed in transit between the server and client. Resolvers can trust a response with the AD bit set only if their communication channel is secure.

DNSSEC Resource Records

Following are the DNSSEC RR types:

- DNS Public Key (DNSKEY) resource records—For information, see [DNSKEY Resource Records](#) on page 735.
- Resource Record Signature (RRSIG) records—For information, see [RRSIG Resource Records](#) on page 737.
- Next Secure (NSEC/NSEC3) records—For information, see [NSEC/NSEC3 Resource Records](#) on page 738.
- NSEC3PARAM records—For information, see [NSEC3PARAM Resource Records](#) on page 739.
- Delegation Signer (DS) resource records—For information, see [DS Resource Records](#) on page 740.

For detailed information about each RR, refer to *RFC 4034, Resource Records for the DNS Security Extensions* and *RFC 5155, DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*.

Note: The appliance supports IDNs for DNSKEY records, DS records, NSEC records, NSEC3PARAM records, and RRSIG records.

DNSKEY Resource Records

When an authoritative name server digitally signs a zone, it typically generates two key pairs, a zone-signing key (ZSK) pair and a key-signing key (KSK) pair. The name server uses the private key of the ZSK pair to sign each RRset in a zone. (An RRset is a group of resource records that are of the same owner, class, and type.) It stores the public key of the ZSK pair in a DNSKEY record. The name server then uses the private key of the KSK pair to sign all DNSKEY records, including its own, and stores the corresponding public key in another DNSKEY record. As a result, a zone typically has two DNSKEY records; a DNSKEY record that holds the public key of the ZSK pair, and another DNSKEY record for the public key of the KSK pair.

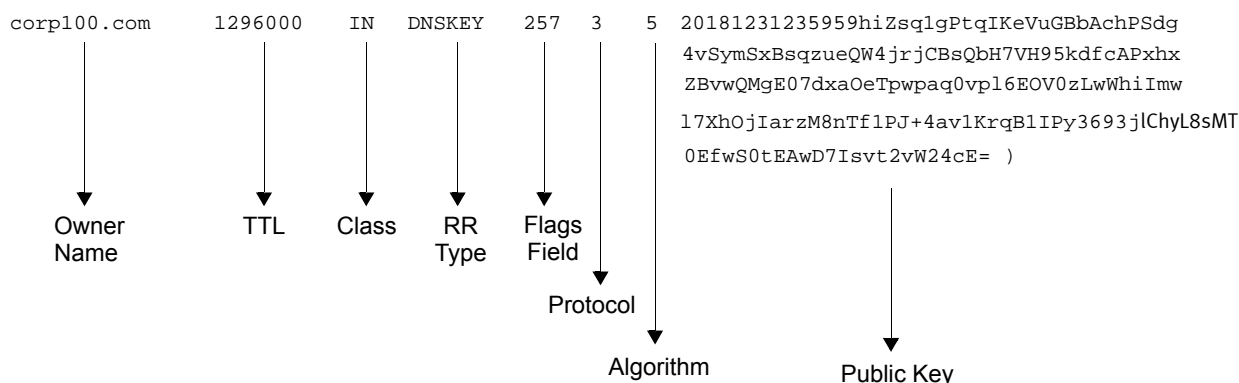
Note: For the remainder of this chapter, the DNSKEY record that holds the public key of the ZSK pair is referred to as the ZSK and the DNSKEY record that holds the public key of the KSK is referred to as the KSK.

The purpose of the KSK is two-fold. First, it is referenced in the Delegation Signer (DS) RR that is stored in a parent zone. The DS record is used to authenticate the KSK of the child zone, so a resolver can establish a chain of trust from the parent zone to its child zone. (For more information about the DS RR, see [DS Resource Records](#) on page 740).

Second, if a zone does not have a chain of trust from a parent zone, security aware resolvers can configure the KSK as a trust anchor; that is, the starting point from which it can build a chain of trust from that zone to its child zones.

Note that though the two key pairs, KSK and ZSK, are used in most DNSSEC environments, their use is not required by the RFCs. A zone administrator can use a single private/public key pair to sign all zone data. (Note that Infoblox appliances require two key pairs.)

Following is an example of a DNSKEY RR:



The first four fields specify the domain name of the zone that owns the key, the resource record TTL, class, and RR type. The succeeding fields are:

- **Flags Field:** In its wire format, this field is two bytes long. (The wire format is used in DNS queries and responses.) Bits 0 through 6 and 8 through 14 are reserved, and have a value of 0. Bit 7 indicates if the record holds a DNS zone key. Bit 15 is the Secure Entry Point (SEP) flag, which serves as a hint that indicates whether the DNSKEY record contains a ZSK or a KSK, as described in *RFC 3757, DNSKEY RR SEP Flag*. Zone administrators typically set the SEP flag of a DNSKEY record of a zone when it contains the KSK, to indicate that it can be used as a trust anchor. However, a DNSKEY record that does not have the SEP flag set can also be used as a trust anchor.

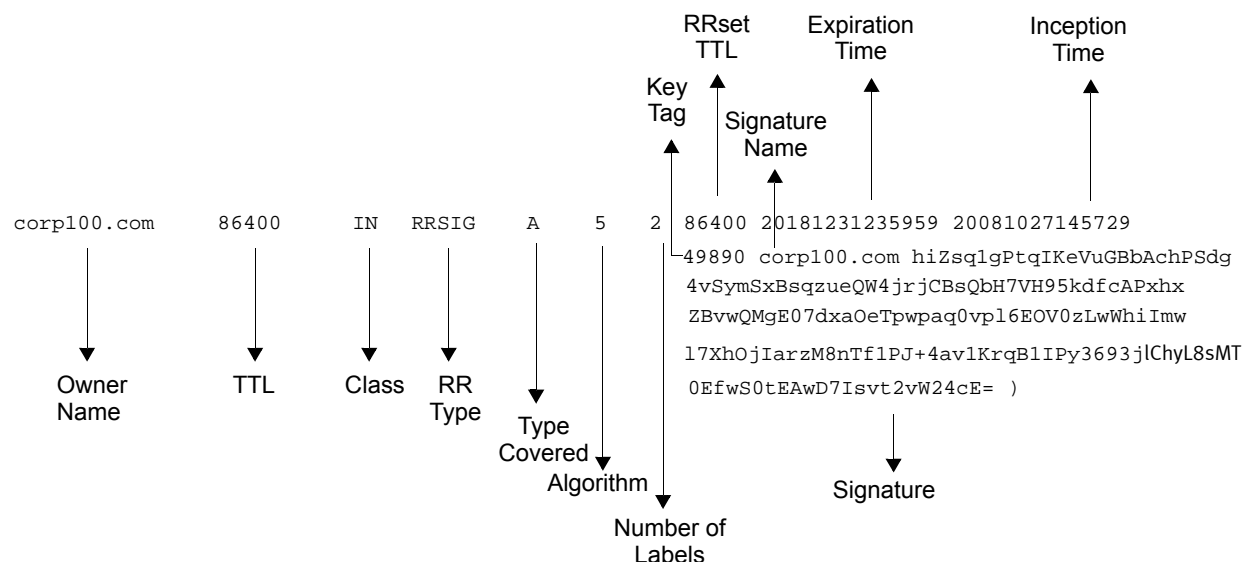
Given the currently defined flags, in its text format, the flags field is represented as an unsigned decimal integer with the possible values of 0, 256 and 257. A value of 256 indicates that the DNSKEY record holds the ZSK and a value of 257 indicates that it contains the KSK. In general, this field contains an odd number when the DNSKEY record holds the KSK.

- **Protocol:** This always has a value of 3, for DNSSEC.
- **Algorithm:** Identifies the public key's cryptographic algorithm. The available types are:
 - 1 = RSA/MD5
 - 2 = Diffie-Hellman (This is not supported by BIND and Infoblox appliances.)
 - 3 = DSA
 - 4 = Reserved
 - 5 = RSA/SHA1
 - 6 = DSA/SHA1/NSEC3
 - 7 = RSA/SHA1/NSEC3
 - 8 = RSA/SHA-256
 - 10 = RSA/SHA-512
- **Public Key:** The public key encoded in Base64.

RRSIG Resource Records

A signed zone has multiple RRsets, one for each record type and owner name. (The owner is the domain name of the RRset.) When an authoritative name server uses the private key of the ZSK pair to sign each RRset in a zone, the digital signature on each RRset is stored in an RRSIG record. Therefore, a signed zone contains an RRSIG record for each RRset.

Following is an example of an RRSIG record:




The first four fields specify the owner name, TTL, class, and RR type. The succeeding fields are:

- **Type Covered:** The RR type covered by the RRSIG record. The RRSIG record in the example covers the A records for corp100.com.
- **Algorithm:** The cryptographic algorithm that was used to create the signature. It uses the same algorithm types as the DNSKEY record indicated in the Key Tag field.
- **Number of Labels:** Indicates the number of labels in the owner name of the signed records. There are two labels in the example, corp100 and com.
- **RRset TTL:** The TTL value of the RRset covered by the RRSIG record.
- **Expiration Time:** The signature expiration time in UTC format.
- **Inception Time:** The signature inception time in UTC format.
- **Key Tag:** The key tag value of the DNSKEY RR that validates the signature.
- **Signature Name:** The zone name of the RRset.
- **Public Key:** The Base64 encoding of the signature.

NSEC/NSEC3 Resource Records

When a name server receives a request for a domain name that does not exist in a zone, the name server sends an authenticated negative response in the form of an NSEC or NSEC3 RR. NSEC and NSEC3 records contain the next secure domain name in a zone and list the RR types present at the NSEC or NSEC3 RR's owner name. The difference between an NSEC and NSEC3 RRs is that the owner name in an NSEC3 RR is a cryptographic hash of the original owner name prepended to the name of the zone. NSEC3 RRs protect against zone enumeration.

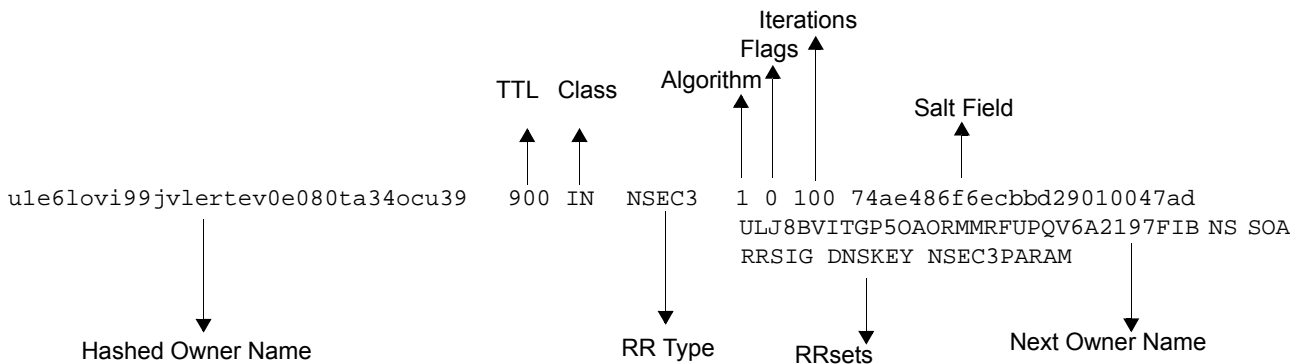
Following is an example of an NSEC record:



The first four fields specify the owner name, TTL, class and RR type. The succeeding fields are:

- **Next Owner Name:** In the canonical order of the zone, the next owner name that has authoritative data or that contains a delegation point NS record.
- **RRsets:** The RRsets that exist at the owner name of the NSEC record, which are NS, SOA, RRSIG, NSEC, and DNSKEY in the example.

Following is an example of an NSEC3 RR:



The first field contains the hashed owner name. It is followed by the TTL, class and RR type. The fields after the RR type are:

- **Algorithm:** The hash algorithm that was used. The currently supported algorithm is SHA-1, which is represented by a value of 1.
- **Flags Field:** Contains 8 one-bit flags, of which only one flag, the Opt-Out flag, is defined by RFC 5155. The Opt-Out flag indicates whether the NSEC3 record covers unsigned delegations.
- **Iterations:** The number of times the hash function was performed.
- **Salt Field:** A series of case-insensitive hexadecimal digits. It is appended to the original owner name as protection against pre-calculated dictionary attacks.
- **Next Owner Name:** Displays the next hashed owner name.
- **RRsets:** The RR types that are at the owner name.

NSEC3PARAM Resource Records

An authoritative DNS server uses NSEC3PARAM RRs to determine which NSEC3 records it includes in its negative responses. An NSEC3PARAM RR contains the parameters that an authoritative server needs to calculate hashed owner names. As stated in RFC 5155, the presence of an NSEC3PARAM RR at a zone apex indicates that the specified parameters may be used by authoritative servers to choose an appropriate set of NSEC3 RRs for negative responses. Following is an example of an NSEC3PARAM record:

corp100.com	900	IN	NSEC3PARAM	1	0	100	74ae486f6ecbbd29010047ad
↓	↓	↓	↓	↓	↓	↓	↓
Owner Name	TTL	Class	RR Type	Algorithm	Flags	Iterations	Salt Field

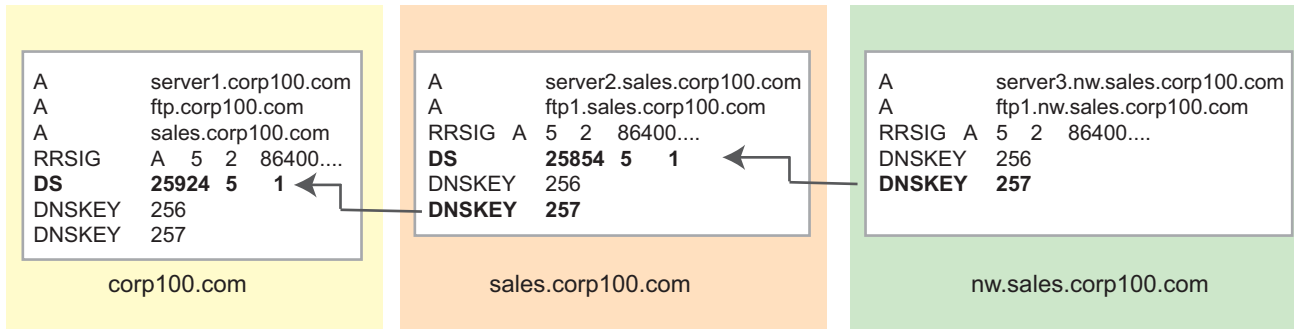
The first four fields specify the owner name, TTL , class and RR type. The succeeding fields are:

- **Algorithm:** The hash algorithm that was used. The currently supported algorithm is SHA-1, which is represented by a value of 1.
- **Flags Field:** Contains 8 one-bit flags, of which only one flag, the Opt-Out flag, is defined by RFC 5155. The Opt-Out flag indicates whether the NSEC3 record covers unsigned delegations.
- **Iterations:** The number of times the hash function was performed. The number of NSEC3 iterations is set to 10.
- **Salt Field:** A series of case-insensitive hexadecimal digits. It is appended to the original owner name as protection against pre-calculated dictionary attacks. New salt value is generated when the ZSK rolls over, for which the user can control the period. For random salt, the selected length is between one and 15 octets.

DS Resource Records

A DS RR contains a hash of a child zone's KSK and can be used as a trust anchor in some security-aware resolvers and to create a secure delegation point for a signed subzone in DNS servers. As illustrated in [Figure 21.1](#), the DS RR in the parent zone corp100.com contains a hash of the KSK of the child zone sales.corp100.com, which in turn has a DS record that contains a hash of the KSK of its child zone, nw.sales.corp100.com.

Figure 21.1



Following is an example of the DS RR:

```
corp100.com      86400      IN      DS      25924      5      1      49D2801B50E25D59440F1FF1A8012B568435
                                     B622B1F8709F33D744C4C6D71EA2
```

Labels for the fields above:

- corp100.com → Owner Name
- 86400 → TTL
- IN → Class
- DS → RR Type
- 25924 → Key Tag
- 5 → Algorithm
- 1 → Digest Type
- 49D2801B50E25D59440F1FF1A8012B568435B622B1F8709F33D744C4C6D71EA2 → Digest

The first four fields specify the owner name, TTL, class and RR type. The succeeding fields are as follows:

- **Key Tag:** The key tag value that is used to determine which key to use to verify signatures.
- **Algorithm:** Identifies the algorithm of the DNSKEY RR to which this DS RR refers. It uses the same algorithm values and types as the corresponding DNSKEY RR.
- **Digest Type:** Identifies the algorithm used to construct the digest. The supported algorithms are:
 - 1 = SHA-1
 - 2 = SHA-256
- **Digest:** If SHA-1 is the digest type, this field contains a 20 octet digest. If SHA-256 is the digest type, this field contains a 32 octet digest.

CONFIGURING DNSSEC ON A GRID

You can configure the name servers in a Grid to support DNSSEC. You can configure the Grid Master as the primary server for a signed zone and the Grid members as secondary servers. (For more information, see [Configuring Grid Members to Support DNSSEC as Secondary Servers](#) on page 754.) Note that only the Grid Master can serve as the primary server for a signed zone.

You can enable the Grid Master to sign zones and manage the DNSSEC keys, or you can configure the Grid Master as a client to a third-party, network-attached Hardware Security Module (HSM) that performs the key generation, zone signing, and key safekeeping. You must use either the Grid Master or HSM for zone signing and key management; you cannot use both. Note that each method may have different performance implications, depending on the hardware platform, number of zones and other factors. For information about using HSMs, see [About HSM Signing](#) on page 750.

Any authoritative forward-mapping or reverse-mapping zone can be signed according to the following criteria:

- The zone does not contain any bulk host records.
- DNSSEC is enabled on the Grid Master.
- The primary server of the zone must be a Grid member. If the zone is assigned to an NS group, the primary server in the group must be a Grid member that has DNSSEC enabled.

Note that you can use DNS views to separate internal and external zone data, to manage your zones more efficiently and reduce the size of the zones that require signing. For information about DNS views, see [Using Infoblox DNS Views](#) on page 602.

Grid Master as Primary Server

When you sign a zone whose primary server is a Grid member, that member becomes a secondary server and the Grid Master becomes the hidden primary server. If the zone is assigned to an NS group, the Grid Master removes the association with the NS group. The previous primary server becomes a secondary server for the zone.

If a master candidate is promoted to Grid Master and the previous Grid Master was the primary server for signed zones, the new Grid Master becomes the hidden primary server for all signed zones. The previous Grid Master, which was the primary server for the zone, becomes a secondary server for the zone.

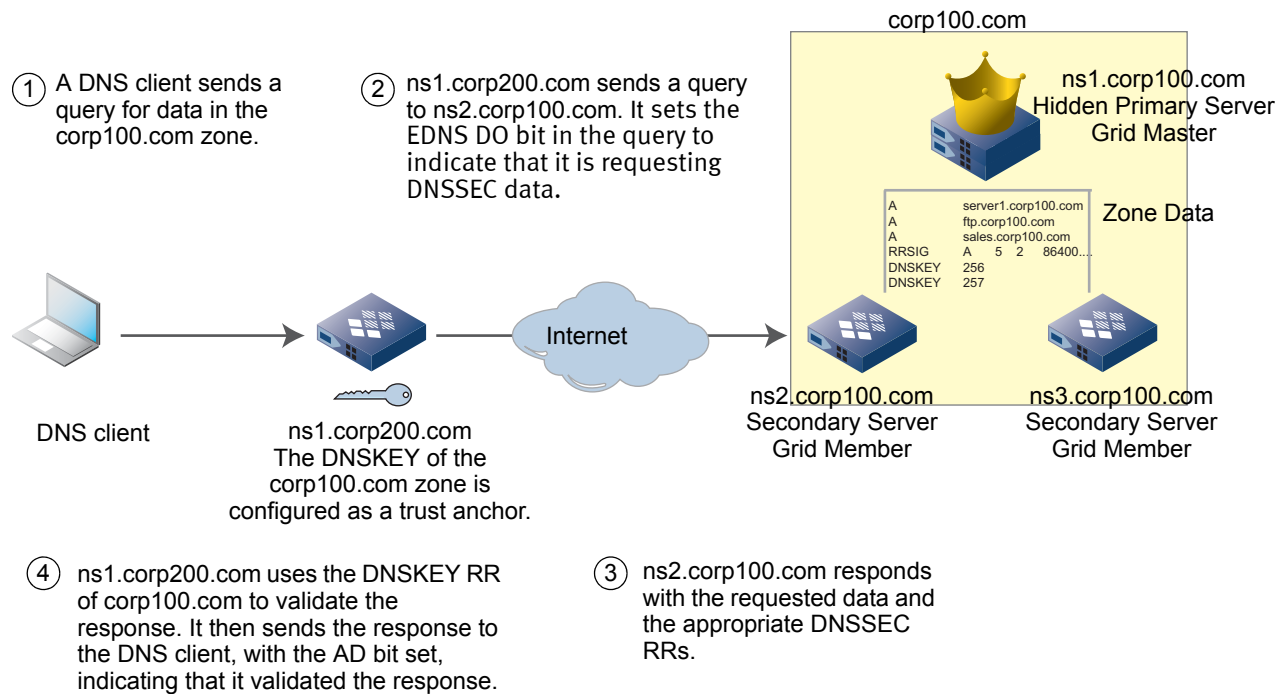
As the primary server, the Grid Master sends zone data to the secondary servers through zone transfers; or, if the secondary servers are Grid members, the Grid Master transfers data to all Grid members through the database replication process, by default. The Grid Master transfers all records in that zone, including all NSEC/NSEC3, RRSIG, DNSKEY and DS records with owner names that belong to that zone. The RRSIG RRs are included in zone transfers of the zone in which they are authoritative data. The Grid Master also performs incremental zone transfers to secondary servers as a result of incremental zone signings.

In addition, the Grid Master automatically performs an incremental signing of the zone data sets when their contents change. Incremental signing refers to signing just those parts of a zone that change when RRs are added, modified, or deleted. The Grid Master uses the private key of the ZSK when it incrementally signs a zone. In addition, the Grid Master adds, modifies or deletes the corresponding RRSIG records and the appropriate NSEC/NSEC3 records.

For example, [Figure 21.2](#) shows a Grid Master as the primary server of a signed zone and its Grid members as secondary servers. The Grid Master, ns1.corp100.com, is the hidden primary DNS server for the corp100.com zone. As the hidden primary name server for corp100.com, the Grid Master does not respond to queries from other name servers. Instead, it provides data to its secondary servers, ns2.corp100.com and ns3.corp100.com, which use this data to respond to DNS queries. Because the secondary servers are Grid members, they receive zone data from the Grid Master through the Grid database replication process.

The name server ns1.corp200.com is a recursive name server. It has configured the DNSKEY of the corp100.com zone as a trust anchor. Therefore, it is able to validate the data it receives when it sends a query for the corp100.com zone.

Figure 21.2



Following are the tasks to configure the Grid Master to sign zones:

1. Create the zones. For information, see [Configuring Authoritative Zones](#) on page 616.
 - Specify the Grid Master as the primary server.
2. Enable DNSSEC, as described in [Enabling DNSSEC](#).
3. Optionally, change the default DNSSEC settings. For information, see [Setting DNSSEC Parameters](#) on page 743.
4. Sign the zone. The appliance automatically generates the DNSSEC RRs when you sign a zone. For information, see [Signing a Zone](#) on page 746.

ENABLING DNSSEC

You can enable DNSSEC on a Grid, individual members, and DNS views. Because only Grid Masters can serve as primary servers for signed zones, you must enable DNSSEC on the Grid Master before you can sign zones. You must also enable DNSSEC on any Grid member that serves as a secondary server for signed zones.

When you enable DNSSEC on a Grid, you can set certain parameters that control the DNSSEC RRs, as described in [Setting DNSSEC Parameters](#) on page 743.

When you enable DNSSEC on a Grid member or DNS view, you can set parameters that affect its operations as a secondary server, as described in [Configuring Grid Members to Support DNSSEC as Secondary Servers](#) on page 754.

To enable DNSSEC on a Grid, member or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **Members** tab -> *member* check box and click the Edit icon.
DNS View: From the **Data Management** tab, select the **Zones** tab -> *dns_view* check box and click the Edit icon.
2. In the editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **DNSSEC**.
4. In the **DNSSEC** tab, select **Enable DNSSEC**.

Note: When you disable EDNS0, all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled. This is due to the restriction of the UDP packet length when you disable EDNS0. For information about EDNS0, see [Using Extension Mechanisms for DNS \(EDNS0\)](#) on page 564.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

SETTING DNSSEC PARAMETERS

The Grid Master uses certain default parameters when it signs a zone and generates the DNSSEC RRs. You can change these defaults for the entire Grid and for individual zones, in case you want to use different parameters for certain zones. The following sections describe the different parameters that you can set:

- [About the DNSKEY Algorithm](#)
- [About Key Rollovers](#)
- [RRSIG Signatures](#) on page 745

For information on setting these parameters, see [Configuring DNSSEC Parameters](#) on page 745.

About the DNSKEY Algorithm

You can select the cryptographic algorithm that the Grid Master uses when it generates the KSK and ZSK. By default, it uses RSA/SHA1 and generates NSEC RRs. If you want the Grid Master to generate NSEC3 RRs, you must select DSA/NSEC3, RSA/SHA1/NSEC3, RSA/SHA-256/NSEC3 or RSA/SHA-512/NSEC3 as the algorithm for both the KSK and ZSK.

When you select an algorithm for the KSK, the Grid Master automatically assigns the same algorithm to the ZSK. You can change this algorithm, but the algorithms used by the KSK and ZSK must generate the same type of NSEC record. A zone cannot contain both NSEC and NSEC3 RRs.

You can select the DNSKey algorithm for HSMs. Thales HSMs don't support DSA. All other parameters are not used by HSMs.

About Key Rollovers

To reduce the probability of their being compromised, ZSKs and KSKs must be periodically changed. The time within which a key pair is effective is its rollover period. The rollover period starts as soon as a zone is signed. After a rollover period starts, you cannot interrupt or restart it unless you unsign the zone.

Zone-Signing Key Rollover

ZSK rollovers occur automatically on the Grid Master, using the double signature rollover method described in RFC 4641. This method provides for a grace period, which is half of the rollover period. The default ZSK rollover period is 30 days; thus the default grace period is 15 days.

At the end of a rollover period of a ZSK, the Grid Master generates a new ZSK key pair. It signs the zone with the private key of the new ZSK key pair, and consequently generates new RRSIG RRs with the new signatures. However, the Grid Master also retains the old ZSK key pair and RRSIG RRs. Thus during the grace period, the data in the zone is signed by the private keys of both the old and new ZSKs. Their corresponding public keys (stored in DNSSEC RRs) can be used to verify both the old and new RRSIGs.

The grace period also allows the data that exists in remote caches to expire and during this time, the updated zone data can be propagated to all authoritative name servers.

The Grid Master removes the old ZSK and its RRSIGs when the rollover grace period elapses.

Key-Signing Key Rollover

Unlike ZSK rollovers, which occur automatically, KSK rollovers must be initiated by an admin. When the KSK rollover is overdue or is due within seven days, the Grid Master displays a warning when admins log in. In addition, you can also check which KSKs are due for a rollover as described in [Checking Key-Signing Keys](#) on page 749.

The Grid Master also uses the double signature rollover method described in RFC 4641 for KSK rollovers.

When a user initiates a KSK rollover, the Grid Master sets the grace period to half the KSK rollover period. It generates a new KSK, and signs the DNSKEY records with the new KSK. Thus during the grace period, the DNSKEY records are signed by the private keys of both the old and new KSKs. Both the old and the new KSKs can be used to validate the zone. The grace period allows the old keys in remote caches to expire. In addition, the admin should also export the new KSK and send it to the recursive name servers that use the KSK as trust anchors.

If the KSK rollover is for a child zone and the primary server of the parent zone is a Grid member, the Grid Master also inserts a DS record in the parent zone for the new DNSKEY in the child zone. If the primary server of the parent zone is external to the Grid, the admin must export either the DS record or the new KSK to the admin of the parent zone. For information about exporting a KSK, see [Exporting Trust Anchors](#) on page 748.

The Grid Master then removes the old KSK and its RRSIG records when the grace period for the KSK rollover ends.

About Key Rollovers and DNS TTLs

Note that the KSK and ZSK rollover intervals affect TTLs used by RRs in signed zones.

A grace period is half of the key rollover interval. For example, if the KSK rollover interval is 1 year (365 days), then the grace period is 182.5 days; if the ZSK rollover interval is 30 days, then the grace period is 15 days.

The DNSKEY RRset in the zone is assigned a TTL that is the minimum of the KSK and ZSK grace period. In the preceding example, the minimum or lowest of these is 15 days. Therefore, the TTLs used for the DNSKEY RRset are 15 days (1296000 seconds).

All other RRs in the signed zone are limited to a “zone maximum TTL,” which is the grace period of the ZSK. In the example, this is also 15 days.

When the zone is initially signed, if the TTL of an RR exceeds the zone maximum TTL, the Grid Master reduces the TTL to the zone maximum TTL. Additionally, the TTL settings for the signed zone are set to override; the values are inherited from the Grid DNS properties at that time, and the default TTL setting is reduced to the zone maximum TTL if the Grid property exceeds it. If the zone is later unsigned, the zone DNS properties remain at their overridden settings.

RRSIG Signatures

As shown in the sample RRSIG record in [RRSIG Resource Records](#) on page 737, the signatures have an inception and an expiration time. The default validity period of signatures in RRSIG records on the Grid Master is four days. You can change this default, as long as it is not less than one day or more than 3660 days. The Grid Master automatically renews signatures before their expiration date.

Configuring DNSSEC Parameters

To set parameters at the Grid or zone level:

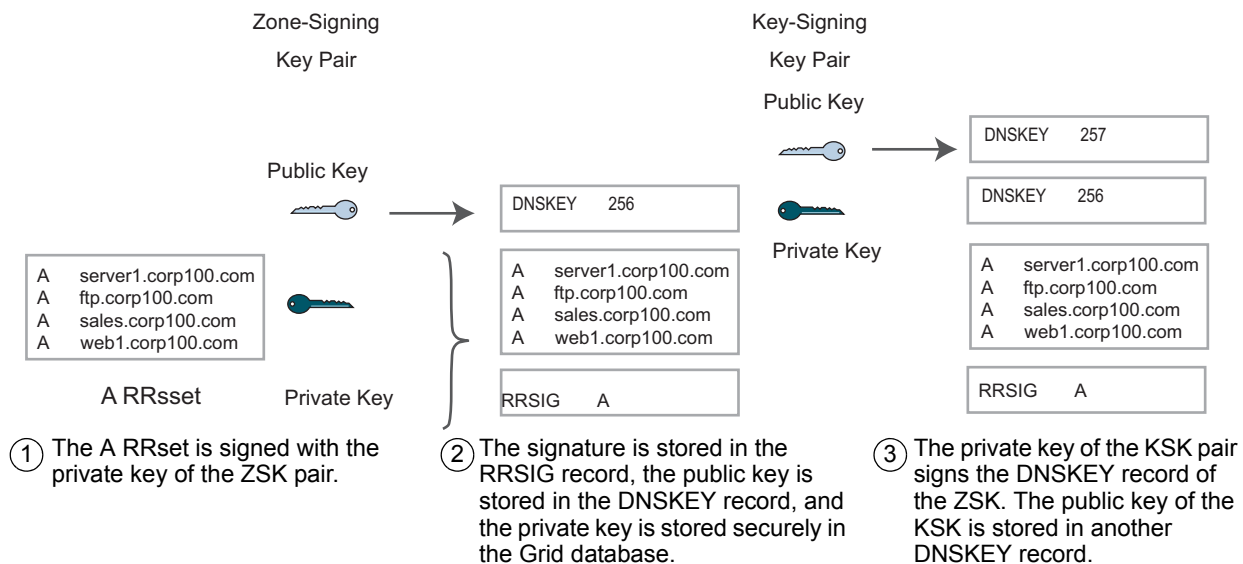
1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.
Zone: From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box, and click the Edit icon. Click **Override** to override the parameters.
2. In the editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **DNSSEC**.
4. In the **DNSSEC** tab, complete the following:
 - **Key-signing Key:** Select the cryptographic algorithm that the Grid Master or HSM uses when it generates the KSK. Note that Thales HSMs do not support DSA. The default is **RSA/SHA1**. Select DSA/NSEC3, RSA/SHA1/NSEC3, RSA/SHA-256/NSEC3 or RSA/SHA-512/NSEC3 to use NSEC3 instead of NSEC records in signed zones. You can also select the default key length for the KSK. Following are the valid values for each algorithm:
 - DSA:** The minimum is 512 bits and the maximum is 1024 bits, which is also the default. The key length must be a multiple of 64.
 - DSA/NSEC3:** The minimum is 512 bits and maximum is 1024 bits, which is also the default. The key length must be a multiple of 64.
 - RSA/MD5:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - RSA/SHA1:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - RSA/SHA1/NSEC3:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - RSA/SHA-256:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - RSA/SHA-256/NSEC3:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - RSA/SHA-512:** The minimum is 1024 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - RSA/SHA-512/NSEC3:** The minimum is 1024 bits, the maximum is 4096 bits, and the default is 2048 bits.
 - **Key-signing Key Rollover Interval:** The minimum value is one day and the maximum is the time remaining to January 2038. The default is one year.
 - **Zone-signing Key:** Select the cryptographic algorithm that the Grid Master or HSM uses when it generates the ZSK. Note that HSMs do not support DSA. When you select an algorithm for the KSK, the Grid Master automatically selects the same algorithm for the ZSK. You can change the algorithm. However, the algorithms used by the KSK and ZSK must use the same type of NSEC record. You can also select the default key length for the zone-signing key. Following are the valid values for each algorithm:
 - DSA:** The minimum is 512 bits and the maximum is 1024 bits. The default is 1024 bits.
 - DSA/NSEC3:** The minimum is 512 bits and maximum is 1024 bits. The default is 1024 bits.
 - RSA/MD5:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.
 - RSA/SHA1:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.
 - RSA/SHA1/NSEC3:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.
 - RSA/SHA-256:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.
 - RSA/SHA-256/NSEC3:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.
 - RSA/SHA-512:** The minimum is 1024 bits, the maximum is 4096 bits, and the default is 1024 bits.
 - RSA/SHA-512/NSEC3:** The minimum is 1024 bits, the maximum is 4096 bits, and the default is 1024 bits.

- **Zone-signing Key Rollover Interval:** The minimum value is one day and the maximum is the time remaining to January 2038. The default is 30 days.
 - **Signature Validity:** Specify the signature validity period for RRSIG RRs. The minimum is one day and the maximum is 3660 days. The default signature validity interval is four days.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

SIGNING A ZONE

When it signs a zone, the Grid Master generates new DNSKEY key pairs. As shown in [Figure 21.3](#), it uses the private key of the ZSK to sign the authoritative RRsets in the zone, and stores the corresponding public key in a DNSKEY record. It then uses the private key of the KSK to sign the DNSKEY records and stores the corresponding public key in another DNSKEY record. It stores the private keys in the Grid database and stores the public keys in the DNSKEY records in the database.

Figure 21.3 Zone Signing Process



The Grid Master also does the following:

- It inserts NSEC/NSEC3 records for each label. The use of NSEC or NSEC3 RRs depends on the algorithm you selected for the KSK and ZSK. When you select DSA/NSEC3, RSA/SHA1/NSEC3, RSA/SHA-256/NSEC3 or RSA/SHA-512/NSEC3, the Grid Master uses NSEC3 records in signed zones. Note that a zone cannot contain both NSEC and NSEC3 RRs. If you want to change the type of NSEC records that a zone uses, you must unsign the zone, change the algorithm for the KSK and ZSK, and then re-sign the zone.
- It increments the SOA serial number and notifies the secondary servers that there is a change to its zone data. When the secondary servers check the serial number and see that it has been incremented, the secondary servers request a zone transfer.
- If the TTL of an RR in the zone exceeds the ZSK grace period, the Grid Master reduces the TTL to the ZSK grace period. (For information about the grace period, see [About Key Rollovers](#) on page 744.) Setting a TTL value that exceeds half of the rollover period is not allowed.
- If the KSK rollover period is less than the ZSK rollover period, the Grid Master sets the TTL of the DNSKEY RR to the KSK rollover period.

When it signs a subzone, the Grid Master automatically inserts DS records for parent zones that are hosted by Grid members. To sign a zone:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Sign Zone**.
3. In the *Sign Zone* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are signing. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Select a zone, and then click **Sign Zone**.
4. When the confirmation dialog displays, click **Yes**.

To view the records of the signed zone, from the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone*. Expand the Records section to list the RRs of the zone, as shown in [Figure 21.4](#).

Figure 21.4

Name	Type	Data	Comment	Site
	SOA Record	Serial 4 MNAME infoblox.localdon RNAME please_set_ema Refresh 10800 Retry 1080 Expire 2419200 Negative caching TTL 900	Auto-created by ...	
	RRSIG Record	3600 DNSKEY 5 1 3600 20091102185...		
	RRSIG Record	3600 DNSKEY 5 1 3600 20091102185...		
	DNSKEY Record	43200 257 5 36337 AwEAdJ8s1fjAl...		
	DNSKEY Record	43200 256 5 2145 AwEAd4PGluY+j...		
	RRSIG Record	28800 SOA 5 1 28800 200911021858...		
	NS Record	docdemo-2.infoblox	Auto-created by ...	
	RRSIG Record	900 NSEC 5 1 900 20091102185818 2...		
	NSEC Record	900 host2.test NS SOA RRSIG NSEC ...		

MANAGING SIGNED ZONES

After you sign a zone, you can do the following:

- You can add a DS RR at the delegation point for a signed subzone when the subzone is hosted on a non-Infoblox DNS server or an Infoblox server that is part of a different Grid. For information, see [Importing a Keyset](#) on page 748.
- Trust anchors can be specified as DNSKEY RRs, DS RRs, and as a BIND trusted-keys statement. You can export any of these as trust anchors. For information, see [Exporting Trust Anchors](#) on page 748.
- You must change the KSK periodically, to ensure its security. For information, see [Checking Key-Signing Keys](#) on page 749 and [Rolling Key-Signing Keys](#) on page 749.
- If, for any reason, the security of the keys are compromised, you can perform an emergency replacement of both the zone-signing and key-signing keys by unsigning the zone, and then re-signing it. For information about unsigning the zone, see [Unsigning a Zone](#) on page 749.

Note that when you re-sign a zone, the Grid Master generates new ZSK and KSK pairs. You must send the new DNSKEY of the KSK to resolvers that use it as a trust anchor and generate new DS records and send them to the parent zones.

- You can move a signed zone to the Recycle Bin, from where you can delete it permanently or restore it. For information, see [Deleting and Restoring Signed Zones](#) on page 749.

In addition, signed zones can accept dynamic DNS updates. For information about configuring zones to accept dynamic DNS updates, see [Configuring DNS Servers for DDNS](#) on page 705.

Importing a Keyset

A keyset is a DS RRset, or a DNSKEY RRset which is used as input to generate the DS RRset. To import a keyset:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Import Keyset**.
3. In the *Import Keyset* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are importing the keyset. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select a zone.
4. Paste the KSK or DS record being imported. It must be a KSK or DS record, and must belong to an immediate subzone of the zone to which the record is being imported.
5. Click **Import**.

If you imported a DNSKEY RRset, the Grid Master uses the SHA-1 algorithm to generate the DS RRset, which it adds to the parent zone. If you imported a DS RRset, the Grid Master adds it to the parent zone. The Grid Master incrementally signs the DS RRset.

Exporting Trust Anchors

A trust anchor is a DNSSEC public key which is used by security-aware resolvers as the starting point for establishing authentication chains. A trust anchor can be specified as a DNSKEY RR or a DS RR, which contains the hash of a DNSKEY RR and can also be used to create a secure delegation point for a signed subzone in DNS servers.

In BIND, trust anchors are configured using the `trusted-keys` directive. A trusted key is a DNSKEY RR without the TTL, class and RR type. You can export the trust anchors for the selected zone in a format that can be used in a BIND `trusted-keys` directive.

To export trust anchors:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Export Trust Anchors**.
3. In the *Export Trust Anchors* dialog box, do the following:
 - The displayed zone name can either be the last selected zone or the zone from which you are exporting trust anchors. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one.
 - Select one of the following: **DNSKEY records**, **DS records**, or **BIND trusted-keys statement**.
4. Click **Export**.
5. Specify the location of the exported file and click **OK**.

If you exported DS records, the exported file contains DS records that use the SHA-1 and SHA-256 algorithms.

Checking Key-Signing Keys

To check which key-signing keys are overdue for a rollover or are due to roll over within a week:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Check Key-Signing Keys**.
3. The *KSK Rollover Due* dialog box lists the key-signing keys that are due to rollover. It includes the domain name of the zone, DNS view (if there are multiple DNS views), and the number of days until the rollover.
4. Click **Close**.

Rolling Key-Signing Keys

Unlike ZSKs, which are automatically rolled over, KSK rollovers must be initiated by an admin. You can initiate a rollover before or after a rollover period, or when you need to replace the KSK for security reasons. You can initiate a rollover at anytime, as long as a KSK rollover is not already in progress for the zone.

To roll over key-signing keys:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Roll Over Key-Signing Key**.
3. In the *Roll Over Key-Signing Key* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are rolling over key-signing keys. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one.
4. Click **Roll Over**.

You can export the new KSK and send it to the security-aware resolvers that use it as a trust anchor.

Unsigning a Zone

When you need to perform an emergency key rollover, you can unsign a zone and then re-sign it to generate new ZSK and KSK key pairs. When you unsign a zone, the Grid Master permanently removes all automatically generated DNSSEC records in the zone and parent zone. It does not remove any DS records associated with a child zone.

To unsign a zone:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Unsign Zone**.
3. In the *Unsign Zone* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are signing. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one. After you have selected a zone, click **Unsign Zone**.
4. When the confirmation dialog displays, click **Yes**.

Deleting and Restoring Signed Zones

When you delete a signed zone, the Grid Master unsigns the zone before moving it to the Recycle Bin. Unsigning the zone effectively deletes all auto-generated DNSSEC RRs; only user-defined DS records are retained and moved to the Recycle Bin as well. The Grid Master also retains the ZSK and KSK in its database, until you permanently delete the zone from the Recycle Bin.

When you restore a signed zone, the Grid Master restores it and re-signs its data sets with the original keys, which are also restored. You can also restore the user-defined DS records. The rollover period of the ZSK and KSK starts when the zone is signed after it is restored.

Note that when you restore a deleted zone from recycle bin on the NIOS server, which is created and signed on the Microsoft Server 2012, then all the DNSSEC records will be deleted, except for the DNSKEY records. The DNSKEY records will only be re-synced. The DNSSEC records are read-only and cannot be regenerated using NIOS. You must recreate the zone manually on the Microsoft Server. When you recreate the zone on the Microsoft Server, new keys will be generated. The signed zone, which is restored, and the DNSSEC keys are synced to Microsoft Server. This zone will be seen as an unsigned zone on the Microsoft Server, as NIOS does not trigger the signing zone request for the corresponding zone. For such zones, the 'DNSSEC' label is not displayed and the value for 'Signed' column is 'No'.

To delete a signed zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the check box of the zone you want to delete.
3. Click the Delete icon.
4. Click **Yes** to confirm the deletion.

To restore a signed zone:

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

ABOUT HSM SIGNING

You can integrate a Grid with third-party, network-attached Hardware Security Modules (HSMs) for secure private key storage and generation, and zone-signing off-loading. Infoblox appliances support integration with either SafeNet HSMs or Thales HSMs. When using a network-attached HSM, you can provide tight physical access control, allowing only selected security personnel to physically access the HSM that stores the DNSSEC keys. When you enable this feature, the HSM performs DNSSEC zone signing, key generation, and key safe keeping.

Note that if you migrate from using the Grid Master to HSMs, HSM signing starts at the next key rollover.

Only a superuser can configure this feature. To configure HSM signing in a Grid, do the following:

1. Create the HSM group and add HSMs to the group. You can create either a SafeNet HSM group or a Thales HSM group. You can use only one group at a time. After you add the HSM group, the Add icon and Add button in the Toolbar are greyed out.
 - For information on adding an SafeNet HSM group, see [Configuring a SafeNet HSM Device](#) on page 750.
 - For information on adding a Thales HSM group, see [Adding and Managing a Thales HSM Group](#) on page 752.

Note that if you delete an HSM or an HSM group, it is permanently deleted. It is not stored in the Recycle Bin.

2. Enable HSM signing. For information, see [Enabling HSM Signing](#) on page 753.

After you enable this feature, you can monitor the HSM group, as described in [Monitoring the HSM Group](#) on page 753. In addition, the Grid sends SNMP traps when zone signing succeeds or fails. For information about these traps, see [Processing and Software Failure Traps](#) on page 1071.

Note that NIOS does not provide key life cycle management functions; these are handled by the HSM and must be configured via the HSM's administrative interface to adhere to corporate policies on key management. The keys (ZSK and KSK) used for DNSSEC are stored securely on the HSM and are not deleted by NIOS when the key is no longer required by the DNSSEC function. However, references to the keys are removed from the appliance.

Configuring a SafeNet HSM Device

You can integrate a Grid with a SafeNet HSM group. The SafeNet HSM group can contain either SafeNet Luna SA 4 or SafeNet Luna SA 5 devices in standalone or HA mode; the group cannot contain a mix of both models. You must first configure each HSM device, as described in [Configuring a SafeNet HSM Device](#); and then create the group and add the devices to the group, as described in [Adding a SafeNet HSM Group](#).

Configuring a SafeNet HSM Device

Do the following for each SafeNet HSM device that you are adding to the group:

1. On the Grid, generate a client certificate for the Grid Master and Grid Master candidate. For information, see [About Client Certificates](#) on page 55.
2. On the SafeNet HSM, do the following:
 - Assign the Grid Master and Grid Master candidate to a partition on the HSM to avoid any service interruptions, in case the Grid Master candidate is promoted to Grid Master.
 - Upload the certificates of the Grid Master and Grid Master candidate to the HSM and register the certificates in the HSM's list of clients. The certificates of the Grid Master and Grid Master candidate are linked to their IP addresses. Therefore, if any of their IP addresses change, you must generate a new client certificate and register it with the HSM.

Note that if the HSM is configured and you replace an appliance that was a Grid Master or Grid Master candidate and you backed up the database of the old appliance and restored it on the replacement appliance, the certificates remain intact. Therefore, you do not need to regenerate a new certificate for the replacement, as long as the IP address does not change.

- Download the HSM certificate.

Note: Make sure that the common name used in the certificates is distinct when you configure HSM servers in HA mode.

For additional information, refer to your SafeNet HSM documentation.

Adding a SafeNet HSM Group

When you configure a SafeNet HSM group, add the SafeNet HSM devices to the group and upload their certificates to the Grid. You can add only one HSM group. To add a SafeNet HSM Group:

1. From the **Grid** tab, select the **HSM Group** tab.
2. Click the Add drop-down list and select **HSM SafeNet Group**.
3. In the *Add HSM SafeNet Group* wizard, complete the following and click **Next**:
 - **Name:** Enter a name for the HSM group.
 - **Partition Password:** Enter the partition password, and re-enter it in the **Confirm Partition Password** field.
 - **Version:** Select the SafeNet HSM version, which is either **LUNA SA 4** or **LUNA SA 5**.
 - **Comment:** You can enter additional information about the HSM.
4. Click the Add icon to add a SafeNet HSM device, and complete the following:
 - **Name or IP Address:** Enter the hostname or IP address of the HSM device.
 - **Partition SN:** Enter the partition serial number (PSN) of the HSM. The **Partition ID** field automatically displays the ID after the configuration is saved and the appliance has successfully connected to the device.
 - **Disabled:** Select this check box to disable use of this HSM.
 - **Server Certificate:** Upload the certificate of the SafeNet HSM.
5. Save the configuration.

After you add the HSM group, the Add icon and Add button in the Toolbar are greyed out. Note that if the HSM is configured in FIPS 140-2 compliant mode, certain key algorithms and key sizes are disallowed. Requests for those key algorithms or key sizes result in an error. The following algorithms are FIPS 140-2 compliant: DSA, DSA/NSEC3, RSA/SHA1, RSA/SHA1/NSEC3, RSA/SHA-256, and RSA/SHA-512. For additional information about selecting key algorithms, see [About the DNSKEY Algorithm](#) on page 743.

You can verify whether the Grid Master candidate is properly registered with the HSM by navigating to the **Grid -> Grid Manager -> Members** page. It's Status icon is yellow if it is not registered with the HSM.

If DNS service is enabled, you can also verify whether the Grid Master was able to contact the SafeNet HSMs by navigating to the **Data Management > DNS > Members** page. If the Grid Master status is yellow, check the status of the HSMs in the **Grid > HSM Group** page. (For more information, see [Monitoring the HSM Group](#) on page 753.) If the status is not green, check the configuration of the HSMs and restart the DNS service.

Adding and Managing a Thales HSM Group

On the Thales HSM, configure the Grid Master and Grid Master candidate as HSM clients. Enroll the IP addresses of both the Grid Master and Grid Master candidate to avoid any service interruptions, in case the Grid Master candidate is promoted to Grid Master. If the Grid Master and Grid Master candidates are HA pairs, you must enroll their VIPs. In addition, you must also set up client cooperation to allow both the Grid Master and Grid Master candidate access to the Remote File Server (RFS). Note that anytime you add a new Grid Master candidate, you must enroll its IP address and set up a client cooperation to allow it access to the RFS. For more information on these procedures, refer to your HSM documentation.

Note that DSA cannot be used as the DNSSEC cryptographic algorithm for Thales HSMs. Therefore, migrating to Thales HSMs is not allowed if the Grid Master uses DSA as the DNSSEC cryptographic algorithm.

You can create one Thales HSM group in the Grid, and then add HSMs to the group. The appliance tries to connect to each of the HSMs in the order that they are listed.

To add a Thales HSM group:

1. From the **Grid** tab, select the **HSM Group** tab and click the Add icon.
2. In the **Add HSM Group** wizard complete the following, and then click **Next**:
 - **Name**: Enter a name for the HSM group.
 - **Protection**: Select the level of protection that the HSM group uses for the DNSSEC key data.
 - **Module**: Select this if the HSM group uses a module-protected key. You do not have to enter a password phrase for this type of key.
 - **Softcard**: Select this if the HSM group uses a softcard-protected key. You must then specify the card name and password.
 - **Card Name**: Enter a name for the softcard.
 - **Password Phrase**: Enter the password and re-enter it in the **Confirm Password Phrase** field.
 - **RFS IP Address**: Enter the remote file server (RFS) IP address. Note that you must ensure that you enter a valid RFS IP address for the Security World. Validation is limited to IP address checking. Infoblox recommends that you use **Test HSM Group** to check the HSM group configuration before proceeding.
 - **RFS Port**: Specify the port of the RFS.
 - **Comment**: Optionally, enter additional information about the group.
3. To add modules to the group, click the Add icon and complete the following:
 - **Remote IP**: Enter the IP address of the HSM.
 - **Remote Port**: Specify the destination port on the HSM. The firewall must be configured to allow connection to this port.
 - **Disabled**: Select this check box to disable use of this HSM.
 - **Keyhash**: Enter the keyhash, which is displayed on the console of the HSM. It can be obtained through an out of band mechanism from the HSM administrator. Note that the appliance validates the keyhash. If the entry is correct, the appliance displays the Electronic Serial Number (ESN) of the HSM when the editor is next launched. If the keyhash is incorrect, the appliance does not connect to the HSM.
 - **ESN**: This is a read-only field that displays the ESN of the HSM after you save the configuration and relaunch the editor. Infoblox strongly recommends that you verify the ESN displayed by the appliance with the one obtained from the HSM administrator to ensure that the appliance is communicating with the correct HSM.
4. Save the configuration.

Monitoring the HSM Group

You can monitor the status of the HSM group and of individual modules in the group by navigating to the **Grid** tab > **HSM Group** panel. To view the status of each HSM, click the arrow beside the group name. This panel displays the following information:

- **Name:** The name of the HSM group or module.
- **Status:** The HSM group status displays the status for all the HSMs in the group. The status icon can be one of the following:

Green: All the HSMs in the group are functioning properly.

Yellow: At least one HSM in the group is not functioning properly.

Red: All the HSMs in the group are not functioning properly.

Black: The status of the HSM devices is unknown.

The status icon for each HSM can be one of the following:

Green: The HSM is functioning properly. For SafeNet Luna SA 5 devices, the status icon can also display **x% used** which refers to the storage capacity of the HSM partition that is assigned to the Grid. Note that when the capacity reaches 100%, new zone signings and key rollovers for existing zones cannot be performed.

Red: The HSM is not functioning properly. For a SafeNet HSM, this indicates that the Grid Master was able to connect to the HSM, but no partition was assigned to the Grid Master.

Black: The status of the HSM device is unknown.

- **FIPS:** This applies to a SafeNet HSM only. It indicates if the HSM is in FIPS compliant mode.
- **Comment:** Any comments that were entered about the HSM group.

You can also do the following in this tab:

- Sort the data in ascending or descending order by column.
- Print and export the data in this tab.

Enabling HSM Signing

When you enable HSM signing, the HSM starts generating the DNSSEC keys at the next key rollover. For information about key rollovers, see [About Key Rollovers](#) on page 744. You can enable this feature at the Grid level only.

To enable HSM signing:

1. From the **Data Management** tab -> **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, Click **Toggle Expert Mode**, if the editor is in Basic mode, and then select the **DNSSEC** tab.
3. In the **DNSSEC** tab, select the **Enable DNSSEC** check box, if it is not selected, and then select the **HSM Signing** check box.
4. Complete the other fields described in [Configuring DNSSEC Parameters](#) on page 745. Note that Thales HSMs do not support DSA.
5. Save the configuration.

Testing the HSM Group

After you configure the HSM group, you can test the HSM signing functionality of the group. Click **Test HSM Group** in the Toolbar, and then click **Yes** when the confirmation dialog displays. The appliance then executes the command to perform a signing test. The feedback panel displays the status of the test in the Grid Manager feedback panel.

Synchronizing the HSM Group

You can click **Resync HSM Group** in the Toolbar to do any of the following:

- For a Thales HSM group, if the RFS security settings change use this function to have the appliance perform an RFS synchronization.
- For a SafeNet HSM group, use this function to synchronize the keys of the HSM members in the group.

CONFIGURING GRID MEMBERS TO SUPPORT DNSSEC AS SECONDARY SERVERS

Any Infoblox Grid member can function as a secondary server for DNSSEC signed zones. It can receive transfers of signed zones from the Grid Master or an external primary server, and from other secondary servers. It can also respond to queries for DNS data in DNSSEC signed zones for which it is a secondary server.

Configuring a Secondary Server for Signed Zones

The following are the tasks to configure an appliance as a secondary server for signed zones:

1. Enable DNSSEC on the appliance. For information, see [Enabling DNSSEC](#) on page 743.
2. Configure the appliance as a secondary server for the zone. For information, see [Specifying a Secondary Server](#) on page 626.
3. If the primary server for the signed zone is external, then you must allow zone transfers to the secondary server. For information, see [Enabling Zone Transfers](#) on page 583. If the primary server is the Grid Master, then the secondary server receives data through the Grid replication process by default.

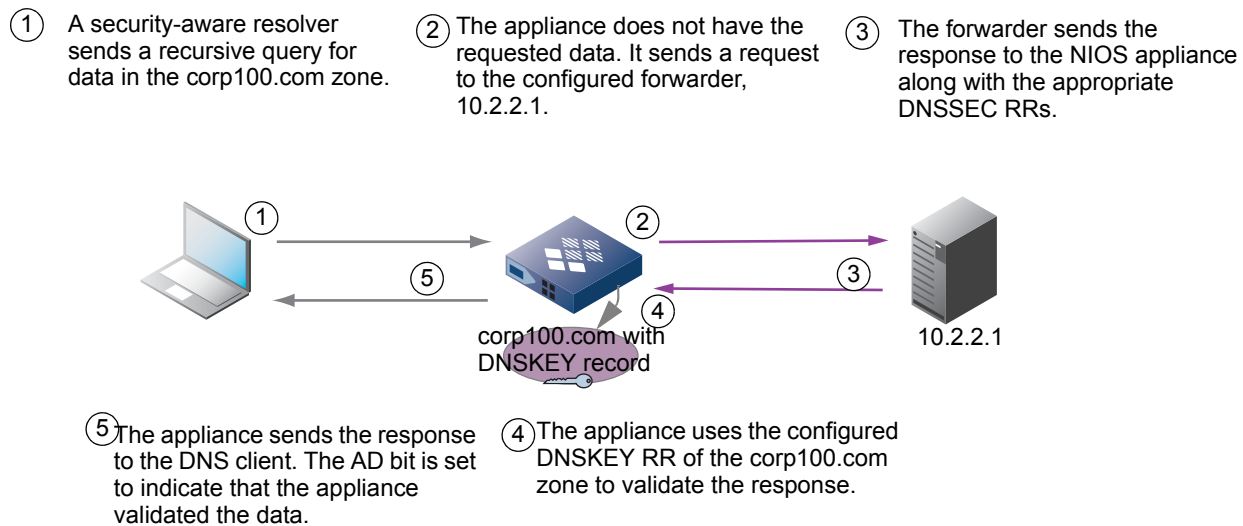
CONFIGURING RECURSION AND VALIDATION FOR SIGNED ZONES

When you enable recursion on a Grid member and it receives a recursive query for DNS data it does not have, it queries remote name servers that you specified in the *Grid DNS Properties* or *Member DNS Properties* editor. It then includes the DNSSEC data it retrieved through recursion in its responses to clients that requested DNSSEC RRs. You can enable the appliance to validate the responses of these servers for certain zones. On the appliance, you specify the zones to validate and configure their DNSKEY records as trust anchors. When the appliance validates a response for a zone configured with a trust anchor or for any of its child zones, the appliance starts with the DNSKEY that you configured and proceeds recursively down the DNS tree.

In the example shown in [Figure 21.5](#), the following was configured on the NIOS appliance:

- Forwarder with the following IP address: 10.2.2.1
- Recursion was enabled
- DNSSEC and validation were enabled
- The corp100.com zone and its DNSKEY record were configured

Figure 21.5



Enabling Recursion and Validation for Signed Zones

The following are the tasks to enable recursion and validate recursively derived data:

1. Enable DNSSEC on the appliance. For information, see [Enabling DNSSEC](#) on page 743.
2. Enable validation and configure the trust anchor of each signed zone. For information, see [Enabling DNSSEC Validation](#) on page 755. You must configure at least one trusted DNSKEY RR.
3. Enable recursion on the appliance. For information, see [Enabling Recursive Queries](#) on page 571.
4. Complete any of the following:
 - Configure the forward, delegated, stub or root zones for the signed zones. For information, see [Configuring Delegated, Forward, and Stub Zones](#) on page 638 and [Creating a Root Zone](#) on page 620.
 - Configure global forwarders and custom root name servers, if needed. For information, see [Using Forwarders](#) on page 569 and [About Root Name Servers](#) on page 587.

Enabling DNSSEC Validation

To configure trust anchors and enable Infoblox name servers to validate responses:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.
Member: From the **Data Management** tab, select the **Members** tab -> *membercheck* box and click the Edit icon.
DNS View: From the **Data Management** tab, select the **Zones** tab -> *dns_view* check box and click the Edit icon.
 To override an inherited property, click **Override** next to the property to enable the configuration.
2. In the editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **DNSSEC**.
4. In the **DNSSEC** tab, complete the following:
 - **Enable DNSSEC validation:** If you allow the appliance to respond to recursive queries, you can select this check box to enable the appliance to validate responses to recursive queries for domains that you specify. You must configure the DNSKEY RR of each domain that you specify.

- **Accept expired signatures:** Click this check box to enable the appliance to accept responses with signatures that have expired. Though enabling this feature might be necessary to work temporarily with zones that have not had their signatures updated in a timely fashion, note that it could also increase the vulnerability of your network to replay attacks.
- **Trust Anchors:** Configure the DNSKEY record that holds the KSK as a trust anchor for each zone for which the Grid member returns validated data. Click the Add icon and complete the following:
 - **Zone:** Enter the FQDN of the domain for which the member validates responses to recursive queries.
 - **Secure Entry Point (SEP):** This check box is enabled by default to indicate that you are configuring a KSK.
 - **Algorithm:** Select the algorithm of the DNSKEY record: RSA/SHA1(5), DSA (3), DSA/NSEC3 (6), RSA/MD5 (1), RSA/SHA1/NSEC3 (7), RSA/SHA-256 (8), or RSA-SHA-512 (10). This must be the same algorithm that was used to generate the keys that were used to sign the zones.
 - **Public Key:** Paste the key into this text box. You can use either of the following commands to retrieve the key:
 - `dig . dnskey +multiline`
The above command retrieves root zone keys and is the only public key you require for full chain of trust validation.
 - `dig [@server_address] <zone> dnskey +multiline +dnssec`
The above command retrieves public keys from the zone you specify on the server and can be used if the parent zone is not signed.

Note that the aforementioned command provides you with a key you need to cross validate against other servers to ensure you have an identical key.

As an alternative, you can use <http://data.iana.org/root-anchors/> to retrieve signed public keys. You can find the trust anchors in formats like XML and CSR. For more information, refer to <http://data.iana.org/root-anchors/draft-icann-dnssec-trust-anchor.txt>.

5. Save the configuration and click **Restart** if it appears at the top of the screen.



Chapter 22 Configuring IP Routing Options

You can configure multiple IP addresses and enable anycast addressing on the loopback interface of the NIOS appliance, allowing the appliance to function in different network deployments.

Configuring non-anycast IP addresses on the loopback interface assists in server migration and network address change. Configuring anycast addresses on the appliance allows you to add redundancy and improve reliability for DNS services. You can use OSPF (Open Shortest Path First), BGP (Border Gateway Protocol), or both, as the routing protocol for anycast advertising.

This chapter contains the following sections:

- [*Using the Loopback Interface*](#) on page 758
- [*Configuring IP Addresses on the Loopback Interface*](#) on page 759
 - [*Advertising Loopback Addresses to the Network*](#) on page 760
- [*About Anycast Addressing for DNS*](#) on page 761
 - [*Configuring Anycast Addresses*](#) on page 762
- [*IP Routing Options*](#) on page 763
 - [*Anycast and OSPF*](#) on page 764
 - [*Configuring OSPF on the NIOS Appliance*](#) on page 765
 - [*Anycast and BGP4*](#) on page 767
 - [*Configuring BGP in the NIOS Appliance*](#) on page 769

USING THE LOOPBACK INTERFACE

The loopback interface is a virtual network interface on the appliance. You can do the following on the loopback interface:

- Configure IP addresses to consolidate DNS servers for migration purposes. For information, see [Configuring IP Addresses on the Loopback Interface](#) on page 759.
- Add anycast addresses to improve the reliability and performance of DNS services in multiple locations. For information, see [About Anycast Addressing for DNS](#) on page 761.
- Separate DNS traffic by assigning an IP address as the source port for DNS queries. For information, see [Specifying Source Ports](#) on page 563.

When you use the loopback interface for anycast addressing, the upstream and neighboring routers can continue to advertise anycast addresses without being affected by hardware malfunctions.

To configure non-anycast addresses on the loopback interface, complete the following:

1. Add IP addresses to the loopback interface. For information, see [Configuring IP Addresses on the Loopback Interface](#) on page 759.
2. Enable DNS services on the loopback addresses. For information, see [Specifying Port Settings for DNS](#) on page 562 and its subtopic, [Specifying Source Ports](#).

To configure DNS anycast addresses and their advertising protocols, complete the following:

1. Add anycast addresses to the loopback interface. For information, see [Configuring Anycast Addresses](#) on page 762.
2. Configure anycast addressing protocols. For information, see [Configuring OSPF on the NIOS Appliance](#) on page 765 and [Configuring BGP in the NIOS Appliance](#) on page 769. This is the primary application for routing protocols in the NIOS appliance.
3. Enable the DNS anycast addresses. For information, see [Specifying Port Settings for DNS](#) on page 562 and its subtopic, [Specifying Source Ports](#).

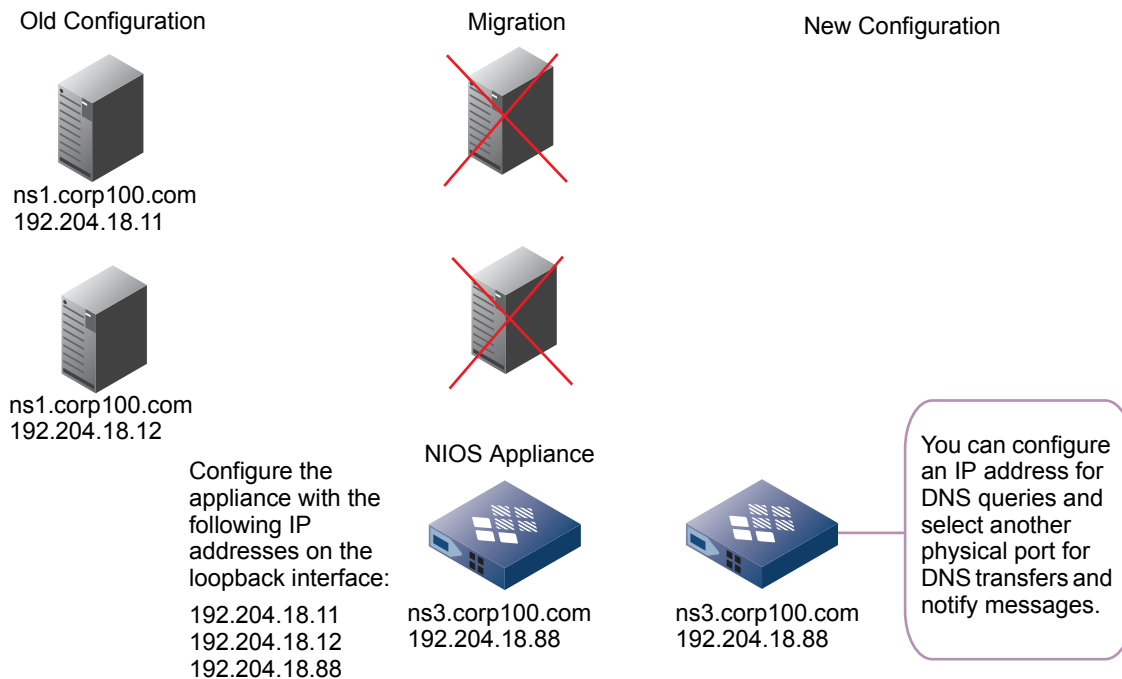
To separate DNS queries from DNS transfers and notify messages, complete the following:

1. Add an IP address of the source port for DNS queries. For information, see [Configuring IP Addresses on the Loopback Interface](#) on page 759.
2. Select the source IP for DNS queries. For information, see [Specifying Source Ports](#) on page 563.

CONFIGURING IP ADDRESSES ON THE LOOPBACK INTERFACE

You can configure IP addresses on the loopback interface to minimize service downtime during a server migration. As illustrated in [Figure 22.1](#), you have two existing DNS servers (ns1.corp100.com 192.204.18.11 and ns2.corp100.com 192.204.18.12) and you want to replace these servers with a new one (ns3.corp100.com 192.204.18.88). The migration takes a few weeks and you want DNS services to be available on all three addresses during the migration. You can add all three IP addresses to the loopback interface of a NIOS appliance, and then configure the appliance to provide DNS services on all addresses. After the server migration, you can shut down the old servers and use the new one for services.

Figure 22.1 DNS Server Migration Using the Loopback Interface



You can also add an IP address that is used solely for DNS queries, to separate the DNS traffic. You first add an IP address you want to use for DNS queries on the loopback interface. You then configure the appliance to listen for DNS queries solely on this address. For information, see [Specifying Source Ports](#) on page 563.

When you configure non-anycast addresses on the loopback interface, ensure that you establish a static route between the appliance and the router so queries to these addresses are routed correctly. For information, see [Advertising Loopback Addresses to the Network](#) on page 760.

To configure an IP address on the loopback interface:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box -> Edit icon.
2. In the *Grid Member Properties Editor*, select the **Network** tab -> **Basic** tab.
You can add an IPv4 or IPv6 address on the loopback. You define each type in their own table.
3. Under **IPv4 Additional Address (loopback)** or **IPv6 Additional Address (loopback)**, click Add and then **Additional Address (Loopback)**.

The appliance adds a row to the table. Complete the following:

- **Interface:** Displays **Additional Address (loopback)**. This value cannot be modified.
- **Address:** Enter the IP address you want to add to the loopback interface.

- **Subnet Mask:** You cannot change the netmask of the loopback interface. It is set to 255.255.255.255, or /32; for an IPv6 address the mask is set to 128 and cannot be modified.

Note: You cannot configure the gateway address and port settings.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

To add multiple IP addresses on the loopback interface, repeat the steps for each IP address.

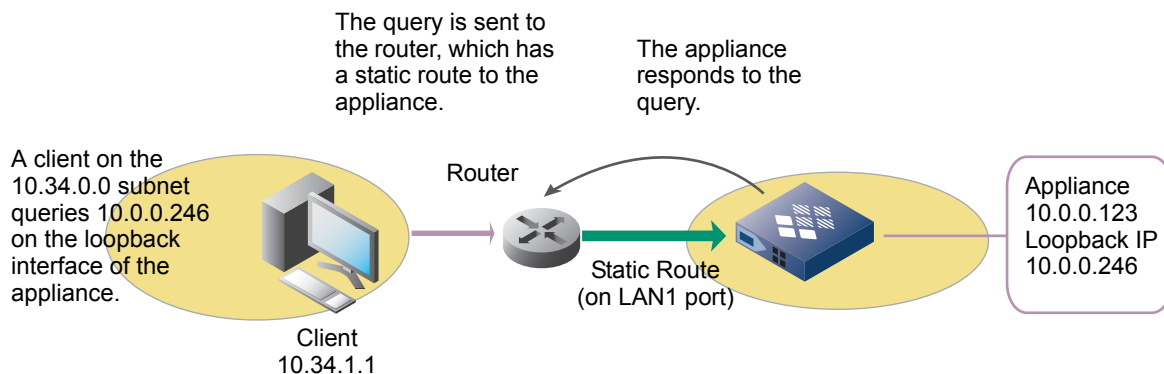
Note: If you are configuring the loopback interface on a Grid Master, the Grid is temporarily disrupted upon saving the configuration and restarting services on the appliance. The Grid reconnects automatically and the appliance regains the role as Grid Master after a short delay.

Advertising Loopback Addresses to the Network

Advertising IP addresses on the loopback interface relies on the upstream router to populate routes to the loopback interface. As illustrated in [Figure 22.2](#), when a client on a different subnet queries an IP address on the loopback interface, it sends the request to the router. If the IP address on the loopback interface is not advertised to the router, the request cannot reach the appliance. Therefore, when you configure non-anycast addresses on the loopback interface, or if OSPF or BGP is not configured within your network, you must configure the upstream router to reach the NIOS appliance through a static route on the LAN1 interface.

Note that when an appliance is configured for both authoritative and recursive queries, you should connect your internet interface through the LAN1 port to allow for maximum flexibility while using auxiliary LAN2 and MGMT ports. Consult with your network administrator for information about configuring static routes from the router to the additional IP addresses on the loopback interface.

Figure 22.2 Static Route for Loopback IP Addresses



When you configure DNS anycast addresses on the loopback interface, you can select OSPF, BGP, or both, to advertise the addresses to upstream and neighboring routers, without establishing a static route. For information, see [About Anycast Addressing for DNS](#) on page 761.

ABOUT ANYCAST ADDRESSING FOR DNS

Note: This feature is not supported on vNIOs appliances for Riverbed.

Four types of communications are utilized within an IP network:

- **Unicast** describes a one-to-one network communication between a single sender and a single recipient. The routing protocol determines the path through the network from the sender to the recipient based on the specific protocol or routing scheme. Unicast also describes the address type assigned to the recipient.
- **Multicast** describes a one-to-many network communication between a single sender and a specific group of recipients. All members within the group are intended recipients and each member receives a copy of the data from the sender. Multicast also describes the address type assigned to a group of recipients, used by the routing protocol to determine the path to the group.
- **Broadcast** is similar to multicast, the exception being that data is sent to every possible destination regardless of the groups or subnetwork. There is no specific group of recipients.
- **Anycast** describes a one-to-nearest communication between a single sender and the nearest recipient within a group. The routing protocol chooses one recipient within a target group based on the routing algorithm for the specific protocol, and sends data to that recipient only.

The NIOS appliance provides the following support for DNS anycast addressing:

- You can configure up to 20 anycast IP addresses on the loopback interface of each Grid member.
- Anycast IP addresses can be in IPv4 or IPv6 format. For all anycast IP addresses, the subnet mask value is always set to /32 for an IPv4 anycast IP or 128 for a 128-bit IPv6 address. These values are separate and distinct from the IP configuration on the NIOS appliance LAN port.
- The appliance advertises routing information of the anycast addresses through OSPF or BGP, or (seldom) both, depending on the deployment. Routers use the configured routing protocols to determine the best path to the nearest server. The appliance advertises the route information to the upstream or neighboring router, a router that forwards data on the network link and determines the forwarding path to destinations. For information, see [IP Routing Options](#) on page 763.
- The appliance advertises and withdraws route information based on reachability information to DNS servers sent by the IP route advertisements.
- When you configure DNS anycast addressing on an appliance and use it as an NTP server, the appliance can answer NTP requests through the anycast IP address. For information about how to configure an appliance as an NTP server, see [Configuring a NIOS Appliance as an NTP Server](#) on page 320.

Anycast addressing for DNS provides the following benefits:

- **Improved Reliability:** Anycast provides improved reliability because DNS queries are sent to an anycast IP address that is defined on multiple DNS servers in the NIOS Grid. If the nearest server somehow goes offline, then the router forwards the request to the next nearest DNS server advertising the target anycast IP address (see [Figure 22.4](#) for an example).
- **Load Distribution:** Anycast distributes the load across multiple DNS servers based on network topology.
- **Improved Performance:** The NIOS appliance uses OSPF or BGP, depending on your configuration, to advertise anycast routing information to the upstream and neighboring routers. The routers determine the best route to the nearest DNS server. Anycast enables the queries to reach the nearest server more quickly, providing faster responses to DNS queries.

Note: For more information about anycast addressing, refer to *RFC 1546 “Host Anycasting Service”*.

Configuring Anycast Addresses

Note: Anycast addressing is supported on loopback interfaces on the NIOS appliance.

IP configuration must be defined on the LAN1 interface before configuring DNS anycast addresses. Before creating IPv6 anycast IPs on the loopback interface, IPv6 must be enabled and configured on the LAN1 interface for the NIOS appliance, including the correct IPv6 gateway IP address.

To enable and configure anycast addressing:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box -> Edit icon.
2. Click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Anycast** tab.
4. Click the Add icon and choose **IPv4 Address** or **IPv6 Address**.
5. In the **Anycast Interfaces** list, enter the values or select the options for the new entry:
 - **Anycast Interface:** Anycast addressing is supported on the **loopback** only. This value is filled in automatically.
 - **Address:** Enter the IP address you want to assign as the anycast address to the loopback interface. Specify an IPv4 Address or an IPv6 Address based on the chosen type of address.
Subnet Mask: You cannot change the subnet mask of a loopback interface. The netmask is automatically set to 255.255.255.255, or /32; or 128 for IPv6.
 - **OSPF:** Select if you want the appliance to use OSPF to advertise the anycast address, and if necessary configure the OSPF settings. For information, see [Configuring OSPF on the NIOS Appliance](#) on page 765. IPv4 and IPv6 options are configurable for this protocol.
 - **BGP:** Select this if you want the appliance to use BGP to advertise the anycast address, and then configure the BGP settings.

Note: You must configure at least one routing method for DNS anycast. You can configure OSPF, BGP, or both (in most cases only one protocol will be used). The appliance cannot save the anycast address if you do not complete at least one routing configuration. Anycast cannot be used without dynamic routing.

- **Comments:** Enter a text string to help identify this interface and IP address.
6. *If using OSPF for the current appliance:* Under **OSPF Area Configuration**, click the Add icon. A new configuration block appears in the properties editor.
 - Enter the values for the OSPF configuration as described in the section [Configuring OSPF on the NIOS Appliance](#).
 - Click the **Add** down arrow icon in the **OSPF Area Configuration** section. The new OSPF configuration is saved into a table.
 7. *If using BGP for the current appliance:* In the properties editor, scroll down to the configuration block for **BGP Configuration**. For information, see [Configuring BGP in the NIOS Appliance](#) on page 769.
 - In the **ASN** field, enter the Autonomous System ID number in which the NIOS appliance resides.
 - If necessary, modify the **BGP Timer Keep Alive** and **Hold Down** values. In most circumstances these values should be left at their defaults. Check your network's defined policies for the desired values if necessary.
 - Click the Add icon.
 - Enter the **Neighbor Router** IP address. This can be an IPv4 address or an IPv6 address.
 - Enter the **Remote ASN** (Autonomous System ID number) for the adjacent router.
 8. Save the configuration. The system will warn that you must restart the appliance services in order to use the new configuration.
 9. Log back in to the appliance.

10. From the **Data Management** tab, select the **DNS** tab -> **Members/Servers** tab -> *Grid_member* check box -> Edit icon.
11. Select **Toggle Advanced Mode** (if necessary), click **General** and the **Advanced** tab.
12. Under **Listen on these additional IP addresses**, click the Add button. The list of one or more previously created IPv4 and IPv6 addresses for the loopback interfaces (created in Step 4) appear in this table. (If the Add button is not active here, this indicates that you have not configured any loopback interfaces with their IP addresses.)

Note: Should you need to configure other DNS properties on this page, see the topics in [Configuring DNS Services](#).

13. Click **Save and Close**.

Configured anycast interfaces are now enabled to carry DNS traffic. For further information, see [Specifying Source Ports](#) on page 563.

Best Practices for Configuring Anycast Addresses

Infoblox highly recommends that you do the following before you configure an Anycast address:


- Enable the Anycast feature in the NIOS application.
- Install a valid DNS license, enable DNS and ensure that the DNS service is active.
- When you configure OSPF or OSPFv6, ensure that the OSPF monitor runs every four seconds.
- You must configure an IP address on the loopback interface.

IP ROUTING OPTIONS

For routing purposes, the internet is divided into ASs (Autonomous Systems). Data is routed within an AS using an IGP (Interior Gateway Protocol) and routed between different ASs using an EGP (Exterior Gateway Protocol). NIOS appliances support OSPFv2 (for IPv4) and OSPFv3 (for IPv6) for a routing IGP, and BGP4 to advertise DNS anycast addresses in the larger internetwork.

As noted in the section [Configuring Anycast Addresses](#), you configure OSPF or BGP4 to advertise anycast addresses, which configured on the loopback interface of NIOS appliances. Use of either protocol depends on the network topology, based on whether the advertisements will propagate only within a single AS or between more than one AS. [Figure 22.3](#) shows a simplified example.

Figure 22.3 OSPF and BGP Routing Example



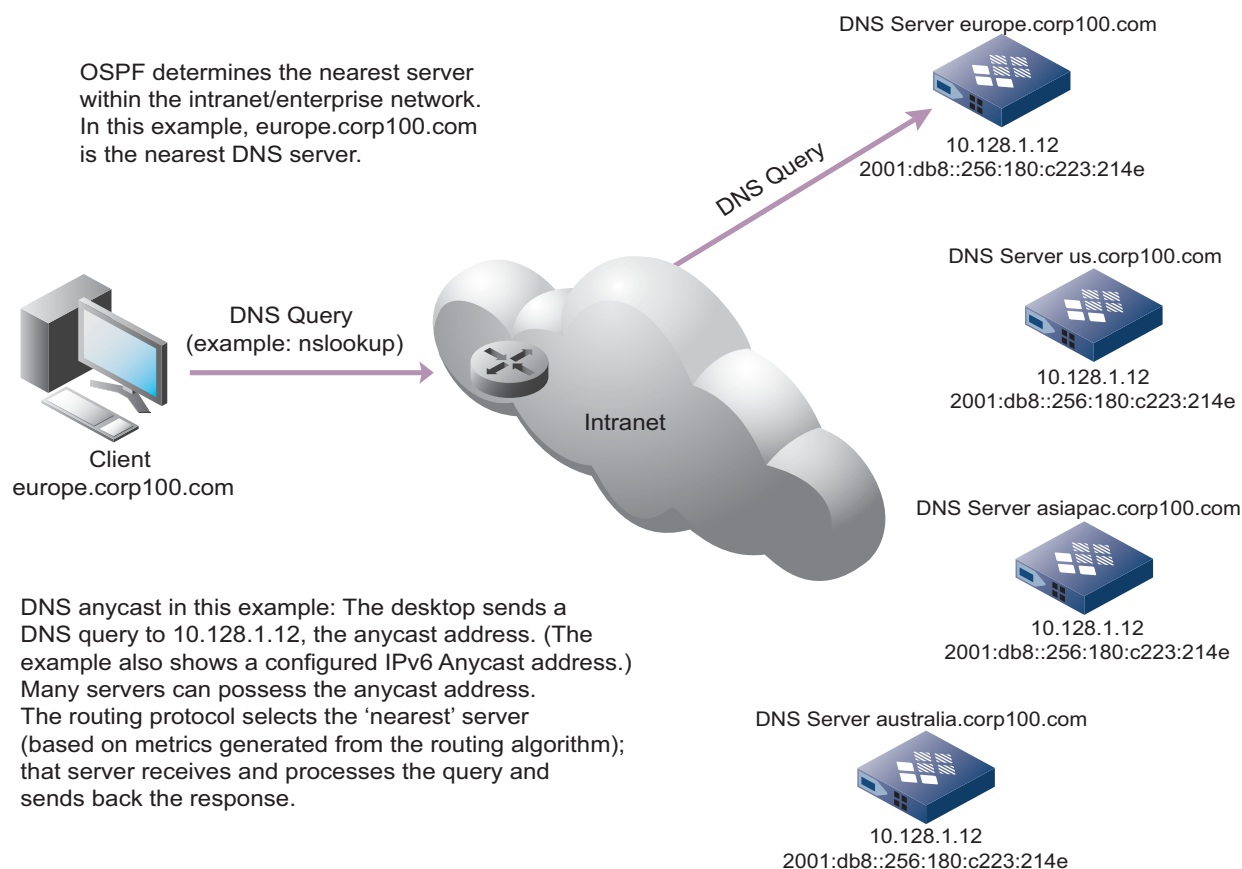
Within each AS, OSPF is the protocol used to forward anycast advertisements. Between ASs, BGP is the protocol selected to advertise anycast addresses. Using this technique, DNS servers in diverse locations can operate together to ensure continuous service.

Anycast and OSPF

NIOS appliances can use the OSPF routing protocol to advertise routes for DNS anycast addresses to an upstream router within the autonomous system. The upstream router uses the OSPF advertisement to determine the nearest DNS server from a group of servers within the internetwork. In practice, the NIOS appliance relies upon OSPF to determine the best route for DNS queries to take to the nearest DNS server. The upstream router then forwards the query to the chosen DNS server.

As illustrated in [Figure 22.4](#), to enable anycast for DNS queries, you configure two or more DNS servers within the AS routing domain with the same anycast address on their loopback interfaces. When you select OSPF as the routing protocol, the upstream router determines the nearest server within the group of servers configured with that anycast address. (The “nearest” DNS server may not necessarily be the geographically closest DNS server; it is the DNS server with the lowest cost associated with its reachability from the current node. This is calculated through the OSPF routing algorithm, a discussion of which is far beyond the scope of this manual.) The nearest DNS server configured with the correct anycast address then responds to the DNS query. In the case where the nearest server becomes unavailable, the next nearest server responds to the query. OSPF anycast provides a dynamically routed failover to ensure that DNS can always resolve client requests within the AS. From the client perspective, anycasting is transparent and the group of DNS servers with the anycast address appears to be a single DNS server.

Figure 22.4 Anycast Addressing for DNS Using OSPF



After you configure or change DNS anycast settings, you must restart the DNS services for the settings to take effect. When you enter any OSPF command and wait for the interface to return more information, the appliance disconnects your CLI session after you restart services or make other OSPF configuration changes through Grid Manager. Re-enter your credentials to log back in to the CLI. (For information, refer to the *Infoblox CLI Guide*.)

To enable the appliance to support OSPF and advertising anycast addresses on OSPF from the loopback, you must first configure the LAN1 or LAN1 (VLAN) interface as an OSPF advertising interface. For information about VLAN, see [About Virtual LANs](#) on page 346.

You can also configure authentication for OSPF advertisements to ensure that the routing information received from a neighbor is authentic and the reachability information is accurate. This process can be implemented for OSPF over IPv4 networks but is not supported for IPv6/OSPFv3. For information, see [Configuring OSPF on the NIOS Appliance](#).

Note: For more information about the OSPF routing protocol, refer to *RFC 2328 "OSPFv2"* and *RFC 5340 "OSPF for IPv6"*.

Configuring OSPF on the NIOS Appliance

Note: Use the CLI command `show ospf` or `show ipv6_ospf` to display configuration and statistical information about the OSPF protocol running on the appliance. You can also use the `set ospf` or `set ipv6_ospf` command to write OSPF statistical information to the syslog. In the NIOS appliance, configuration of OSPF is limited to Syslog and the DNS anycast application.

To support DNS anycast and other routing-dependent applications on NIOS appliances, you must first configure the LAN1, LAN1 (VLAN), or HA (for HA pairs only) interface as an OSPF advertising interface, and then assign an area ID on the interface to associate it with a specific OSPF area. The interface advertises the OSPF routing information to the network so that routers can determine the best server to query. For anycasting, the advertising interface sends out routing advertisements about the anycast address into the network out to upstream routers.

To configure the LAN(HA) or LAN1(VLAN) interface to be an OSPF advertising interface, perform the following tasks:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. Select the **Anycast** tab in the *Grid Member Properties* editor.

The **Anycast Interfaces** appear in a table. You can add new anycast interfaces when needed.

3. Click the Add icon of the OSPF Area Configuration table and choose **IPv4 Configuration** or **IPv6 Configuration** to define a new OSPF Area. The OSPF Area Configuration will show a similar set of **Add (IPv4/IPv6) OSPF Area** configuration settings based on the protocol type. Enter the following information to configure the LAN1, LAN1 (VLAN), or HA interface as the OSPF advertising interface:

- **Advertising Interface:** Displays the interface that sends out OSPF routing advertisements. OSPF advertisements are supported on the LAN1, LAN1(VLAN), and the HA interface, depending on whether the appliance is an HA pair.
- **Area ID:** Enter the OSPF area identifier of the network containing the upstream routers, in either an IP address format or a decimal format. All network devices configured with the same OSPF area ID belong to the same OSPF area. The area ID configured on the Grid member must match the area ID of the upstream router configuration. Area ID numbers are defined in the same format for IPv6 and IPv4.
- **Area Type:** Select the type of OSPF area to associate with the advertising interface from the drop-down list. The area type configured on the Grid member must match the area type of the upstream router configuration. The supported area types are described as follows:
 - **Standard:** A standard area has no restrictions on routing advertisements, and connects to the backbone area (Area 0) and accepts both internal and external link-state advertisements.
 - **Stub:** A stub area is an area that does not receive external routes.
 - **Not-so-stubby:** A not-so-stubby area (NSSA) imports autonomous system (AS) external routes and sends them to the backbone, but cannot receive AS external routes from the backbone or other areas.

Note: OSPF for IPv6 (known as OSPFv3) configuration does not support OSPF authentication options.

- **Authentication Type:** Select the authentication method to use to verify OSPF routing advertisements on the interface. The authentication type configured on the Grid member must match the authentication type of the upstream router configuration. The supported authentication types are described as follows:
 - **None:** No authentication for OSPF advertisement.
 - **Simple:** A simple password for OSPF advertisement authentication, in clear text.
 - **MD5:** An MD5 hash algorithm to authenticate OSPF advertisements. This is the most secure option.
- **Authentication Key ID:** Enter the key identifier to use to specify the correct hash algorithm after you select **MD** as your OSPF authentication type. The authentication key ID configured on the Grid member must match the authentication key ID of the upstream router configuration.
- **Authentication Key:** Enter the authentication password to use to verify OSPF advertisements after you select **Simple** or **MD** as your OSPF authentication type. Specify a key string between 1 to 8 characters for Simple authentication, and a string between 1 to 16 characters for MD5 authentication. The authentication key configured on the Grid member must match the authentication key of the upstream router configuration.
- **Cost:** Select one of the following:
 - **Calculate Automatically:** Select this check box to auto generate the cost to associate with the advertising OSPF interface to the appliance. If this check box is not selected, then you specify the cost value explicitly. Calculate the cost as 100,000,000 (reference bandwidth) divided by the interface bandwidth. For example, a 100Mb interface has a cost of 1, and a 10Mb interface has a cost of 10.
 - **Fixed Metric:** Enter the cost to associate with the advertising OSPF interface to the appliance.

- **Hello Interval:** Specify how often to send OSPF hello advertisements out from the appliance interface, in seconds. Specify any number from 1 through 65,535. The default value is 10 seconds. The hello interval configured on the Grid member must match the hello interval of the upstream router configuration.
- **Dead Interval:** Specify how long to wait before declaring that the NIOS appliance is unavailable and down, in seconds. Specify any number from 1 through 65,535. The default value is 40 seconds. The dead interval configured on the Grid member must match the dead interval of the upstream router configuration.
- **Retransmit Interval:** Specify how long to wait before retransmitting OSPF advertisements from the interface, in seconds. Specify any number from 1 through 65,535. The default value is 5 seconds. The retransmit interval configured on the Grid member must match the retransmit interval of the upstream router configuration.
- **Transmit Delay:** Specify how long to wait before sending an advertisement from the interface, in seconds. Specify any number from 1 through 65,535. The default value is 1 second. The transmit interval configured on the Grid member must match the transmit interval of the upstream router configuration.
- Click **Add** to add the interface to the table.

The **Cost**, **Hello Interval**, **Dead Interval**, **Retransmit Interval** and **Transmit Delay** settings can be configured for IPv6 deployments. OSPF authentication is not supported for IPv6 on the NIOS platform.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing OSPF

- OSPF advertises the route when the DNS service starts. The `start dns` command creates an interface and starts the OSPF daemon.
- OSPF stops advertising the route when the DNS service stops. The `stop dns` command stops the OSPF daemon and deletes the interface.
- The NIOS application does not support a route flap. For example, temporary DNS downtime such as restart, does not stop or re-instate the OSPF advertisement.
- The OSPF advertisement stops if DNS service is down for more than 40 seconds.

Anycast and BGP4


Note: Use the CLI command `show bgp` or `show ipv6_bgp` to display configuration and statistical information about the Border Gateway Protocol running on the appliance. You can also use the `set bgp` command to write OSPF statistical information to the syslog. In the NIOS appliance, configuration of BGP is limited to Syslog and the DNS anycast application.

BGP4 (henceforth referred to as BGP) is designed to distribute routing information among ASs and exchange routing and reachability information with other BGP systems using a destination-based forwarding paradigm. Unlike OSPF, which calculates routes within a single AS, BGP is a vector routing protocol that distributes routing information among different ASs. A unique ASN (autonomous system number) is allocated to each AS to identify the individual network in BGP routing. A BGP session between two BGP peers is an eBGP (external BGP) session if the BGP peers are in different ASs. A BGP session between two BGP peers is an iBGP (internal BGP) session if the BGP peers are in the same AS.

BGP configuration enables large enterprises using BGP as the internetworking protocol to provide resilient DNS services using the Infoblox solution. While BGP is mostly used by ISPs, it is also used in larger enterprise environments that must interconnect networks that span geographical and administrative boundaries. In these environments, it is required to use BGP to advertise anycast routes. Using BGP allows the appliance to advertise DNS anycast addresses to neighboring routers across multiple ASs that also use BGP as their routing protocols.

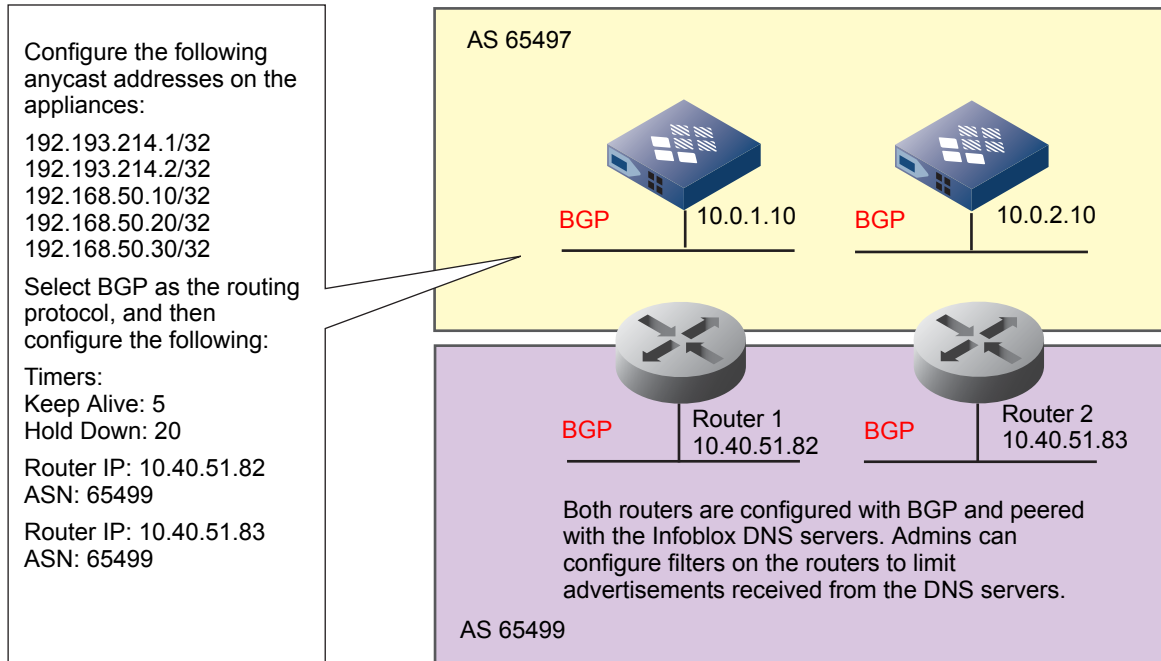
As illustrated in [Figure 22.5](#), to enable anycast for DNS queries among three different networks that span different geographical regions, you can configure two DNS servers with the same DNS anycast addresses in the AS 65497 network. Since other network routers in AS 65498 and AS 65499 also use BGP as the routing protocol, the DNS anycast addresses can be advertised across these networks.

Figure 22.5 Anycast Addressing for DNS using BGP



To enable DNS anycast addressing across different ASs, you configure BGP as the routing protocol on the NIOS appliance. (As illustrated in [Figure 22.6](#), the AS 65497 network contains the Infoblox DNS anycast servers, and the AS 65499 network contains Router 1 and 2. The routers use BGP and are peered with the DNS servers. You can configure anycast addressing on the loopback interface of the DNS servers and select BGP as the protocol to advertise the anycast addresses to Router 1 and 2 in AS 65499. For information, see [Configuring Anycast Addresses](#) on page 762. Once you have configured the DNS servers, the appliances automatically add filters on the advertising interfaces to limit the advertisements to the configured anycast IP addresses. Similarly, BGP filters are applied to ensure that the DNS servers only receive default route advertisements from the neighboring routers.

Figure 22.6 Anycast and BGP Configuration on Infoblox Appliances



BGP uses timers to determine how often the appliance sends keepalive and update messages, and when to declare a neighboring router out of service. You can configure the time intervals for these timers. For information, see [Configuring BGP in the NIOS Appliance](#) on page 769.

The BGP protocol service is automatically configured to send SNMP queries about BGP runtime data. The appliance sends SNMP traps to its neighboring routers when it encounters issues with the protocol. BGP is configured to send SNMP traps as defined in *RFC4273 Definitions of Managed Objects for BGP-4*. You must enable and configure the SNMP trap receiver on the Grid member for the member to send SNMP traps. For information, see [SNMP MIB Hierarchy](#) on page 1048.

You can use the `set bgp` command to set the verbosity levels of the BGP routing service. The appliance writes BGP statistical information to the syslog. After you configure the settings, you must restart the DNS services for the settings to take effect. For information, refer to the *Infoblox CLI Guide*. Note that when you enter any BGP command and wait for the interface to return more information, the appliance disconnects your CLI session if you restart services or make other BGP configuration changes through Grid Manager. You must re-enter your credentials to log back in to the CLI.

You can configure BGP on any interface to advertise anycast addresses across multiple ASs.

Note: NIOS selects the interface for BGP advertisement based on the routing configuration.

The appliance supports BGP version 4. For more information about BGP, refer to *RFC4271, A Border Gateway Protocol 4 (BGP-4)*.

Configuring BGP in the NIOS Appliance

You can configure the appliance as a BGP advertising interface for anycast addresses. The NIOS appliance advertises the BGP routing information to the network so routers can determine the nearest server to query. The NIOS appliance does not perform dynamic routing itself; it can use dynamic routing protocols to advertise its DNS anycast availability. You must define the ASN of the interface and list any neighboring routers that will receive the BGP announcements. On an HA pair, BGP runs only on the active node. In an HA failover, the BGP service resumes on the new active node.

Note: If you encounter Malformed AS_PATH error, then remove the dont-capability-negotiate option. Infoblox doesn't provide an option to create confederation of autonomous systems if the BGP peer is configured by enabling the dont-capability-negotiate option.

To configure BGP for anycast addresses:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the *Grid Member Properties* editor, select the **Anycast** tab.
3. In the BGP Configuration section, complete the following:
 - **Interface Link Detection:** Select this check box to enable link detection when the default connection fails. This enables the router to track the next available route. For example, if LAN1 is set as the default gateway when both LAN1 and LAN2 are working, and LAN1 later fails, the router will switch to LAN2. When LAN1 reconnects, the router will then switch back to LAN1.
 - **ASN:** Enter the autonomous system number of the interface. Enter a number from 1 to 65535. You can configure only one ASN on each Grid member.
 - **BGP Timers:** BGP uses timers to control how often the interface sends KEEPALIVE messages and how long it waits before declaring a neighboring router out of service. The keepalive timer determines the time interval at which the interface sends KEEPALIVE messages to a neighboring router to inform the neighbor that the appliance is alive. The hold down timer determines how long the interface waits to hear a KEEPALIVE or UPDATE message before it assumes its neighbor is out of service. If a neighboring router is down, the interface terminates the BGP session and withdraws all the BGP routing information to the neighbor.
 - **Keep Alive:** Enter the time interval in seconds when the interface sends keepalive messages. You can enter a time from 1 to 21845 seconds. The default is four seconds.
 - **Hold Down:** Enter the time in seconds that the interface waits to hear a keepalive message from its neighbor before declaring the neighbor out of service. You can enter a time from 3 to 65535 seconds. The default is 16 seconds.

Click the Add icon to add a neighboring router to receive BGP advertisements from the NIOS appliance. The appliance adds a new row to the table. Complete the following:

- **Neighbor Router IP:** Enter the IP address (IPv4 or IPv6) of the neighboring BGP router. The neighboring router can be within the same AS (the most likely case) or from a router in an external AS.
- **Remote ASN:** Enter the ASN of the neighboring router. You can enter an ASN number from 1 to 65535.
- **Comment:** Enter useful information about this neighboring router.

Click the Add icon again to add another neighboring router. You can add up to 10 neighboring routers.

4. Save the configuration and click **Restart** if it appears at the top of the screen.
5. Anycast configuration is complete. To activate anycast, see [Specifying Port Settings for DNS](#) on page 562 and its subtopic, [Specifying Source Ports](#).



PART 5 DHCP

This section describes how to configure the Grid to provide DHCP services. It includes the following chapters:

- [Chapter 23, *Infoblox DHCP Services*](#), on page 773
- [Chapter 24, *Configuring DHCP Properties*](#), on page 791
- [Chapter 25, *Managing DHCP Templates*](#), on page 825
- [Chapter 26, *Managing IPv4 DHCP Data*](#), on page 841
- [Chapter 27, *Managing IPv6 DHCP Data*](#), on page 869
- [Chapter 28, *DHCP Failover*](#), on page 883
- [Chapter 29, *Configuring IPv4 DHCP Filters*](#), on page 891
- [Chapter 30, *Authenticated DHCP*](#), on page 915
- [Chapter 31, *Managing Leases*](#), on page 945



Chapter 23 Infoblox DHCP Services

This chapter provides an overview of the Infoblox DHCP services for IPv4 and IPv6. It contains the following sections:

- [About Infoblox DHCP Services](#) on page 774
- [IPv4 DHCP Protocol Overview](#) on page 774
- [IPv6 DHCP Protocol Overview](#) on page 775
 - [IPv6 Address Structure](#) on page 776
- [Configuring DHCP Overview](#) on page 777
- [Managing DHCP Data](#) on page 779
 - [About Networks](#) on page 779
 - [About Shared Networks](#) on page 779
 - [About DHCP Ranges](#) on page 780
 - [About Fixed Addresses](#) on page 780
 - [About Hosts](#) on page 780
- [About DHCP Inheritance](#) on page 782
 - [Overriding DHCP Properties](#) on page 783
 - [Viewing Inherited Values](#) on page 783
- [About Network Views](#) on page 787
 - [Adding Network Views](#) on page 789
 - [Modifying Network Views](#) on page 789
 - [Deleting Network Views](#) on page 790

ABOUT INFOBLOX DHCP SERVICES

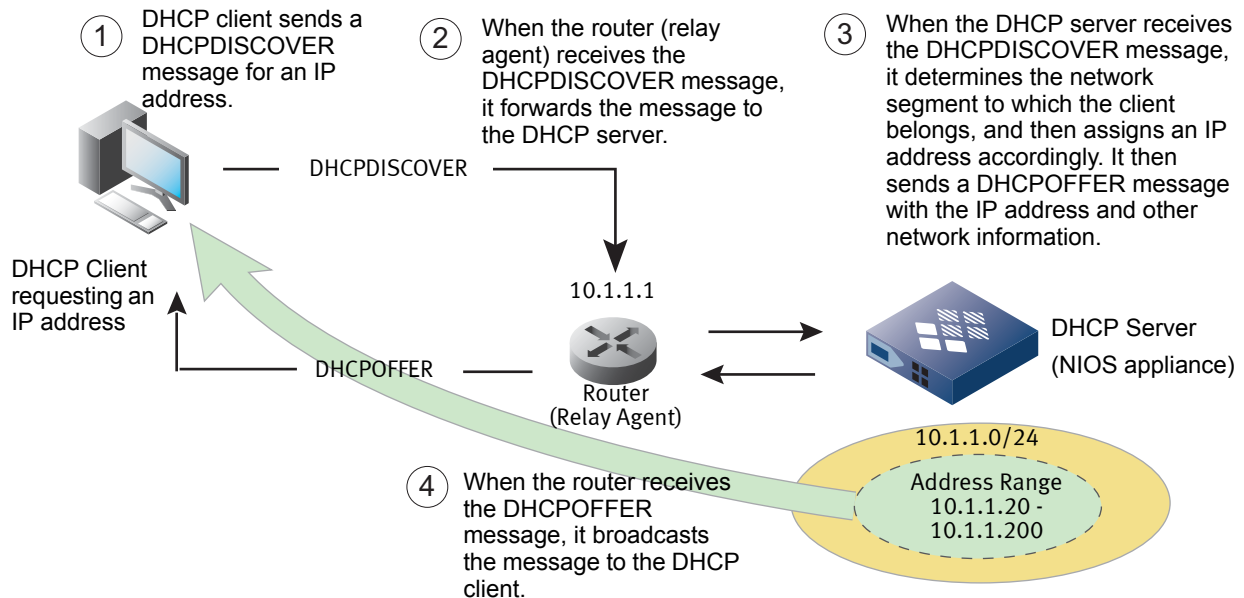
DHCP (Dynamic Host Configuration Protocol) is a network application protocol that automates the assignment of IP addresses and network parameters to DHCP-configured network devices (DHCP clients). When a DHCP client connects to a network, it sends a request to obtain an IP address and configuration information from the DHCP server. The DHCP server manages a pool of IP addresses and configuration information such as default gateway, domain name, and DNS server. Depending on the configuration, the DHCP server either assigns or denies an IP address to a client request. It also sends network configuration parameters to the client.

You can configure a NIOS appliance to provide DHCP service for IPv4 and IPv6. The Infoblox DHCP server complies with a number of DHCP and DHCPv6 RFCs (see Appendix A Product Compliance). Limited-access admin groups can access certain DHCP resources only if their administrative permissions are defined. For information on setting permissions for admin groups, see *Chapter 3, Managing Administrators*.

IPv4 DHCP PROTOCOL OVERVIEW

As illustrated in [Figure 23.1](#), when a DHCP client requests an IP address, it sends a DHCPDISCOVER message to the router, which can act as a relay agent. The router forwards the message to the DHCP server. When the DHCP server receives the DHCPDISCOVER message, it determines the network segment to which the client belongs and assigns an IP address. The DHCP server then sends a DHCPOFFER message that includes the IP address and other network configuration information. When the router receives the DHCPOFFER message, it broadcasts the message to the client that sent the DHCPDISCOVER message.

Figure 23.1 IP Address Allocation Process



IPv6 DHCP PROTOCOL OVERVIEW

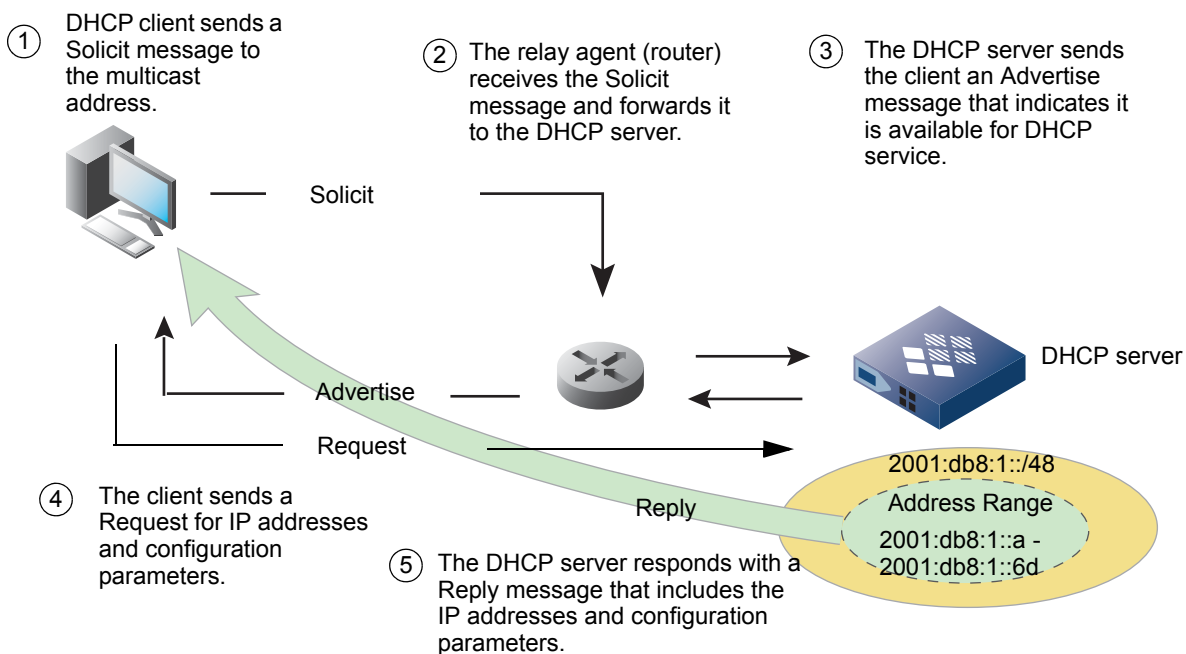
You can configure NIOS appliances to support DHCP for IPv6 (DHCPv6), the protocol for providing DHCP services for IPv6 networks.

The DHCPv6 client-server model is similar to that of IPv4. DHCP clients and servers use a reserved, link-scoped multicast address to exchange DHCP messages. When a DHCP client needs to send messages to a DHCP server that is not attached to the same link, a DHCP relay agent can be used to relay messages between the client and server.

Each IPv6 DHCP server and client has a unique DHCP unique identifier (DUID). DHCP servers use DUIDs to identify clients when providing configuration parameters, and clients use DUIDs to identify the source of the DHCP messages from servers.

As illustrated in [Figure 23.2](#), a DHCP client that needs an IPv6 address sends a Solicit message to the well-known multicast address. DHCPv6 servers then send Advertise messages to the client to indicate that they are available. The client sends a Request message to a specific DHCPv6 server to request IP addresses and configuration parameters. The DHCPv6 server responds with a Reply message that contains the IP addresses and configuration parameters. You can view statistics about the IPv6 messages on the Dashboard.

Figure 23.2 Client DHCPv6 Configuration Workflow



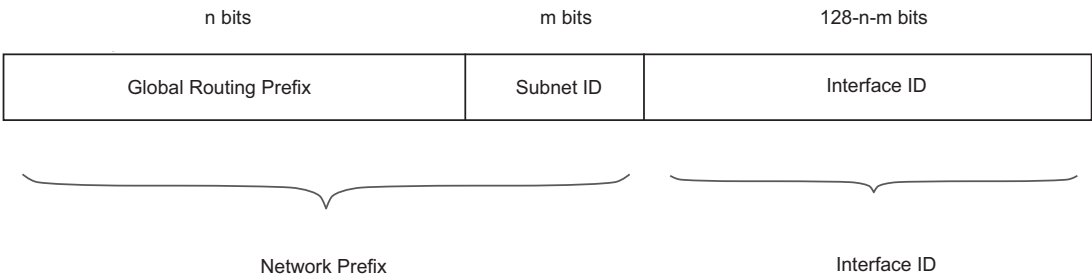
Infoblox DHCP servers also supports stateless configuration in which a DHCP client does not need IP addresses, but needs configuration information only. The DHCP client sends an Information-Request packet to obtain configuration information, and the server sends a Reply message with the requested information. For more information, refer to *RFC 2462, IPv6 Stateless Address Autoconfiguration*.

IPv6 Address Structure

An IPv6 address consists of the following:

- Global Routing Prefix—Global routing prefix is a (typically hierarchically-structured) value assigned to a site. For example, an ISP can delegate a prefix to your site, which you can then divide into subnets.
- Subnet ID—Subnet ID is an identifier of a link within the site.
- Interface ID—Interface Identifier. This portion of the address identifies the interface on the subnet. This is equivalent to the host identifier for IPv4 addresses.

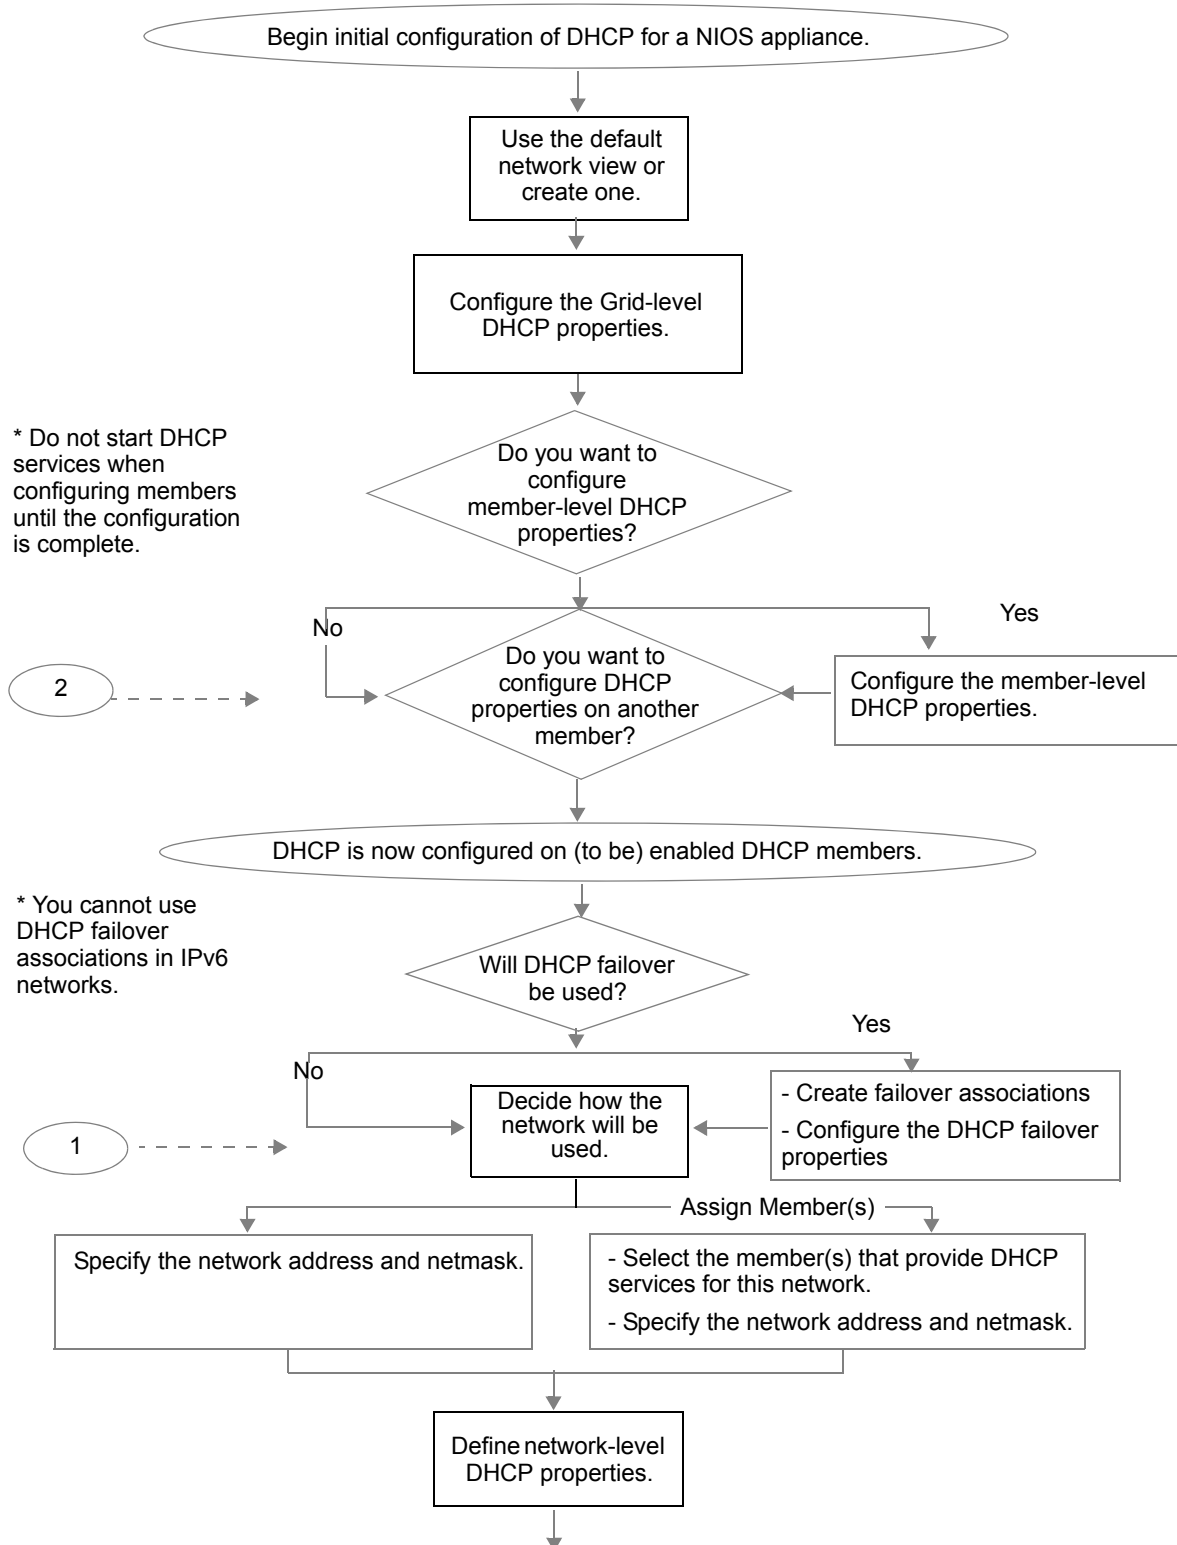
Figure 23.3 IPv6 Address Structure

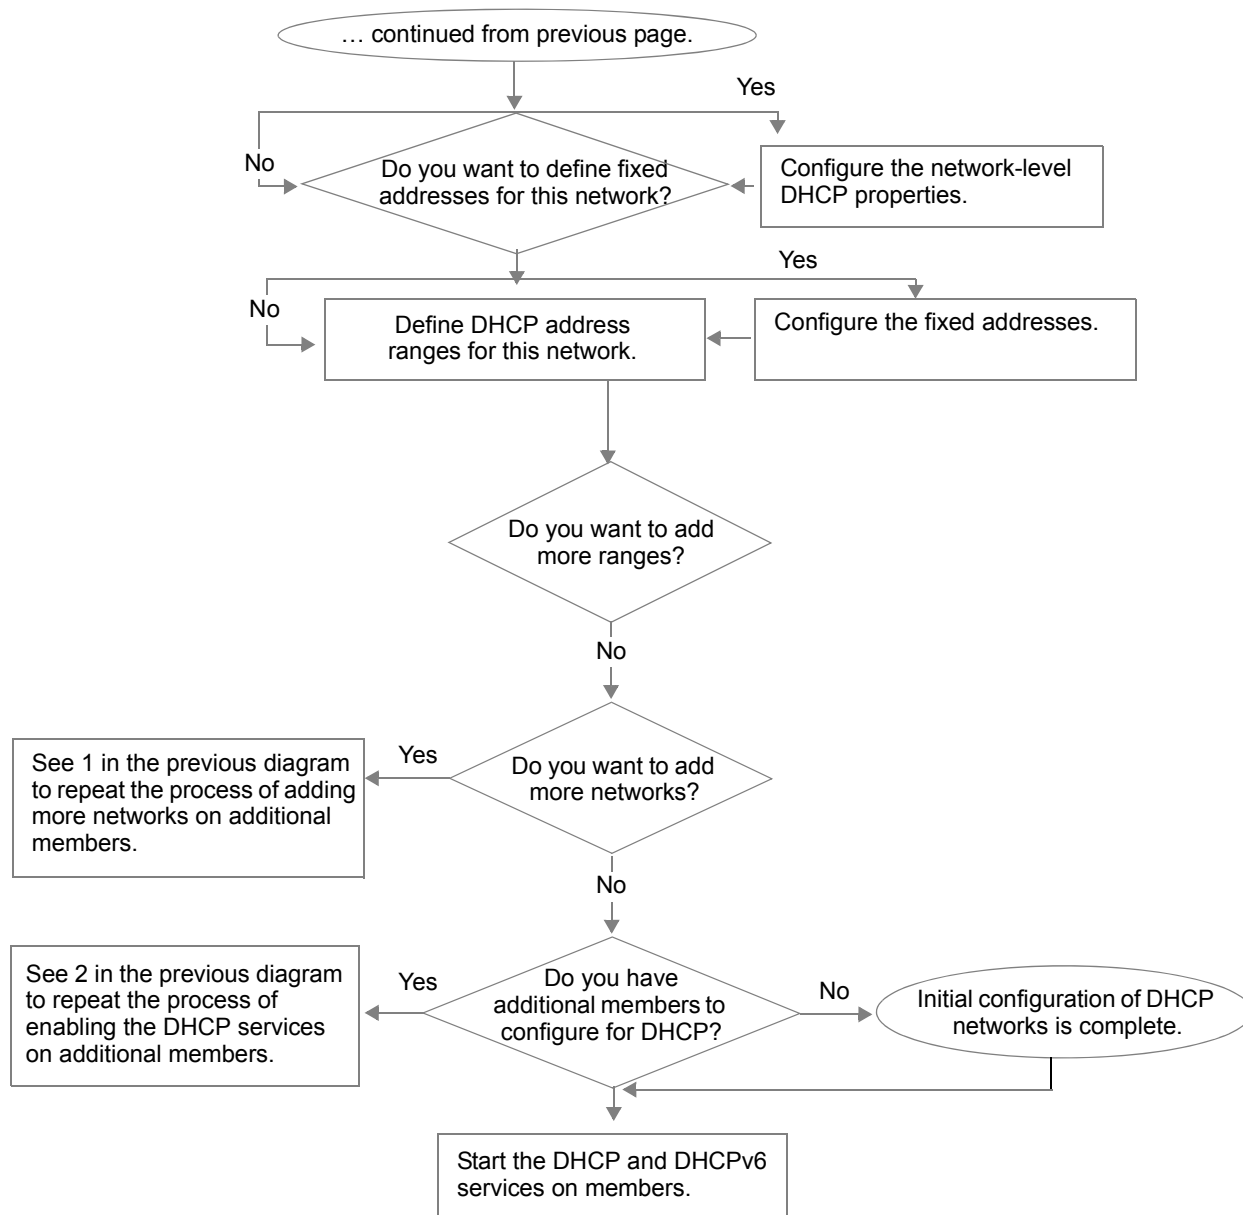


When you enter an IPv6 address in Grid Manager, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered.

CONFIGURING DHCP OVERVIEW

An overview of the complete DHCP configuration process is outlined in the following diagram (and continued on the next page), illustrating the main steps for preparing a NIOS appliance for use. Note that the process for configuring the DHCP server is the same for IPv4 and IPv6 networks, except that failover associations are not supported in IPv6 networks.





Configuring the Connecting Switch

To ensure that VRRP (Virtual Router Redundancy Protocol) works properly, configure the following settings at the port level for all the connecting switch ports (HA, LAN1, and LAN2):

- Spanning Tree Protocol: Disable. For vendor specific information, search for “HA” in the Infoblox Knowledge Base system at <https://support.infoblox.com>.
- Trunking: Disable
- EtherChannel: Disable
- IGMP Snooping: Disable
- DHCP Snooping: Disable or Enable Trust Interface

Note: You must disable DHCP Snooping to successfully run DHCP services on the Grid. For more information about DHCP services, see [About Infoblox DHCP Services](#) on page 774.

- Port Channeling: Disable
- Speed and Duplex settings: Match these settings on both the Infoblox appliance and switch
- Disable other dynamic and proprietary protocols that might interrupt the forwarding of packets

Note: By default, a NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between its LAN1 or LAN1 (VLAN), HA, and MGMT ports and the Ethernet ports on the connecting switch. If the two appliances fail to auto-negotiate the optimal settings, see [Modifying Ethernet Port Settings](#) on page 354 for steps you can take to resolve the problem.

MANAGING DHCP DATA

You can configure a NIOS appliance to provide DHCP service for IPv4 and IPv6, and manage both IPv4 and IPv6 objects. When you define DHCP objects, you can track specific information about a network device by defining extensible attributes. Extensible attributes are fields that you define to track properties such as network locations or device models. For more information, see [About Extensible Attributes](#) on page 322.

About Networks

You can configure DHCP IPv4 and IPv6 properties for the Grid and its members, and then define the IPv4 and IPv6 networks that they serve.

All networks, both IPv4 and IPv6, must belong to a network view. The appliance has one default network view and unless you create additional network views, all networks belong to the default view. Note that because network views are mutually exclusive, you can create networks with overlapping IP address spaces in two different network views. For more information, see [About Network Views](#) on page 787.

Note: The 255.255.255.255 limited broadcast address is reserved. The appliance does not automatically create glue A records for this address. You can however create an NS record without the associated glue records. For more information, see [Changing the Interface IP Address](#) on page 609.

About Shared Networks

A shared network is a network segment to which you assign two or more subnets. When subnets in a shared network contain IP addresses that are available for dynamic allocation, the addresses are put into a common pool for allocation when client requests arise. When you create a shared network, the DHCP server can assign IP addresses to client requests from any subnet (that resides on the same network interface) in the shared network. For example,

when you have networks A, B, and C on the same network interface and you assign them to a shared network, the DHCP server can allocate available IP addresses from any DHCP range within networks A, B, and C even when all the client requests originate from network A. When adding subnets to a shared network, ensure that the subnets are assigned to the same members to avoid DHCP inconsistencies.

Before creating a shared network, you must first create the subnets. For example, you must first create the IPv4 networks 10.32.1.0 and 10.30.0.0 before designating them to a shared network or create the IPv6 networks 2001:db8:1::/48 and 2001:db8:2::/48 before designating them to a shared network.

After you create a network, you can define their DHCP resources such as DHCP ranges, fixed addresses, reservations, host records, and roaming hosts. IPv4 and IPv6 support most of the same DHCP objects, except that IPv6 does not support reservations.

About DHCP Ranges

A DHCP range is a pool of IP addresses from which the appliance allocates IP addresses. You must add a DHCP address range in your network so the appliance can assign IP addresses to DHCP clients within the specified range. IPv6 DHCP ranges can also contain a range of IPv6 prefixes that it delegates to DHCP clients that request them.

You must assign a DHCP range to a Grid member. Note that you can only assign DHCP ranges to members and networks that are in the same network view. If the server is an independent appliance, you must specify this appliance as the member that serves the DHCP range. In addition, you can also assign IPv4 DHCP ranges to failover associations.

About Exclusion Ranges

You can define an exclusion range within a DHCP range. Creating an exclusion range prevents the appliance from assigning the addresses in the exclusion range to clients. IP addresses in an exclusion range are excluded from the pool of IP addresses. You can use exclusions to split a DHCP range into multiple blocks of ranges. You can also use addresses in the exclusion ranges as static IP addresses for network devices such as legacy printers that do not support DHCP. An exclusion in a range can help prevent address conflicts between statically configured devices and dynamically configured devices.

About Fixed Addresses

You can configure fixed addresses for network devices, such as routers and printers, that are not frequently moved from network to network. By creating fixed addresses for network devices, clients can reliably reach them by their domain names. Some network devices, such as web or FTP servers, can benefit from having fixed addresses for this reason. In IPv4 networks, you can also reserve an IP address that is not part of a DHCP range by defining a reservation. For information about creating reservations, see [Configuring IPv4 Reservations](#) on page 860.

About Hosts

Infoblox hosts are data objects that contain DNS, DHCP, and IPAM data of the assigned addresses. You can assign multiple IPv4 and IPv6 addresses to a host. When you create a host, you are specifying the name-to-address and address-to-name mappings for the IP addresses that you assign to the host. For information about Infoblox hosts, see [About Host Records](#) on page 459.

DHCP Configuration Checklists

After you complete the appliance configuration for each member in the Grid, as described in [Chapter 7, Managing Appliance Operations](#), on page 303, you can configure DHCP services.

The following checklist includes the major steps for configuring DHCP service for IPv4:

Table 23.1 IPv4 DHCP Configuration Checklist

Step	For more information
Configure DHCP properties for the Grid and members.	<ul style="list-style-type: none"> • Configuring IPv4 DHCP Properties on page 793 • Chapter 20, Configuring DDNS Updates from DHCP, on page 689 • Configuring DHCP IPv4 and IPv6 Common Properties on page 812 • Configuring the Lease Logging Member on page 815
Decide if you want to configure a DHCP failover association.	<ul style="list-style-type: none"> • Configuring Failover Associations on page 885
Configure networks based on your network requirements and decide if you want to override the Grid or member DHCP configuration for the networks.	<ul style="list-style-type: none"> • Configuring IPv4 Networks on page 845 • Configuring IPv4 Shared Networks on page 852
Decide if you want to configure fixed addresses and reservations, and whether to override the upper level DHCP properties for the fixed addresses and reservations.	<ul style="list-style-type: none"> • Configuring IPv4 Fixed Addresses on page 857 • Configuring IPv4 Reservations on page 860
Define DHCP ranges and decide whether to override the upper level DHCP properties for the ranges.	<ul style="list-style-type: none"> • Configuring IPv4 Address Ranges on page 854
Enable DHCP services on the member.	<ul style="list-style-type: none"> • Starting DHCP Services on a Member on page 822

The following checklist includes the major steps for configuring DHCP service for IPv6:

Table 23.2 IPv6 DHCP Configuration Checklist

Step	For more information
Configure DHCP properties for the Grid and members.	<ul style="list-style-type: none"> • Configuring DHCPv6 Properties on page 809 • Chapter 20, Configuring DDNS Updates from DHCP, on page 689 • Configuring DHCP IPv4 and IPv6 Common Properties on page 812 • Configuring the Lease Logging Member on page 815
Configure networks based on your network requirements and decide if you want to override the Grid or member DHCP configuration for the networks.	<ul style="list-style-type: none"> • Configuring IPv6 Networks on page 870 • About IPv6 Shared Networks on page 875
Decide if you want to configure fixed addresses and reservations, and whether to override the upper level DHCP properties for the fixed addresses and reservations.	<ul style="list-style-type: none"> • Configuring IPv6 Fixed Addresses on page 878
Define DHCP ranges and decide whether to override the upper level DHCP properties for the ranges.	<ul style="list-style-type: none"> • Configuring IPv6 Address Ranges on page 876

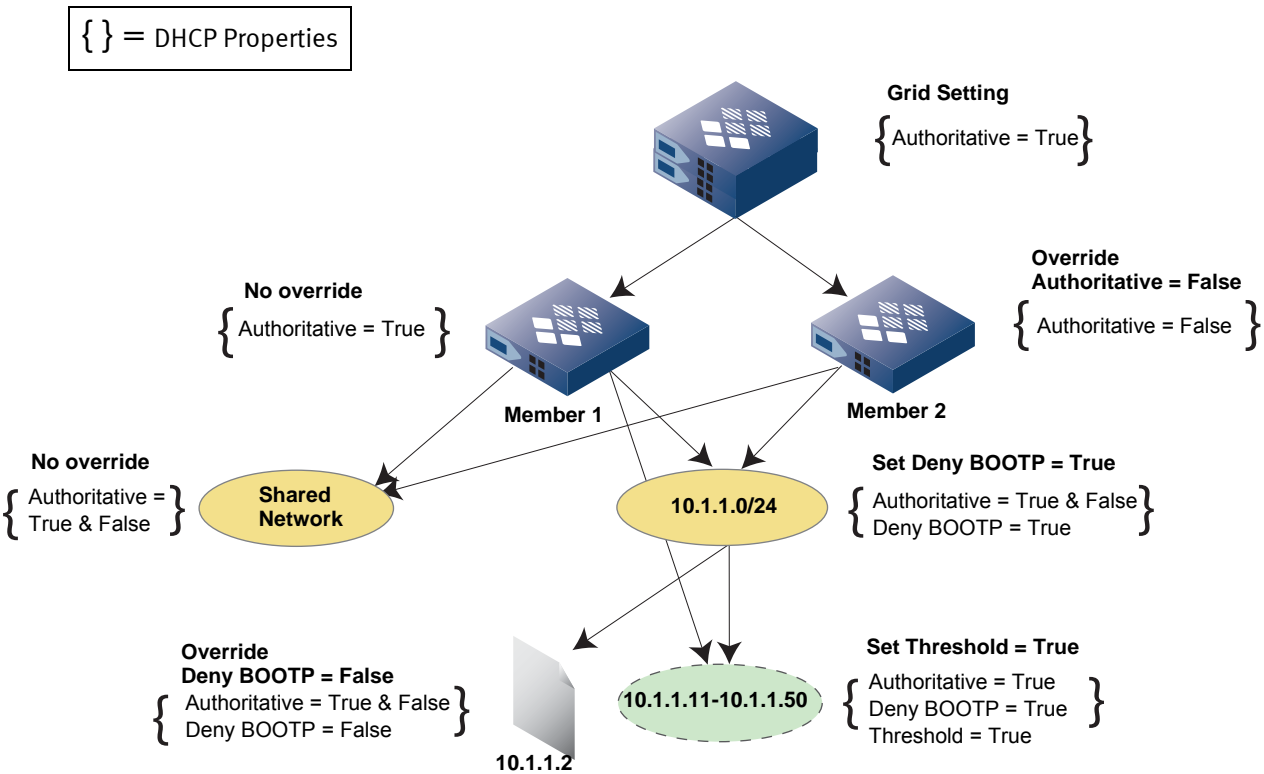
Step	For more information
Enable DHCP services on the member.	<ul style="list-style-type: none"> • Starting DHCP Services on a Member on page 822

ABOUT DHCP INHERITANCE

When you configure DHCP properties for the Grid, members, networks, shared networks, DHCP ranges, fixed addresses, reservations, host addresses, and roaming hosts, the appliance applies the configured properties hierarchically. In addition, IPv4 DHCP objects inherit IPv4 specific properties and IPv6 objects inherit IPv6 specific properties. For example, when you set DHCP IPv4 properties for the Grid, all DHCP IPv4 objects inherit the properties from the Grid unless you override them at a specific level, and the same applies for IPv6 properties and objects. Properties set at the member level override Grid-level settings and apply to the objects that the member serves. Properties set at the network level override member-level settings and apply to the objects within the network. Properties set for a DHCP range override those set at higher levels. You can also set specific properties that apply only to fixed addresses, reservations, host addresses, and roaming hosts.

[Figure](#) illustrates some inheritance scenarios that can occur in a Grid. As shown in the figure, the authoritative server configuration set for the Grid is inherited by the members. Since Member 1 has no overrides and Member 2 overrides the authoritative server configuration, they have different DHCP configurations. Grid Manager applies DHCP properties hierarchically from the Grid down. Therefore, a DHCP object below the member level can inherit DHCP properties with multiple values from multiple sources. In [Figure](#), network 10.1.1.0/24 inherits multiple values (True and False) from the members for the authoritative server configuration. The shared network, which includes 10.1.1.0/24, inherits DHCP properties from both members. For DHCP range 10.1.1.11 - 10.1.1.50, since Member 1 is the assigned member, it inherits properties from Member 1 and the network. The fixed address 10.1.1.2 overrides the BOOTP settings and inherits the authoritative server configuration from both members and the network.

Table 23.3 Inheritance Hierarchy in a Grid



When a DHCP property contains inherited values from different sources, the appliance displays the corresponding information when you create or modify an object. Based on the information provided, you can then decide whether to override or keep the inherited values. You must have read/write permissions to the DHCP resources to override inherited values. You can only view inherited values and paths if you have read-only permissions.

Overriding DHCP Properties

DHCP properties configured at the Grid level apply to the entire Grid. You can choose to keep the inherited properties or override them when you configure the properties for a member, network, shared network, DHCP range, fixed address, host address, or roaming host. For example, you can override the values of DHCP properties inherited from a member and enter unique values for a network that is configured for DHCP.

To override an inherited value:

- 1. In a wizard or editor, click **Override** next to a property to enable the configuration. The **Override** button changes to **Inherit**.
- 2. Enter a new value to override the inherited value.

Viewing Inherited Values

When you configure DHCP properties that contain inherited values, the appliance displays the information based on the inheritance sources. The following table summaries what the appliance can display:

When you see...	it means...	For details, see...
Inherited From <i><object></i>	the DHCP property has a definite value from an inheritance source.	Simple Inheritance .

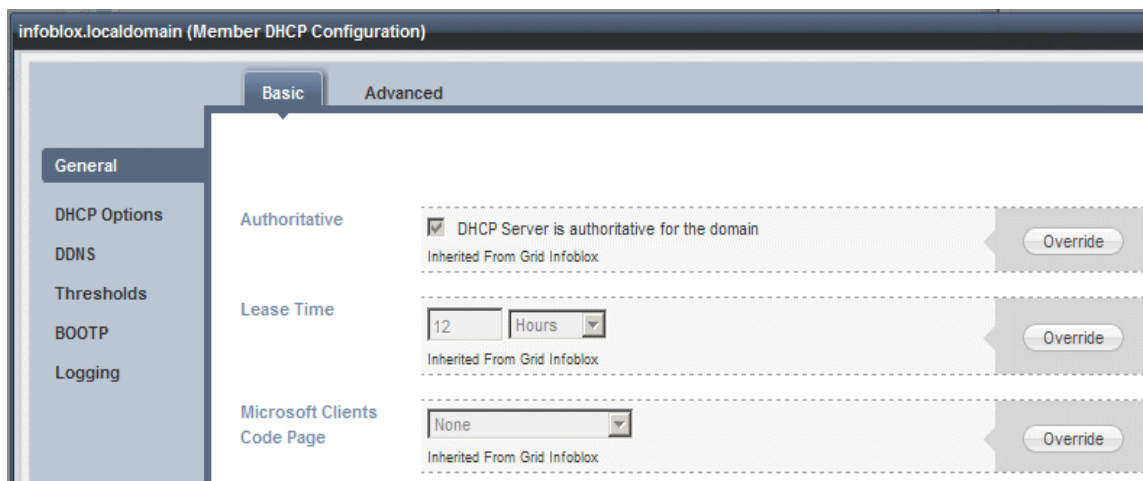
When you see...	it means...	For details, see...
Inherited From Upper Level	the appliance cannot determine the inherited value or inheritance source for the DHCP property.	Unknown Inheritance on page 784.
Inherited From Multiple	the DHCP property has the same value that it inherits from multiple sources.	Multiple Inheritance on page 785.
Settings Inherited from Multiple Ancestors, View Multiple Inheritance Scenarios	the DHCP property has multiple values that it inherits from multiple sources, and you can view the values and their corresponding sources by clicking the View Multiple Inheritance Scenarios link.	Multiple Inheritance on page 785.

Simple Inheritance

When a DHCP property has an inherited value from a specific source, the appliance displays the value. It also displays **Inherited From <object>** (where <object> can be the Grid, member, network, shared network, or DHCP range) to indicate the source from which the value is inherited.

For example, when you set DHCP properties at the Grid level and do not override the properties at any level, the members, networks, shared networks, DHCP ranges, fixed addresses, reservations, host addresses, and roaming hosts inherit these properties from the Grid. The appliance displays the property value and **Inherited From Grid Infoblox** for each configured DHCP property, as shown in [Figure 23.4](#).

Figure 23.4 Simple Inheritance



Unknown Inheritance

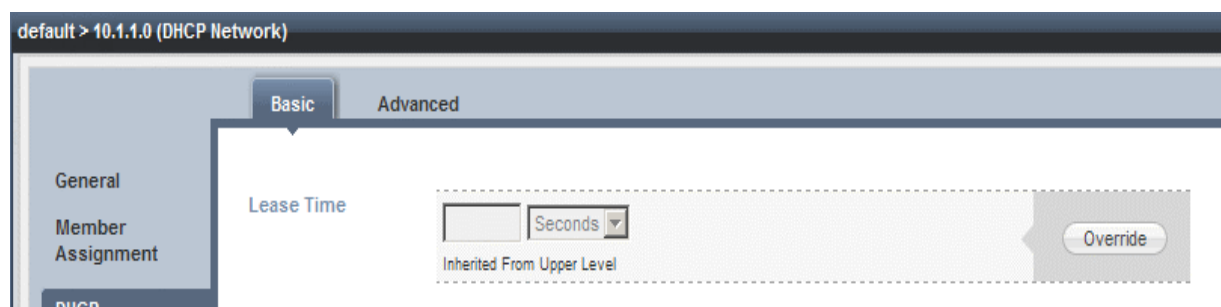
In some cases, DHCP properties may not have definite inherited values and inheritance sources. The following are examples of unknown inheritance:

- The appliance cannot determine the inheritance sources of the DHCP properties in a template until you use the template to create an object.
- When a network or a DHCP range does not have an assigned member, it does not have a clear definition of an inheritance source because a network or a DHCP range inherits properties from a member.
- When individual networks in a shared network do not have member assignments, the shared network has unknown inheritance because the shared network inherits DHCP properties from a member and its networks.

- All roaming hosts have unknown inheritance because the DHCP properties can be inherited from different DHCP ranges within a network view.

In cases where the source of the inheritance is unknown, the appliance displays **Inherited From Upper Level** as the inheritance source. As shown in [Figure 23.5](#), network 10.1.1.0 has unknown lease time value because it does not have any assigned member.

Figure 23.5 Unknown Inheritance

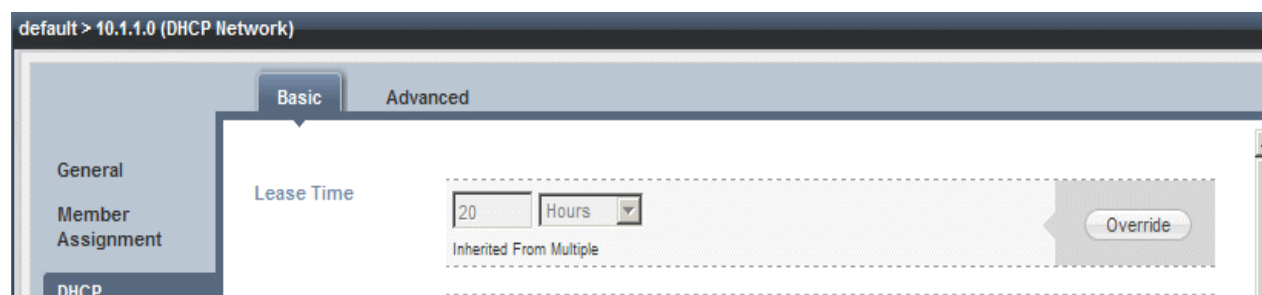


Multiple Inheritance

As illustrated in [Figure](#), a network can have multiple inherited values and inheritance sources for DHCP properties when it is served by multiple members. When an object inherits a DHCP property from different sources, the property value can be the same from all sources or it can be different. When the value is the same, the appliance displays the value in the property field. When there are multiple values inherited from multiple paths, the appliance displays the information to indicate so.

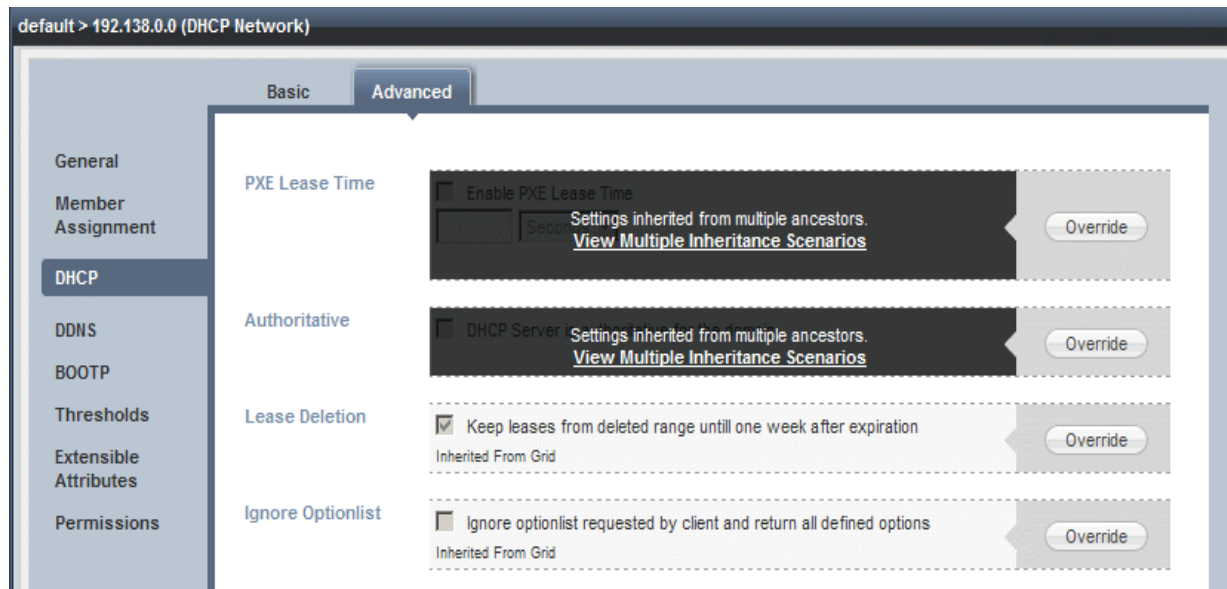
In a Grid, when two members serve the same network, the network inherits DHCP properties from both associated members. If both members have the same configured DHCP property, the network inherits the same value from both members. For example, when DHCP network 10.1.1.0 has two associated members and both members have the lease time set for 20 hours, the appliance displays the lease time value and **Inherited From Multiple** to indicate the value is inherited from multiple sources, as shown in [Figure 23.6](#).

Figure 23.6 Multiple Inherited Paths with the Same Inherited Value



In the same Grid with the two members serving the same network, the network inherits different values for the same properties if you override the Grid configuration on one member but not on the other. For example, you can configure different PXE lease times for the members and configure a member as an authoritative DHCP server for the domain and the other not. In this case, the appliance displays **Settings inherited from multiple ancestors** and provides a **View Multiple Inheritance Scenarios** link so you can view the inherited values and paths, as shown in [Figure 23.7](#).

Figure 23.7 Multiple Inheritance Sources with Multiple Values



For example, to view the multiple inherited values of the **Authoritative** field, click **View Multiple Inheritance Scenarios**, and the *Multiple Inheritance Viewer* displays the inherited values from the two members. Since member1.foo.net does not have a configured value for this field, the viewer displays **Not Set**, as shown in [Figure 23.9](#). You can use this information to determine whether you want to keep the inherited values or configure new ones.

Figure 23.8 Multiple Inheritance Viewer

Multiple Inheritance Viewer		
	infoblox.localdomain	member1.foo.net
Authoritative	false	Not Set

Another scenario of multiple inherited levels is when you have multiple DHCP properties that can inherit the same or multiple values from different sources. For example, when you configure multiple DHCP custom options, each of the options can inherit the same or multiple values from multiple paths. You can override the inherited options and configure new ones at a specific level other than the Grid level. Though these options are grouped under *DHCP Custom Options*, the appliance treats each of them as a separate property. The appliance groups the inherited options at the top, as shown in [Figure 23.9](#). You can override these options but you cannot delete them. For multiple values inherited from multiple sources, you can view the values in the *Multiple Inheritance Viewer* by clicking **View Inheritance**, as shown in [Figure 23.10](#).

Figure 23.9 DHCP Custom Options with Multiple Inheritance Sources

Name	Value	
subnet-mask (1) Inherited From Infoblox	255.255.255.0	Override
host-name (12) Inherited From Infoblox	hostname	Override
root-path (17) Inherited From Multiple	Settings inherited from multiple sources View Inheritance	Override

Figure 23.10 Multiple Inheritance Viewers for Options

	infoblox.localdomain	infoblox2.localdomain
root-path (17)	www.google.com	www.infoblox.com

When you configure email notification for the Grid or Grid member from the **Data Management** tab -> **Grid** tab, the email address you enter there is inherited by the DHCP configuration for the Grid, members, networks, and DHCP ranges unless you override it at a specific level. The appliance uses this email address to send notification for a DHCP range when the DHCP usage crosses either the effective watermark threshold. For information, see [Configuring Thresholds for DHCP Ranges](#) on page 807.

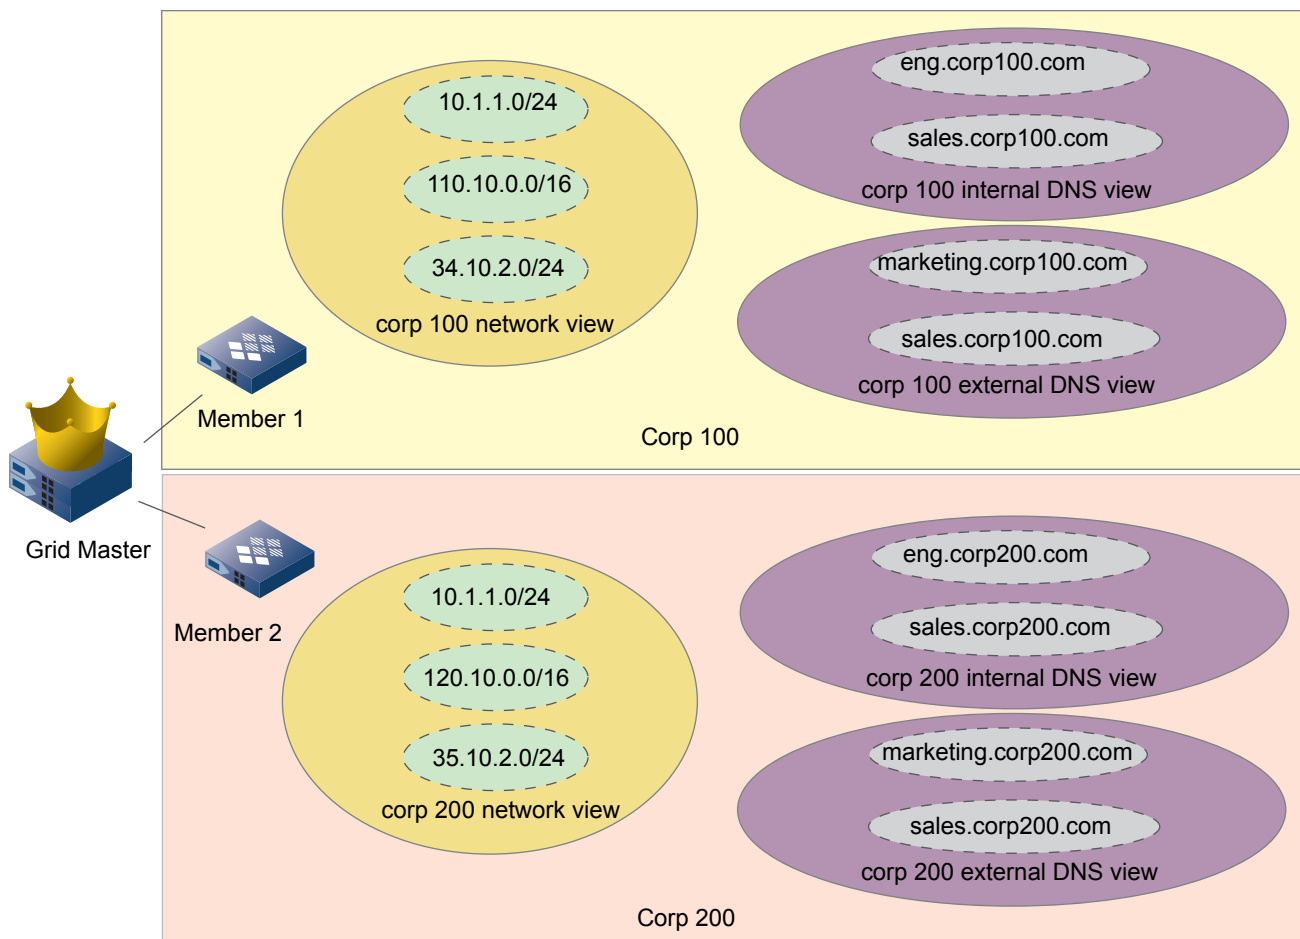
ABOUT NETWORK VIEWS

A network view is a single routing domain with its own networks and shared networks. A network view can contain both IPv4 and IPv6 networks. All networks must belong to a network view.

You can manage the networks in one network view independently of the other network views. Changes in one network view are not reflected in other network views. Because network views are mutually exclusive, the networks in each view can have overlapping address spaces with multiple duplicate IP addresses without impacting network integrity. For example, two corporations, Corp 100 and Corp 200, merge. They each have their own networks and DNS domains. They also have their own private IP address spaces in the 10.0.0.0/24 network. Both corporations have DHCP and DNS servers, and use dynamic DNS updates. The DHCP servers of each corporation serve IP addresses for networks in their respective corporations. The DHCP clients in each corporation update DNS zones within their DNS domains. They plan to migrate the networks and hosts in Corp 200 to the Corp 100 address space and the corp100.com domain. To support both networks in the meantime and to facilitate the migration, you can configure an Infoblox Grid to centrally manage the networks and domains of both corporations. As shown in [Figure 23.11](#), you can configure network views for each corporation and manage their networks independently of the other.

Member 1 serves DNS and DHCP to Corp 100. The networks of Corp 100 are contained in the corp 100 network view, which is associated with both the internal and external DNS views of the corp100.com domain. Member 2 serves DNS and DHCP to Corp 200. The networks of Corp 200 are in the corp 200 network view, which is associated with both the internal and external DNS views of the corp200.com domain. The two corporations have one overlapping network, 10.1.1.0/24.

Figure 23.11 Two Network Views Managed by a Grid



A Grid member can serve one network view only, but a network view can be served by multiple Grid members. DHCP failover associations must be defined within a single network view, and both the primary and secondary peer must serve the same network view.

The NIOS appliance provides one default network view. You can rename the default view and change its settings, but you cannot delete it. There must always be at least one network view in the appliance. If you do not need to manage overlapping IP address spaces in your organization, you can use the system-defined network view for all your networks. You do not need to create additional network views. But if there are overlapping IP address spaces and you need more than one network view, you can create up to 100 network views.

Each network view must be associated with at least one DNS view. The default network view is always associated with the default DNS view, which also cannot be deleted. When you create a network view, the appliance automatically creates a corresponding DNS view with the same name as the network view, but with “default” prepended to the name. You can then rename that system-defined DNS view, but you cannot delete it.

A network view can be associated with multiple DNS views (as shown in [Figure 23.11](#)), but a DNS view cannot be associated with more than one network view. Each network view must be associated with a unique set of DNS views. You can initiate a network discovery in only one network view at a time. When you run a discovery task, the appliance sends updates to all DNS views associated with the network view. (For information about network discoveries, see [Chapter 13, Network Discovery](#), on page 493.)

Adding Network Views

All networks must belong to a network view. You can use the default network view on the appliance and create additional network views, as needed. If you plan to enable DDNS (dynamic DNS) updates on any of the networks, DHCP ranges and fixed addresses in the network view, you must set parameters that specify which DNS view is updated for each network view.

To create a network view:

1. From the **Administration** tab, select the **Network Views** tab, and then click the Add icon.
2. In the *Network View* wizard, do the following:
 - **Name:** Enter the name of the network view.
 - **Comment:** Enter useful information about the network view.
3. Click **Next** to enter values for required extensible attributes or add optional extensible attributes for the network view. For information, see [About Extensible Attributes](#) on page 322.
4. Click **Next**, and then save the configuration or select:

Configure DDNS Properties: Configure the DNS zones that are associated with the network view to receive DDNS updates. When you select this option, the *Configure DDNS Properties* dialog box appears. The appliance saves the network view entry before it opens the *Configure DDNS Properties* dialog box. For information, see [Configuring DDNS Updates from DHCP](#) on page 689.

Modifying Network Views

1. From the **Administration** tab, select the **Network Views** tab -> *network_view* check box, and then click the Edit icon.
2. The *Network View* editor provides the following tabs from which you can edit data:
 - **General:** You can modify all the fields in this tab.
 - **Members:** This tab displays the members that provide DHCP services for the networks in this network view. You cannot modify information in this tab. It displays the following:
 - **Name:** The name of the DHCP member.
 - **IP Address:** The IP address of the DHCP member.
 - **Failover Association:** The name of the failover association to which the DHCP member belongs. If there are multiple failover associations, only the first one is displayed.
 - **Comment:** The information that you entered for the DHCP member.

You can sort the information in the table by column. You can also print and export the information.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network view. You can also modify the values of extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 - **Permissions:** This tab displays only if you belong to a superuser admin group. For information, see [Administrative Permissions for DHCP Resources](#) on page 205.

Deleting Network Views

You can delete any network view, except for the default network view. You can delete a network view that has only one DNS view associated with it. You cannot delete a network view that has more than one DNS view associated with it. When you delete a network view, the appliance deletes all the networks and records within the network view.

To delete a network view:

1. From the **Administration** tab, select the **Network Views** tab -> *network_view* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

The appliance removes the network view and its associated DNS views. You can restore the network view from the Recycle Bin, if enabled. If you restore a network view, the appliance restores the associated DNS views as well. For information about the Recycle Bin, see [Using the Recycle Bin](#) on page 64.



Chapter 24 Configuring DHCP Properties

This chapter explains how to configure DHCP IPv4 and IPv6 properties. It contains the following sections:

- [About DHCP Properties](#) on page 793
- [Configuring IPv4 DHCP Properties](#) on page 793
- [Configuring General IPv4 DHCP Properties](#) on page 793
 - [Specifying Authoritative](#) on page 794
 - [Defining Lease Times](#) on page 794
 - [Scavenging Leases](#) on page 794
- [Configuring Ping Settings](#) on page 795
- [Configuring One Lease per Client](#) on page 797
- [Configuring IPv4 BOOTP and PXE Properties](#) on page 798
- [About IPv4 DHCP Options](#) on page 800
 - [DHCP Option Data Types](#) on page 800
 - [Configuring IPv4 DHCP Options](#) on page 801
 - [Defining Basic IPv4 Options](#) on page 801
 - [Defining IPv4 Option Spaces](#) on page 802
 - [Configuring Custom DHCP Options](#) on page 803
 - [Applying DHCP Options](#) on page 804
 - [Configuration Example: Defining a Custom Option](#) on page 805
 - [Defining Option 60 Match Rules](#) on page 805
 - [About the DHCP Relay Agent Option \(Option 82\)](#) on page 806
- [Configuring Thresholds for DHCP Ranges](#) on page 807
- [Configuring DHCPv6 Properties](#) on page 809
 - [Defining General IPv6 Properties](#) on page 809
- [About DHCPv6 Options](#) on page 810
 - [Configuring DHCPv6 Options](#) on page 810
 - [Defining IPv6 Option Spaces](#) on page 810
 - [Configuring Custom IPv6 DHCP Options](#) on page 811
 - [Applying DHCPv6 Options](#) on page 811
- [Configuring DHCP IPv4 and IPv6 Common Properties](#) on page 812
 - [Configuring UTF-8 Encoding for Hostnames](#) on page 812
 - [Associating Networks with Zones](#) on page 813

- [*Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges*](#) on page 814
 - [*Configuring Fixed Address Leases For Display*](#) on page 814
- [*Configuring DHCP Logging*](#) on page 815
 - [*Configuring the Lease Logging Member*](#) on page 815
- [*About IF-MAP*](#) on page 816
 - [*Configuring a Grid to Support IF-MAP*](#) on page 817
 - [*Validating the IF-MAP Server Certificate*](#) on page 818
 - [*Configuring Members as IF-MAP Clients*](#) on page 818
 - [*Creating IF-MAP Client Certificates*](#) on page 819
 - [*Overriding IF-MAP Publishing Settings*](#) on page 820
 - [*Deleting Data from the IF-MAP Server*](#) on page 821
- [*Starting DHCP Services on a Member*](#) on page 822
- [*Viewing DHCP Member Status*](#) on page 822
 - [*Viewing DHCP Configuration Files*](#) on page 824

ABOUT DHCP PROPERTIES

When you configure a NIOS appliance to function as a DHCP server, you can set DHCP properties that control how the appliance operates and enable DHCP service for IPv4 and IPv6.

You can also specify configuration information the appliance includes in its IPv4 and IPv6 DHCP messages. When a DHCP server assigns an IP address to a client, it can include information the client needs to connect to the network and communicate with other hosts and devices on the network. You can set these properties at the Grid level and override them for a member, network, shared network, DHCP range, fixed address, IPv4 reservation, host address, or roaming host.

When you configure a DHCP object that has inherited DHCP properties, you can either keep the inherited properties or override them. The appliance displays the inherited values and the levels from which the DHCP properties are inherited. For information, see [About DHCP Inheritance](#) on page 782.

CONFIGURING IPv4 DHCP PROPERTIES

The following sections describe how to configure properties that apply to IPv4 DHCP objects only. You can configure and define the following DHCP properties:

- General properties, as described in [Configuring General IPv4 DHCP Properties](#) on page 793.
- Ping settings, as described in [Configuring Ping Settings](#) on page 795.
- One-lease-per-client settings, as described in [Configuring One Lease per Client](#) on page 797.
- BOOTP and PXE properties, as described in [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.
- Custom DHCP options, as described in [About IPv4 DHCP Options](#) on page 800.
- DDNS settings, as described in [Chapter 20, Configuring DDNS Updates from DHCP](#), on page 689.
- Ignore DHCP client identifiers, as described in [Ignoring DHCP Client Identifiers](#) on page 797.
- Thresholds for DHCP ranges, as described in [Configuring Thresholds for DHCP Ranges](#) on page 807.

For information on configuring properties that apply to IPv4 and IPv6 DHCP objects, see [Configuring DHCP IPv4 and IPv6 Common Properties](#) on page 812.

Note: Limited-access admin groups can access certain DHCP resources only if their administrative permissions are defined. For information on setting permissions for admin groups, see [Administrative Permissions for DHCP Resources](#) on page 205.

CONFIGURING GENERAL IPv4 DHCP PROPERTIES

When you configure general IPv4 DHCP properties at the Grid level, the configuration applies to the entire Grid. Though you can set DHCP properties at the Grid level, you can enable DHCP services at the member level only. Infoblox recommends that you configure the DHCP properties before you enable DHCP on the appliance. Depending on the properties, you can override some of them for the members, networks, DHCP ranges, fixed addresses, reservations, host addresses, and roaming hosts. To override an inherited DHCP property, click **Override** next to the property to enable the configuration.

Specifying Authoritative

Only authoritative DHCP servers can send clients DHCPNAK messages when they request invalid IP addresses. For example, a client moves to a new subnet and broadcasts a DHCPREQUEST message for its old IP address. An authoritative DHCP server responds with a DHCPNAK, causing the client to move to the INIT state and to send a DHCPDISCOVER message for a new IP address. Authoritative servers also respond to DHCPINFORM messages from clients that receive their IP addresses from the DHCP server and require additional options after the initial leases have been granted.

Defining Lease Times

When you configure DHCP general properties, you can specify the length of time the DHCP server leases an IP address to a client. The default on the appliance is 12 hours, and you can change this default according to your network requirements. There are a number of factors to consider when setting the lease time for IP addresses, such as the types of resources and clients on the network, and impact to traffic and performance. With NIOS appliances, you can set lease times at different levels, based on these factors. You can set a default lease time at the Grid level and then override this setting for specific members, networks, IP address ranges or fixed addresses when appropriate.

Scavenging Leases

You can enable member DHCP servers to automatically delete free and backup leases that remain in the database beyond a specified period of time. When you enable this feature, the appliance permanently deletes the free and backup leases, and you can no longer view or retrieve the lease information.

Note that the period of time that you specify is the number of days or weeks after the expiration date of a lease, not its release date. For example, you specify a time period of 5 days when you enable this feature. If the lease time of an IP address is 10 days, but the lease is released after five days, the appliance still deletes the lease from the database after 15 days because the IP address has been leased.

You can enable this option globally at the Grid level, and more specifically for a member, shared network, network or DHCP range. You can also enable this option for a network template or DHCP range template.

Note: If you plan to enable this feature after upgrading from a previous NIOS version, Infoblox recommends that you enable it during off-peak hours, as it may impact DHCP services.

To configure general IPv4 properties:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
Network: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
DHCP Range: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.
Fixed Address: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed_address* check box, and then click the Edit icon.
Reservation: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *reservation* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor of a Grid or member, select the **General Basic** tab. For all other objects, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, select the **IPv4 DHCP Options Advanced** tab.
3. Complete the following:
 - **Authoritative:** Select **DHCP server is authoritative** to set the DHCP server as authoritative for the domain. This can be set for the Grid, member, network and range.
 - **Lease Time:** Enter the lease time and select the time unit from the drop-down list. The default is 12 hours.

To set all other properties for a Grid or member, toggle to the advanced mode and select the **General Advanced** tab to complete the following:

- **DHCP Lease Scavenging:** Select the **Scavenge free/backup leases that persist longer than** check box and specify the number of days or weeks that free and backup leases remain in the database before they are automatically deleted. This can be set for the Grid, member, network and range.
- **Ignore Optionlist:** Select **Ignore optionlist requested by client and return all defined options** if you want the appliance to ignore the requested list of options in the DHCPREQUEST messages it receives from DHCP clients, and to include all the configured options in the DHCPACK and DHCPOFFER messages it sends back to the clients.
- **LEASEQUERY:** Select **Allow LEASEQUERY** to enable the DHCP server to respond to DHCPLEASEQUERY messages.

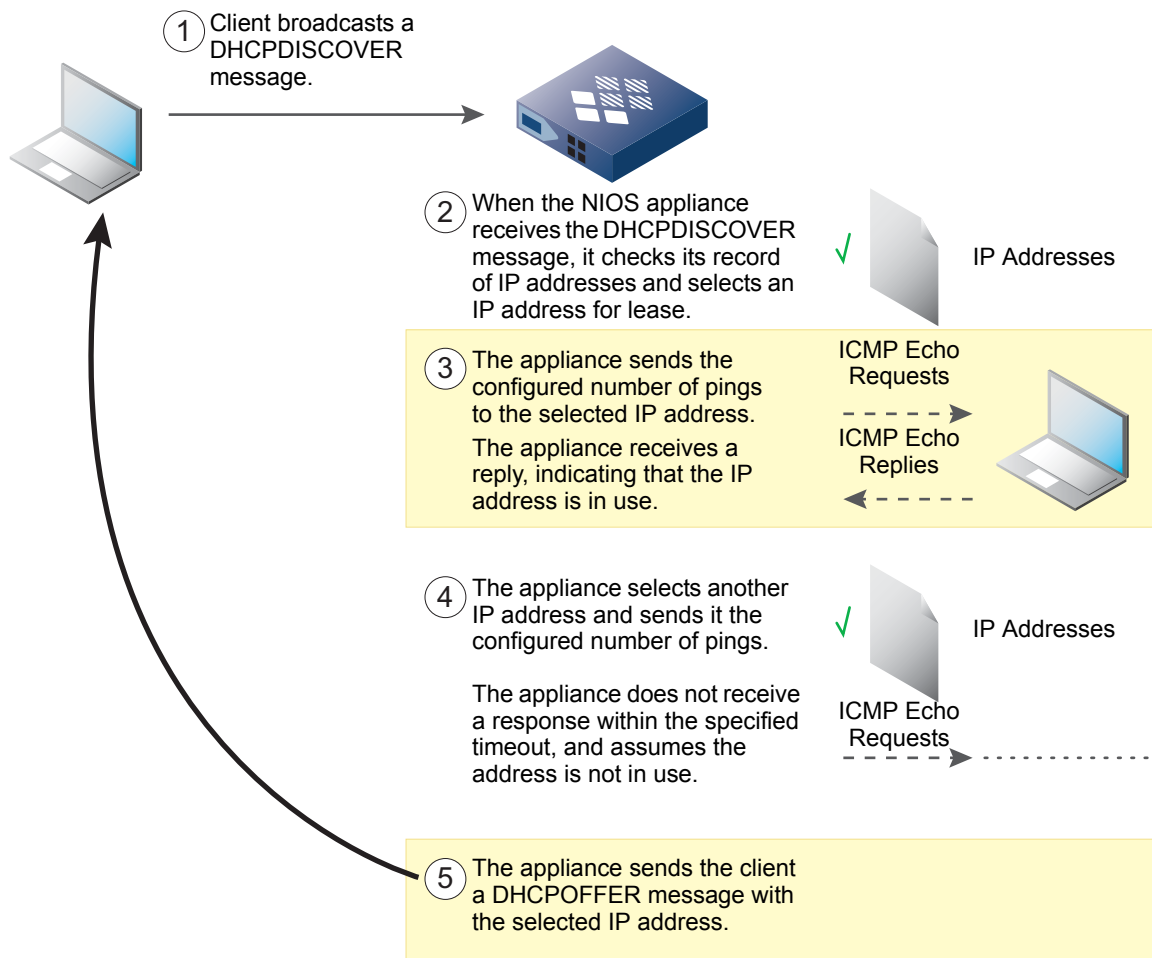
4. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING PING SETTINGS

When a DHCP client first tries to connect to a network, it broadcasts its request for an IP address. When the appliance receives such a request, it checks its record of assigned IP addresses and leases. Because there are a limited number of IP addresses available, the appliance reassigns IP addresses whose leases might have expired. Therefore, once the appliance selects a candidate IP address for lease, it sends an ICMP echo request (or ping) to the IP address to verify that it is not in use.

If the appliance receives a response, this indicates that the IP address is still in use. Note that the lease status for this IP address is **Abandoned**. The appliance then selects another candidate IP address and sends it a ping. The appliance continues this process until it finds an IP address that does not respond to the ping. The appliance then sends a DHCPOFFER message with the unused IP address to the DHCP client.

Figure 24.1 Ping Overview



By default, the appliance pings the candidate IP address once and waits one second for the response. You can change these default settings to better suit your environment. Though you can increase the ping or timeout value to accommodate delays caused by problems in the network, increasing any of these values increases the delay a client experiences when acquiring a lease. You can also disable the appliance from sending pings by changing the number of pings to 0.

You can define ping settings for an entire Grid, and when necessary, define different ping settings for a member. Settings at the member level override settings at the Grid level.

To configure ping settings:

- Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
- In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **General** tab -> **Advanced** tab and complete the following:
 - Number of Ping Requests:** Enter the number of pings the appliance sends to an IP address to verify that it is not in use. The range is 0 to 10, inclusive. Enter 0 to disable DHCP pings.
 - Ping Timeout:** Select the ping timeout value from the drop-down list.
- Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING ONE LEASE PER CLIENT

You can enable one-lease-per-client to ensure that each DHCP IPv4 client receives only one lease at any given time. When you enable one-lease-per-client and a DHCP client sends a DHCPREQUEST for a particular lease, the appliance releases other leases that the client holds, on the interface that the client is currently using. Note that this feature supports only DHCP IPv4 clients.

Enabling one-lease-per-client is useful when you want to control the number of leases on your subnets and ensure that each DHCP client receives only one lease at a time. Typically, you enable one-lease-per-client for a DHCP client that moves around a lot within different subnets and uses long leases.

When you enable one-lease-per-client at the Grid level, all members inherit the setting. You can override the Grid setting for each member. By default, one-lease-per-client is disabled.

Note: This feature enables a single lease per client on a per member basis, not on a Grid wide basis. Lease information is not replicated among members.

To configure one-lease-per-client:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the *Edit* icon.
2. In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. Click the **General** tab -> **Advanced** tab and complete the following:
 - **One Lease Per Client:** Select this check box to enable one-lease-per-client. By default, this check box is not selected.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

IGNORING DHCP CLIENT IDENTIFIERS

You can set the DHCP server to ignore the UID (unique client identifier) of a DHCP client when it places a request to the DHCP server for a new lease. When you enable “Ignore DHCP Client ID” and a DHCP client sends a DHCPREQUEST for a lease, the DHCP server identifies the DHCP client using the physical MAC address of the appliance while the UID is ignored. The DHCP server then allocates an IP address based on the MAC address of the DHCP client.

For example, when a DHCP client places a request for a new lease, the DHCP server identifies the DHCP client with the MAC address and allocates the same IP address that was previously allocated for that MAC address.

You can define this feature at the Grid level, which is inherited at the member, shared network, IPv4 network and range level. This feature is disabled by default.

Note: This feature is applicable only to dynamic leases and does not have any effect on the static lease generated for fixed addresses or roaming hosts.

To ignore the client identifier of DHCP clients:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the **Toolbar**.
Member: From the **Data Management** tab, select the **DHCP** tab -> and click the **Members** tab -> *member* check box -> *Edit* icon.
Standalone DHCP: From the **Data Management** tab, select the **DHCP** tab, and then click **System DHCP Properties**.
Shared Network Editor: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared_network* check box, and then click the *Edit* icon.
IPv4 Network Editor: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* check box, and then click the *Edit* icon.

IPv4 Range Editor: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> click on the network address. Select the *IP address range* check box, and then click the *Edit* icon.

2. In the *DHCP Properties* editor, select the **General** tab -> **Advanced** tab (or click **Toggle Advanced Mode**) and then complete the following:
 - **Ignore DHCP Client ID:** By default, this check box is deselected at the Grid level. Select this check box to ignore the client identifier of a DHCP client while placing a request to the DHCP server for a new lease. The DHCP server then will only identify the MAC address and ignores the client identifier. DHCP clients requesting leases with different client UIDs receive the same IP address based on the MAC address.
The initial default state is inherited from the Grid level. Click **Override** to modify the inherited setting. To retain the same state as the Grid, click **Inherit** at the member, IPv4 network and range, and shared network level.
3. Save the configuration and click **Restart** at the top of the screen.

Limitations of the Ignore Client ID Feature on DHCP Failover Associations

- You cannot assign a DHCP range that has the ignore DHCP client ID feature enabled to a DHCP failover association if:
 - one of the members is an external DHCP server in the failover association.
 - one of the members is running a NIOS version earlier than 6.6.
- The DHCP failover association does not work if a DHCP range having multiple inherited values has the ignore DHCP client ID feature enabled on one server and disabled on the other.
- The range assigned to a DHCP failover association and the member (failover peer) must have the same DHCP range setting. The DHCP failover association does not work if a range associated with it does not have the same ignore DHCP client ID setting as the member.

CONFIGURING IPv4 BOOTP AND PXE PROPERTIES

You can configure the DHCP server to support IPv4 clients that use BOOTP (bootstrap protocol) or that include the TFTP server name option and boot file name option in their DHCPREQUEST messages. You can specify the name or IP address of the boot server and the name of the file the host needs to boot.

In addition, you can configure the DHCP server to support hosts that use PXE (Preboot Execution Environment) to boot remotely from a server. When such a host starts up, it first requests an IP address so it can connect to a server on the network and download the file it needs to boot. After it downloads the file, the host reboots and sends another IP address request. To better manage your IP resources, set a different lease time for PXE boot requests. You can configure the DHCP server to allocate an IP address with a shorter lease time to hosts that send PXE boot requests, so IP addresses are not leased longer than necessary.

You can configure BOOTP and PXE properties at the Grid level and override them for members, IPv4 networks, DHCP ranges, fixed addresses, and reservations, host addresses, and roaming hosts. You cannot configure BOOTP and PXE properties for IPv6 DHCP objects.

To configure or override BOOTP and PXE properties:

1. **Grid Level:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.
Member Level: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
Network Level: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
DHCP Range Level: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.

Fixed Address Level: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed_address* check box, and then click the Edit icon.

Reservation: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *reservation* check box, and then click the Edit icon.

2. In the *DHCP Properties* editor, select the **BOOTP/PXE** tab and complete the following

- **PXE Lease Time:** Select **Enable PXE Lease Time** if you want the DHCP server to use a different lease time for PXE clients. You can specify the duration of time it takes a host to connect to a boot server, such as a TFTP server, and download the file it needs to boot. For example, set a longer lease time if the client downloads an OS (operating system) or configuration file, or set a shorter lease time if the client downloads only configuration changes. Enter the lease time for the preboot execution environment for hosts to boot remotely from a server.
- **Deny BOOTP Requests:** Select **Deny BOOTP Requests** to disable the BOOTP settings and deny BOOTP boot requests.
- Complete the following in the **BOOTP Settings** section:
 - **Boot File:** Enter the name of the boot file the client must download.
 - **Next Server:** Enter the IP address or hostname of the boot file server where the boot file is stored. Complete this field if the hosts in your network send requests for the IP address of the boot server. If the TFTP server is the NIOS appliance that is also serving DHCP, enter the IP address of the appliance.
 - **Boot Server:** Enter the name of the server on which the boot file is stored. Clients can request for either the boot server name or IP address. Complete this field if the hosts in your network send requests for the boot server name. If the TFTP server is the appliance that is also serving DHCP, enter the name of the appliance.

Note: Enter values in both the **Next Server** and **Boot Server** fields if some hosts on your network require the boot server name and others require the boot server IP address.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note that a few characters need manual escaping when you configure a DHCP boot file name, in order to keep the *dhcpcd.conf* file consistent. If you do not use appropriate escape characters, then it might lead to a non working boot file name. The following characters require manual escaping:

- '\t' – Tabulation character
- '\r' – Carriage return
- '\n' – New line
- '\b' – Bell
- '\xYY' – YY hex-number (a-f, 0-9)

For example, if you set the 'Boot File' to:

```
'\x86\topdir\subdir\file.img'
```

You might need to enter \x and \t as the manual escape characters:

```
'\\x86\\topdir\\subdir\\file.img'
```

You can also specify all \ as the manual escape character:

```
'\\x86\\topdir\\subdir\\file.img'
```

The above commands result in the underlying *dhcpcd.conf* file:

```
'\x5cx86\x5ctopdir\\subdir\\file.img'
```

or

```
'\x5cx86\x5ctopdir\x5csubdir\x5cfile.img'
```

ABOUT IPv4 DHCP OPTIONS

DHCP options provide specific configuration and service information to DHCP clients. These options appear as variable-length fields at the end of the DHCP messages that DHCP servers and clients exchange. For example, DHCP option 3 is used to list the available routers in the network of the client and option 6 is used to list the available DNS servers.

An option space is a collection of options. ISC (Internet Systems Consortium) DHCP has five predefined option spaces: dhcp, agent, server, nwip, and fqdn. The NIOS appliance supports only the predefined DHCP option space, which contains the industry standard options as well as additional options you can configure as needed:

- **Predefined options:** These are option codes 1 to 125. They are allocated by the IANA and defined by IETF standards. The DHCP server knows these standard options, and they are predefined on the server. You cannot redefine these options or delete them from the DHCP option space.
- **Custom options:** These are option codes 126 to 254. They are not defined by IETF standards and are available for private use. You can use these option codes to provide configuration or service information that none of the predefined options provide.

You can also create option spaces to define new groups of options. For example, you can create additional option spaces to define vendor specific options, which are encapsulated in option 43. When a DHCP client requests vendor specific options, it makes a request using the vendor identifier set in option 60 and a list of requested vendor specific options (option 43). The DHCP server then responds with the list of replies for the various options encapsulated into option 43.

Note that custom options defined in the DHCP option space are included in the options section of the DHCP messages that DHCP servers and clients exchange. Custom options defined in a user-defined option space are always encapsulated in option 43 in DHCP messages.

You can apply options globally at the Grid level, or more specifically at the member, network, range, host and roaming host levels.

You can also create an option filter the appliance uses to filter address requests by the DHCP options of requesting hosts. The filter instructs the appliance to either grant or deny an address request if the requesting host matches the filter. For information, see [Defining Option Filters](#) on page 902.

The DHCP option configuration conforms to the following RFCs:

- *RFC 2132, DHCP Options and BOOTP Vendor Extension*
- *RFC 3046, DHCP Relay Agent Information Option.* The supported options include option 60 (Client Identifier), 21 (Policy Filter), 22 (Maximum Datagram Reassembly Size), 23 (Default IP Time-to-Live), and 82 (Support for Routed Bridge Encapsulation).
- *RFC 3925, Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4)*
- *RFC 2939, Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types*

DHCP Option Data Types

Each DHCP option is identified by a name and an option code number, and specifies a data type. The data type for some options is predefined. For example, in the DHCP option space, the data type for option 1: subnet-mask is an IP address. You cannot change the data type for this option. The data type for some options is user-defined and can be in one of the formats shown in [Table 24.1](#).

Table 24.1 DHCP Option Data Types

Data type	Specifies
String	An ASCII text string (the same as the <i>text</i> data type) or a list of hexadecimal characters separated by colons Formatting to distinguish an ASCII text string from a hexadecimal string is important. For details, see the following section

Data type	Specifies
Boolean	A flag with a value of either <i>true</i> or <i>false</i> (or <i>on</i> or <i>off</i>)
IP address	A single IP address
Array of IP addresses	A series of IP addresses, separated by commas You can optionally include a space after each comma
Text	An ASCII text string
8-, 16-, or 32-bit unsigned integer	A numeric range of the following possible values 8-bit unsigned integer: from 0 to 255 16-bit unsigned integer: from 0 to 65,535 32-bit unsigned integer: from 0 to 4,294,967,295
8-, 16-, or 32-bit signed integer	A numeric range of the following possible values 8-bit signed integer: from -128 to 127 16-bit signed integer: from -32,768 to 32,767 32-bit signed integer: from -2,147,483,648 to 2,147,483,647
Domain name	A list of domain names, separated by spaces

When defining a hexadecimal string for a DHCP option (such as option 43, vendor encapsulated options), use only hexadecimal characters (0-9, a-f, or A-F) without spaces and separated by colons. The accepted form for a hexadecimal string, as presented in a regular expression, is `[0-9a-fA-F]{1,2}(:[0-9a-fA-F]{1,2})*`

Two examples of correctly written hexadecimal strings:

- aa:de:89:1b:34
- 1C:8:22:A3 (Note that the DHCP module treats a single hexadecimal character, such as “8” as “08”.)

A few examples of incorrectly written hexadecimal strings:

- :bb:45:d2:1f – Problem: The string erroneously begins with a colon.
- bb:45:d2:1f: – Problem: The string erroneously ends with a colon.
- bb:4 5:d2:1f – Problem: The string erroneously includes a space between two characters (“4” and “5”).
- bb:45:d2:1g – Problem: The string erroneously includes a nonhexadecimal character (“g”).

The DHCP module treats incorrectly written hexadecimal strings as simple text strings, not hexadecimal strings. If the string appears in quotes, it is a text string.

Configuring IPv4 DHCP Options

To use DHCP options, you can do the following:

- Define basic DHCP options, as described in [Defining Basic IPv4 Options](#).
- Configure one or more option spaces, as described in the next section [Defining IPv4 Option Spaces](#).
- Define custom options in the predefined DHCP option space or add options to an option space that you configured. For more information, see [Configuring Custom DHCP Options](#) on page 803.
- Specify values for the options and apply them to the Grid, or to a member, network, range, fixed address, reservation, host, or roaming host. For more information, see [Applying DHCP Options](#) on page 804.

Defining Basic IPv4 Options

You can define basic DHCP options that the DHCP server uses to provide configuration information to DHCP clients. The server includes these options in its DHCP messages.

To define DHCP options:

1. **Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
DHCP Range: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *DHCP_range* check box, and then click the Edit icon.
Fixed Address: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed_address* check box, and then click the Edit icon.
Reservation: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *reservation* check box, and then click the Edit icon.
Host Address: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *host_record* check box, and then click the Edit icon. Select the host IP address, and then click the Edit icon.
Roaming Host: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming_host* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **DHCP Options Basic** or **DHCP Basic** tab and complete the following:
 - **Routers:** Click the Add icon. Grid Manager adds a row to the table. In the table, enter the IP address of the router that is connected to the same network as the DHCP client. When configuring this for a template, enter the offset value of the IP address of the router. The DHCP server includes this information in its DHCP OFFER and DHCP ACK messages.
 - **Domain Name:** Enter the name of the domain for which the Grid serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCP OFFER packet to a DHCP DISCOVER packet from a client. If DDNS is enabled on the DHCP server, it combines the host name from the client and this domain name to create the FQDN (fully-qualified domain name) that it uses to update DNS. For information about DDNS, see [Chapter 20, Configuring DDNS Updates from DHCP](#), on page 689.
 - **DNS Servers:** Click the Add icon. Grid Manager adds a row to the table. In the table, enter the IP address of the DNS server to which the DHCP client sends name resolution requests. The DHCP server includes this information in the DHCP OFFER and DHCP ACK messages.
 - **Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached. When configuring this for a template, enter the offset value of the broadcast IP address of the network.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Defining IPv4 Option Spaces

DHCP members support the DHCP option space by default. You can create additional option spaces to provide additional configuration or service information. Note that custom options defined in a user-defined option space are always encapsulated in option 43 in DHCP messages.

To add a custom option space:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Click the Add icon -> **IPv4 Option Space**.
3. In the *Option Space* wizard, do the following:
 - **Name:** Enter the name of the option space.
 - **Comment:** Enter useful information about the option space.
 - **Options:** Click the Add icon to add options. For additional information, see the next section, [Configuring Custom DHCP Options](#) on page 803.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create an option space and add options to it, you can apply the options as described in [Applying DHCP Options](#) on page 804.

Configuring Custom DHCP Options

You can define custom options in the DHCP option space or in an option space that you configured, as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Select either the **DHCP** option space or an IPv4 option space that you configured, and then click the Edit icon.
3. In the *Option Space* editor, click the Add icon to add a custom option. In the new row, complete the following:
 - **Name:** Enter the name of the custom DHCP option.
 - **Code:** Select an option code from the drop-down list. Select a number between 126 and 254 if you are adding custom options to the **DHCP** option space. If you are adding custom options to an IPv4 option space you configured, you can enter a number between 1 and 254.
 - **Type:** Select the option type (such as ip-address, text, boolean, and string as described in [Table 24.1](#)).
For example, to create an option that defines the IP addresses of Solaris root servers, enter the name SrootIP4, select option code 126, and then select the type as ip-address.Click the Add icon to add more options.
4. Save the configuration.

Applying DHCP Options

Some options may apply to all networks and some may apply to specific ranges and even hosts. When you apply an option, you select the object to which the option is applied, such as the Grid member, or network, and then specify a value for the option.

Use the following guidelines when specifying option values:

- Enter **false** or **true** for a Boolean Flag type value.
- Enter an ASCII text string, or enter a series of octets specified in hex, separated by colons.
- Separate multiple values by commas. For example, to enter multiple IP addresses for netbios-name-servers, enter a comma between each IP address.

Here are some examples of option names and correctly formatted values:

Option name	Value	Comment
option 61 dhcp-client-identifier	MyPC	Double quotes are no longer needed for string type values
dhcp-client-identifier	43:4c:49:45:54:2d:46:4f:4f	Series of octets specified in hex, separated by colons for a Data-string type value
netbios-name-servers	10.1.1.5,10.1.1.10	Multiple IP addresses separated by commas
option-80	ABC123	Custom option number 80 set to the string ABC123.

To apply DHCP options:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.

Member: From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Members** → *member* check box, and then click the Edit icon.

Network: From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* check box, and then click the Edit icon.

DHCP Range: From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* → *addr_range* check box, and then click the Edit icon.

Fixed Address: From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* → *fixed_address* check box, and then click the Edit icon.

Reservation: From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* → *reservation* check box, and then click the Edit icon.

Host Address: From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* → *host_record* check box, and then click the Edit icon. Select the host IP address, and then click the Edit icon.

Roaming Host: From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Roaming Hosts** → *roaming_host* check box, and then click the Edit icon.

2. In the *DHCP Properties* editor, select the **IPv4 DHCP Options** or **DHCP** tab and complete the following:
 - The **Custom DHCP Options** section displays two fields. The first field displays **Choose option**. Click the arrow and select an option from the list. In the second field, enter a value for the selected option. Note that certain options have predefined data types and their values must be entered in a specific format. For information about the data types, see [DHCP Option Data Types](#) on page 800.
 - Click **+** to add another option, or click **-** to delete a previously specified option. When overriding an option, enter the new value for the selected option.
 - Note that if you created an option space as described in [Defining IPv4 Option Spaces](#) on page 802, this section displays a list of option spaces in the first drop-down menu, so you can select the option space of the option you want to define.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuration Example: Defining a Custom Option

In this example, you configure two custom options in the DHCP option space, and apply them to a DHCP range in the network 192.168.2.0/24.

Add the custom options to the DHCP options space:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters/Option Spaces** tab.
2. Click the **Option Spaces** subtab to display the panel, click the **DHCP** check box, and then click the Edit icon.
3. In the *Option Space* editor, click the Add icon. In the new row, complete the following:
 - **Name:** Enter **tftp-server**.
 - **Code:** Enter **150**.
 - **Type:** Select **array of ip-address**.
4. Click the Add icon to add another option. In the new row, complete the following:
 - **Name:** Enter **pxe-configfile**.
 - **Code:** Enter **209**.
 - **Type:** Select **text**.
5. Click **Save & Close**.

Enter values for the newly defined custom options and apply them to a DHCP range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** subtab, and click the 192.168.2.0/24 network.
2. Click the 192.168.2.10 - 100 check box, and then click the Edit icon.
3. In the *DHCP Properties* editor, select the **DHCP** tab and complete the following in the **Custom DHCP Options** section:
 - From the drop-down list of options, select **tftp-server (150) array of address**. In the second field, enter **192.168.1.2**.
Click + to add another option.
 - From the drop-down list of options, select **pxe-configfile (209) text**. In the second field, enter **pxe.config**, which is the file name of the boot image.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

The member then includes options 150 and 209 in its DHCP messages to clients that are allocated IP addresses from the DHCP range 192.168.2.10 - 100.

Defining Option 60 Match Rules

The appliance uses option 60 (vendor-class-identifier) to forward client requests to the DHCP server for services that the clients require. You can define option 60 match rules and filter on these rules. You can set these rules for the Grid and override for a member.

To define option 60 for the Grid or member:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **DHCP Options** tab -> **Advanced** tab and complete the following:
To override the Grid configuration for a member, click **Override** next to the property. Grid Manager hides the Grid configuration. You can then add new values for the member.

- **Option 60 (Vendor Class Identifier) Match Rules:** Click the Add icon if you want to add a match rule to a vendor class option. The appliance adds a row to the table. Complete the following:
 - **Option Space:** Select an option space from the drop-down list. This field appears only when you have custom option spaces. The appliance uses the default **DHCP** option space if you do not have custom option spaces.
 - **Match Value:** Enter the value you want the appliance to use when matching vendor class options.
 - **Is Substring:** Select this check box if the match value is a substring of the option data.
 - **Substring Offset:** Enter the number of characters at which the match value substring starts in the option data. Enter 0 to start at the beginning of the option data, enter 1 for the second position, and so on. For example, when you enter 2 here and have a match value of RAS, the appliance matches the value RAS starting at the third character of the option data.
 - **Substring length:** Enter the length of the match value. For example, if the match value is SUNW, the length is 4.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

About the DHCP Relay Agent Option (Option 82)

The typical relationship between a DHCP client, relay agent, and server (that is, the NIOS appliance) on a network is as follows:

1. A DHCP client broadcasts a DHCPDISCOVER message on its network segment.
2. A DHCP relay agent on that segment receives the message and forwards it as a unicast message to one or more DHCP servers (such as NIOS appliances).
3. If the NIOS appliance accepts the address request, it responds to the relay agent with a DHCPOFFER message. If the appliance denies the request, it does not send any response in case other DHCP servers that might be involved respond instead.
4. The relay agent forwards the response to the client, usually as a broadcast message.

The situation is different for individual hosts connecting to the Internet through an ISP, usually over a circuit-switched data network.

1. A host connects to its ISP's circuit access concentration point, authenticates itself, and requests an IP address.
2. The circuit access unit relays the address request to a DHCP server, which responds with a DHCPOFFER message.

To avoid broadcasting the DHCPOFFER over the network segment on which the host made the request, the relay agent sends the response directly to the host over the established circuit.

Option 82 assists the agent in forwarding address assignments across the proper circuit. When a relay agent receives a DHCPDISCOVER message, it can add one or two agent IDs in the DHCP option 82 suboption fields to the message.

The two relay agent IDs are:

- **Circuit ID:** This identifies the circuit between the remote host and the relay agent. For example, the identifier can be the ingress interface number of the circuit access unit (perhaps concatenated with the unit ID number and slot number). The circuit ID can also be an ATM virtual circuit ID or cable data virtual circuit ID.
- **Remote ID:** This identifies the remote host. The ID can be the caller ID telephone number for a dial-up connection, a user name for logging in to the ISP, a modem ID, and so on. Because the remote ID is defined on the relay agent, which is presumed to have a trusted relationship with the DHCP server, and not on the untrusted DHCP client, the remote ID is also presumably a trusted identifier.

Note: For information about the relay agent option, refer to *RFC3046, DHCP Relay Agent Information Option*.

The NIOS appliance can screen address requests through a relay agent filter you set up using option 82. For information, see [About Relay Agent Filters](#) on page 899.

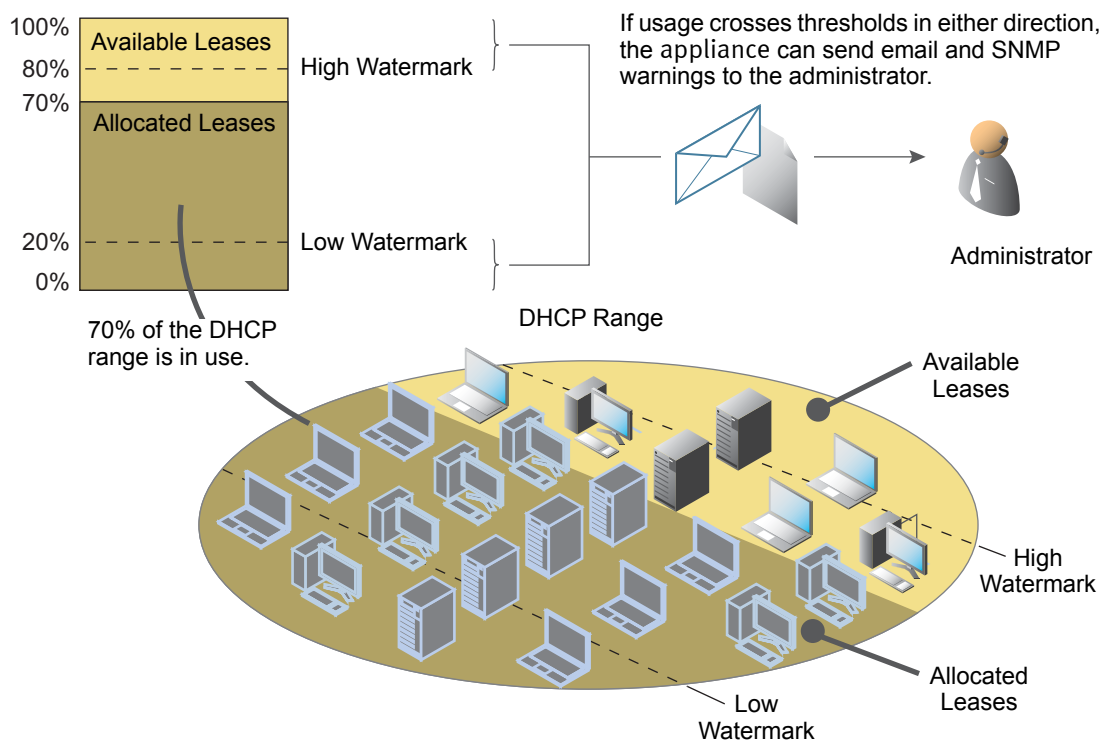
You can also use the relay agent information (circuit ID or remote ID) as a host identifier when configuring a fixed address, though you cannot do so in a host record. For information about how to configure a circuit ID or remote ID as an identifier, see [Adding IPv4 Fixed Addresses](#) on page 858.

CONFIGURING THRESHOLDS FOR DHCP RANGES

Grid Manager can provide a view of the current overall DHCP range usage for the DHCP ranges defined on each Grid member. The view is in the form of a percent: address leases in use/total addresses for each network. Such information can indicate if there is a sufficient number of available addresses at each of these levels. It can also provide information about the distribution of address resources, indicating if there are too many unused addresses in one location while all the addresses in another are in use.

In addition to viewing the percent of addresses in use, you can also apply high and low thresholds for each DHCP range. These watermarks represent thresholds above or below which DHCP range usage is unexpected and might warrant your attention. For example, usage falling below a low threshold might indicate network issues preventing the renewal of leases. When usage for a DHCP range crosses a threshold, the appliance makes a syslog entry and—if configured to do so—sends the administrator alerts as SNMP traps and email notifications. [Figure 24.2](#) illustrates the relationship of allocated and available addresses to high and low watermarks in a DHCP range.


Figure 24.2 Overall DHCP Address Usage for a DHCP Range



You can define watermarks at the Grid, member, network, and DHCP range levels, but the appliance applies them solely to DHCP ranges. Because the appliance applies settings hierarchically in a parent-child structure, by defining watermarks once at a higher level, DHCP ranges can then inherit these settings without your needing to redefine them for each range. For example, if you set high and low watermarks for a Grid, then each Grid member, each network, and each DHCP range inherits these settings. However, if you override these settings at the member level, then the network and DHCP ranges for that member inherit its settings. If you override the Grid member settings at the network level, then that network and any DHCP ranges within that network inherit the network-level settings. Finally, you can set high and low watermarks for an individual DHCP range, which override anything set at a higher level.

[Figure 24.3](#) shows different high and low watermark settings at different levels. Although you can set thresholds at four levels (Grid, Grid member, network, and DHCP range), the NIOS appliance applies them to DHCP ranges.

Figure 24.3 High and Low Watermarks



Address usage in a DHCP range can trigger an event and an email notification when it crosses a watermark. You must enable DHCP threshold and email warnings to receive events and notifications. The following are actions that do and do not trigger an address usage event and notification:

Address usage triggers an event and the appliance sends a notification when the percentage of the allocated addresses in the DHCP range:

- Exceeds the high watermark
- Drops below or equals to the high watermark after exceeding it
- Drops below the low watermark
- Exceeds the low watermark after dropping below it

Address usage does not trigger an event when the percentage of the allocated addresses in the DHCP range:

- Never exceeds the low watermark
- Initially exceeds the low watermark
- Reaches a watermark but does not cross it

Note: You can effectively disable address usage events for a DHCP range by setting its high watermark at 100% and the low watermark at 0% (default setting for the low watermark). Because address usage cannot cross these watermarks, no events can occur.

You can configure the threshold settings at the Grid level and override them at the member, network, and DHCP range levels. To override an inherited DHCP property, click **Override** next to the property to enable the configuration. For information, see [Overriding DHCP Properties](#) on page 783.

To configure thresholds:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.

Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.

Network: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.

DHCP Range: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.

2. In the *DHCP Properties* editor, select the **IPv4 Thresholds** tab and complete the following:
 - **Enable DHCP Thresholds:** Select **Enable DHCP Thresholds** to enable the DHCP threshold feature.
 - **High:** Enter a number between 0 and 100. Enter Trigger and Reset values. If the percentage of allocated addresses in a DHCP range exceeds the Trigger value, the appliance makes a syslog entry and—if configured to do so—sends an SNMP trap and an email notification to a designated destination. When the percentage first reaches the Reset value after it hit the Trigger value, the appliance sends an SNMP trap and an email notification to a designated destination. The default Trigger value is 95, and the default Reset value is 85.
 - **Low:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below the Trigger value, the appliance makes a syslog entry and—if configured to do so—sends an SNMP trap and an email notification to a designated destination. When the percentage first reaches the Reset value after it hit the Trigger value, the appliance sends an SNMP trap. The default Trigger value is 0 and the default Reset value is 10.
 - **Enable SNMP Warnings:** Select this for the appliance to send an SNMP trap to the trap receiver that you define for the Grid when the address usage in a DHCP range crosses a high or low mark threshold.
 - **Enable Email Warnings:** Select this for the appliance to send an email notification to an administrator if the address usage in a DHCP range crosses a high or low mark threshold.
 - **Email Addresses:** Click **Override** to override the Grid administrator email address configured in the **Data Management** tab -> **Grid** tab. This address is not hierarchically inherited from the Grid DHCP configuration. Click the Add icon, and then enter an email address to which you want the appliance to send email notifications when the DHCP range for the network crosses a threshold. You can create a list of email addresses.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING DHCPV6 PROPERTIES

The following sections describe how to configure properties and options that apply to DHCPv6 objects only. You can configure and define the following DHCP properties:

- General properties, as described in the next section, [Defining General IPv6 Properties](#).
- DHCP options, as described in [About DHCPv6 Options](#).

Defining General IPv6 Properties

You can configure general DHCPv6 properties at the Grid level and override them at the member and lower levels.

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.

Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.

Network: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.

Fixed Address: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *Fixed address* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **IPv6 DHCP Options** tab, and complete the following:
 - **Valid Lifetime:** Specify the length of time addresses that are assigned to DHCP clients remain in the valid state. When this time expires, an address becomes invalid and can be assigned to another interface.
 - **Preferred Lifetime:** Specify the length of time that a valid address is preferred. A preferred address can be used with no restrictions. When this time expires, the address becomes deprecated.

- **Domain Name:** Enter the name of the domain for which the Grid serves DHCP data.
- **DNS Servers:** Click the Add icon. Grid Manager adds a row to the table. In the table, enter the IPv6 addresses of DNS recursive name servers to which the DHCP client can send name resolution requests. The DHCP server includes this information in the DNS Recursive Name Server option in Advertise, Rebind, Information-Request, and Reply messages.

3. Save the configuration.

ABOUT DHCPv6 OPTIONS

DHCPv6 options provide configuration and service information to IPv6 clients. Just like IPv4 options, IPv6 options appear as variable length fields at the end of the DHCPv6 messages.

Just as in IPv4, the NIOS appliance supports the following options in the DHCPv6 options space:

- **Predefined options:** These are the option codes defined in RFC 3315. You cannot redefine these options or delete them from the DHCP option space. Option codes 1-48 are reserved and cannot be used to define custom options.
- **Custom options:** These are option codes 49 to 254. They are not defined by IETF standards and are available for private use. You can use these option codes to provide configuration or service information that none of the predefined options provide.

You can also create option spaces to define new groups of options. For example, you can create additional option spaces to define vendor specific options, which are encapsulated in DHCPv6 option 17. When an IPv6 client requests vendor specific options, it makes a request using the vendor specific options (option 17). The DHCP server then responds with the list of replies for the various options encapsulated into option 17.

Note that custom options defined in the DHCP option space are included in the options section of the DHCP messages that DHCP servers and clients exchange.

You can apply options globally at the Grid level, or more specifically at the member, network, range, host and roaming host levels.

Configuring DHCPv6 Options

To use DHCPv6 options, you can do the following:

- Configure one or more option spaces, as described in the next section [Defining IPv6 Option Spaces](#).
- Define custom options in the predefined DHCPv6 option space or add options to an option space that you configured. For more information, see [Configuring Custom IPv6 DHCP Options](#) on page 811.
- Specify values for the options and apply them to the Grid, or to a member, network, fixed address, host, or roaming host. For more information, see [Applying DHCPv6 Options](#) on page 811.

Defining IPv6 Option Spaces

DHCP members support the DHCPv6 option space by default. You can create additional option spaces to provide additional configuration or service information.

To add a custom option space:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Click the Add icon -> **IPv6 Option Space**.
3. In the *IPv6 Option Space* wizard, do the following:
 - **Name:** Enter the name of the option space.
 - **Enterprise Number:** Enter the vendor's Enterprise Number that is registered with IANA.
 - **Comment:** Enter useful information about the option space.

- **Options:** Click the Add icon to add options. For additional information, see the next section, [Configuring Custom DHCP Options](#) on page 803.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create an option space and add options to it, you can apply the options as described in [Applying DHCP Options](#) on page 804.

Configuring Custom IPv6 DHCP Options

You can define custom options in the DHCP option space or in an option space that you configured, as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Select either the **DHCPv6** option space or an IPv6 option space that you configured, and then click the Edit icon.
3. In the *Option Space* editor, click the Add icon to add a custom option. In the new row, complete the following:
 - **Name:** Enter the name of the custom DHCP option.
 - **Code:** Enter a number from 1 to 65535 to add a custom option in the DHCP option space or in an IPv6 option space that you have configured.
 - **Type:** Select the option type (such as ipv6-address, text, boolean, and string as described in [Table 24.1](#)).

Click the Add icon to add more options.

4. Save the configuration.

Applying DHCPv6 Options

You can apply some options at the Grid or member level, and some options to specific networks, shared networks, fixed addresses and roaming hosts. When you apply an option, you select the object to which the option is applied, such as the Grid, member, or network, and then specify a value for the option.

Use the following guidelines when specifying option values:

- Enter **false** or **true** for a Boolean Flag type value.
- Enter an ASCII text string, or enter a series of octets specified in hex, separated by colons.
- Separate multiple values by commas. For example, to enter multiple IP addresses for netbios-name-servers, enter a comma between each IP address.

DHCPv6 options support the same data types as DHCP IPv4 options. For more information about the data types, see [DHCP Option Data Types](#) on page 800.

To apply DHCP options:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then select **Grid DHCP Properties** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
Network: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
Fixed Address: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed_address* check box, and then click the Edit icon.
Host Address: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *host_record* check box, and then click the Edit icon. Select the host IP address, and then click the Edit icon.
Roaming Host: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming_host* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **IPv6 DHCP Options** or **DHCP** tab and complete the following:

- The **Custom IPv6 DHCP Options** section displays two fields. The first field displays **Choose option**. Click the arrow and select an option from the list. In the second field, enter a value for the selected option. Note that certain options have predefined data types and their values must be entered in a specific format. For information about the data types, see [DHCP Option Data Types](#) on page 800.

Click **+** to add another option, or click **-** to delete a previously specified option. When overriding an option, enter the new value for the selected option.

Note that if you created an option space as described in [Defining IPv4 Option Spaces](#) on page 802, this section displays a list of option spaces in the first drop-down menu, so you can select the option space of the option you want to define.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

CONFIGURING DHCP IPV4 AND IPV6 COMMON PROPERTIES

This section describes DHCP properties that apply to both IPv4 and IPv6. It includes the following sections:

- [Configuring UTF-8 Encoding for Hostnames](#)
- [Associating Networks with Zones](#) on page 813
- [Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges](#) on page 814
- [Configuring Fixed Address Leases For Display](#) on page 814

Configuring UTF-8 Encoding for Hostnames

When you configure the appliance as a DHCP server, the appliance supports UTF-8 encoding of hostnames that are encoded with Microsoft Windows code pages. You can configure the DHCP services on the appliance to convert these client hostnames to UTF-8 characters. The appliance stores the UTF-8 encoded hostnames in the database. If you also configure the DHCP services on the appliance to perform DDNS updates, the appliance sends the UTF-8 encoded host names in the DDNS updates. You can configure the UTF-8 encoding of host names at the Grid DHCP service and member DHCP service levels. For information on UTF-8 encoding, see [Printing from Grid Manager](#) on page 91.

The appliance displays the host names in their original characters in the following:

- DHCP lease history
- DHCP lease details
- IP address management
- Syslog
- Audit log

To configure UTF-8 encoding for hostnames:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then select **Grid DHCP Properties** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **General Basic** tab and complete the following:
 - **IPv4 Properties**
 - **Microsoft Clients Code Page:** From the drop-down list, select the code page with which the host names are encoded when the appliance converts the Microsoft code page encoded host names to UTF-8 characters.

— IPv6 Properties

- **Microsoft Clients Code Page:** From the drop-down list, select the code page with which the host names are encoded when the appliance converts the Microsoft code page encoded host names to UTF-8 characters.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Associating Networks with Zones

You can associate IPv4 and IPv6 networks with DNS zones to limit the zones that admins can use when they create DNS records for IP addresses in the networks. When a network is associated with one or more zones and an admin creates a DNS record for one of its IP addresses, Grid Manager allows the admin to create the DNS record in the associated zones only. For example, if you associate the 10.1.0.0/16 network with the corp100.com zone, admins are allowed to create DNS records in the corp100.com zone only for IP addresses in the 10.1.0.0/16 network; or if you associate the 2001:db8:1::/48 network with the corp200.com zone, admins are allowed to create DNS records in the corp200.com zone only for IP addresses in the 2001:db8:1::/48 network.

This feature applies to A, AAAA and host records only. It does not apply to records in a shared record group. If you are creating a host record with multiple IP addresses in different networks, the networks must be associated with the zone of the host record.

If a network is not associated with a zone, admins can create DNS records for its IP addresses only in zones with no network associations as well.

You can associate a network with any authoritative zone whose primary server is a Grid member or a Microsoft server, or is unassigned. You cannot associate networks with zones that have external primary servers.

You can associate a network with multiple zones, and associate a zone with more than one network. You can associate IPv4 and IPv6 network containers and networks with zones. When you associate a network container with zones, its networks inherit the zone associations. You can override the zone associations at the network level.

If you split a network, the resulting subnets inherit the zone associations. If you join networks, the resulting network retains the zone associations of the network that you selected when you performed the join operation. You can override the inherited zone associations of individual networks. Subzones do not inherit the network associations of their parent zones.

When you import data into a zone that is associated with a list of networks, the imported A, AAAA and host records must have IP addresses in the associated networks. Grid Manager does not allow you to import A, AAAA and host records with IP addresses in unassociated networks.

When you associate a network with a zone, the DNS records created before the association are not affected. But if you edit an A, AAAA or host record after the association, Grid Manager does not allow you to save the record if its IP address is not in an associated network.

To associate an IPv4 or IPv6 network with a zone:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
2. In the *DHCP Network* editor, click **Toggle Advanced Mode** if the editor is in basic mode.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. Click the Add icon and select the zone you want to associate with the network.
 - Optionally, select a default zone. When you create or edit an A, AAAA or host record from a network in the **IPAM** tab, Grid Manager automatically selects the default zone that is assigned to the network.
5. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Viewing the Networks Associated with a Zone

You can view the IPv4 or IPv6 networks associated with a zone from the zone editor. The tab to display network associations in zone editors is visible only if the primary server is a Grid member, a Microsoft server, or unassigned.

To view the network associations of a zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* check box, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Toggle Advanced Mode** if the editor is in basic mode.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.

The Network Associations table lists the networks and their corresponding comments. You cannot change the network associations in this editor. Navigate to the *DHCP Network* editor of the network, to change the zone associations.

Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges

You can configure the DHCP server to store leases in a deleted DHCP range for up to one week after the leases expire. When you add a new DHCP range that includes the IP addresses of these leases or assign the DHCP range to another member within the Grid, the appliance automatically restores the active leases. You can configure this feature for the Grid, and override the configuration for members, networks, and DHCP ranges.

To keep active leases in a deleted DHCP range:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
Network: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
Range: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *range* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor of the Grid or member, click **Toggle Advanced Mode** if the editor is in basic mode, and then click the **General** tab -> **Advanced** tab. In the Network editor or Range editor, click **Toggle Advanced Mode** if the editor is in basic mode, and then click **IPv4 DHCP Options** -> **Advanced** or **IPv6 DHCP Options** -> **Advanced**. Complete the following:
 - **IPv4 Properties**
 - **Lease Deletion:** When you select **Keep leases from deleted range until one week after expiration** and delete a DHCP range with active leases, the appliance stores these leases for up to one week after they expire.
 - **IPv6 Properties**
 - **Lease Deletion:** When you select **Keep leases from deleted range until one week after expiration** and delete a DHCP range with active leases, the appliance stores these leases for up to one week after they expire.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring Fixed Address Leases For Display

You can configure the DHCP server to capture the hostname and lease time of a fixed address when you assign an IPv4 or IPv6 fixed address to a client. The appliance displays the hostname, and the start and end time of each fixed address lease in the *Current Leases* panel in Grid Manager.

You can set this at the Grid level only for IPv4 and IPv6 leases.

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **General** tab -> **Advanced** tab and complete the following:
 - **IPv4 Properties**
 - **Fixed Address Lease:** Select **Capture hostname and lease time when assigning Fixed Addresses**. The appliance displays the host name, and the start and end time of each fixed address lease in the *Current Leases* panel. If there are multiple records (A, host, and lease) for the IP address, it also displays the information for the records. This option is available in the Grid Properties editor only.)
 - **IPv6 Properties**

- **Fixed Address Lease:** Select **Capture hostname and lease time when assigning Fixed Addresses**. The appliance displays the host name, and the start and end time of each fixed address lease in the *Current Leases* panel. If there are multiple records (AAAA, host, and lease) for the IP address, it also displays the information for the records. This option is available in the Grid Properties editor only.)

3. Save the configuration.

CONFIGURING DHCP LOGGING

If you have a syslog server operating on your network, you can specify in which facility you want the server to display the DHCP logging messages. You can also select the Grid member on which you want to store the DHCP lease history log, as described in the next section [Configuring the Lease Logging Member](#). You can configure DHCP and lease logging only on the Grid and member levels.

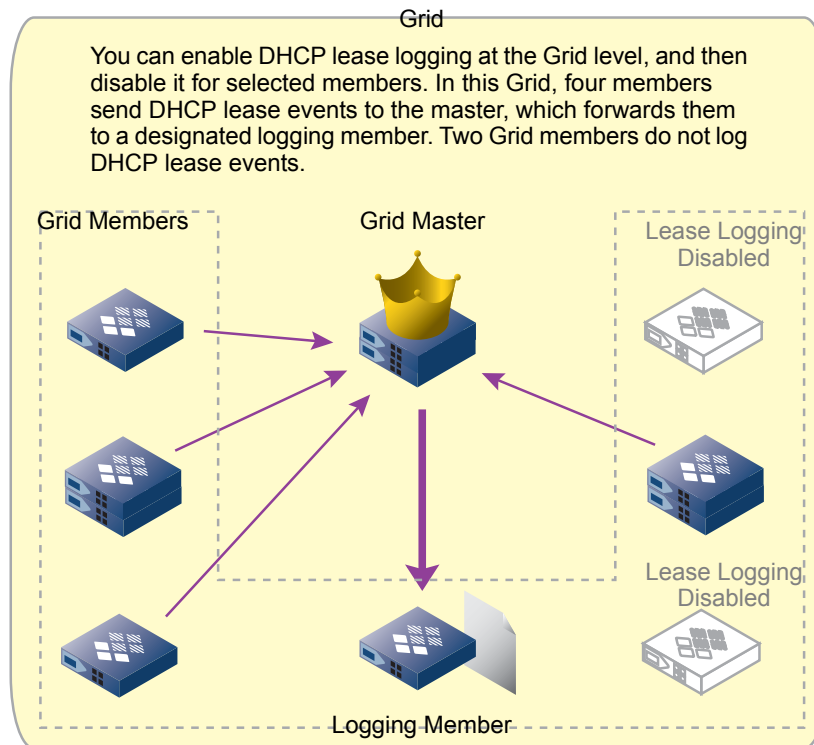
To specify DHCP logging for the Grid or member:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **Logging Basic** tab and complete the following:
 - **Syslog Facility:** From the drop-down list, select the facility that is used to tag syslog messages from the DHCP server. This facility can be used to filter messages on a central syslog server.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring the Lease Logging Member

Logging DHCP lease events makes significant CPU demands, especially when there is heavy DHCP activity. Therefore, Infoblox strongly recommends that you designate a Grid member other than the master as a logging member whenever possible. Another way to manage the increased load that logging introduces is to log selectively per Grid member. For example, you might want to log DHCP leases for members serving critical parts of your network and not keep historical logs for members serving other parts.

Figure 24.4 DHCP Lease History Logging with Member Overrides



To specify lease logging for a member:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Configuration** from the Toolbar.

Member: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.

2. In the **Logging** tab, complete the following:
 - **Lease Logging:** Select **Enable Lease History** (for Grid) or **Log Lease Events from DHCP Server** (for member) to enable DHCP lease logging. To disable DHCP lease logging, clear the check box. You can set member overrides if you want to enable or disable lease logging per member.
 - **Send leases to:** For Grid only. Click **Select**. In the *Select Member* dialog box, select the Grid member on which you want to store the DHCP lease history log. Infoblox recommends that you dedicate a member other than the Grid Master as a logging member. If possible, use this member solely for storing the DHCP lease history log. If you do not select a member, no logging can occur.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: You cannot configure vNIOs Grid members on Riverbed as DHCP lease history logging members.

4. For information about viewing current leases, see [Viewing Current Leases](#) on page 946

ABOUT IF-MAP

You can configure Infoblox DHCP servers to publish DHCP data to an IF-MAP server. The IF-MAP server takes real-time information from different sources and stores it in a shared database from which clients can retrieve information about network devices, their status and activities. For details about the IF-MAP protocol, refer to <http://www.trustedcomputinggroup.org>. For information about the Infoblox IF-MAP server, refer to the *Infoblox Administrator Guide for Infoblox Orchestration Server*.

Each Infoblox DHCP server in a Grid can function as an IF-MAP client, with the ability to publish lease information to an IF-MAP server. For information about how to configure an IF-MAP client, see [Configuring Members as IF-MAP Clients](#) on page 818. You can configure the client to publish ip-mac and ip-duid (for DHCPv6 leases) metadata at the Grid and member levels. You can also configure the client to publish metadata for specific leases by overriding the Grid or member publishing settings at the network (IPv4 and IPv6) or range (IPv4 only) level. The DHCP server sends updates to the IF-MAP server using the XML format and SOAP/HTTPS bindings specified in IF-MAP v1.1r5 and v2.0r26. The DHCP server supports the IF-MAP 2.0 protocol by default. You can also enable the support for IF-MAP 1.1, as described in [Configuring a Grid to Support IF-MAP](#).

When the DHCP server grants an IPv4 lease and sends the DHCPACK packet to the DHCP client, it updates the link in the IF-MAP server between the leased IP address and client MAC address with ip-mac metadata with the following attributes: start-time, end-time, and dhcp-server. The dhcp-server attribute contains the DHCP server hostname. The ip-mac metadata is attached to a link with:

- An ip-address identifier with the type attribute set to IPv4, a value attribute that contains the leased IP address, and the administrative-domain attribute set to the network view to which the IP address belongs.
- A mac-address identifier with a value attribute that contains the client MAC address. It does not have the administrative-domain attribute.

When the DHCP server grants an IPv6 lease and sends the Reply message to the DHCP client, it updates the link in the IF-MAP server between the leased IP address and client DHCP Unique Identifier (DUID) with ip-duid metadata that contains the following attributes: start-time, end-time, and dhcp-server. The dhcp-server attribute contains the DHCP server hostname. The ip-duid metadata is attached to a link with:

- An ip-address identifier with the type attribute set to IPv6, a value attribute that contains the leased IP address, and the administrative-domain attribute set to the network view to which the IP address belongs.
- A duid identifier with a value attribute that contains the client DUID. It does not have the administrative-domain attribute.

The Infoblox DHCP server also publishes data when an IPv4 or IPv6 lease changes. When a lease is released or when an active lease expires, the DHCP server sends a “publish delete” request to the IF-MAP server.

You can define how the IF-MAP server handles the existing ip-mac and ip-duid information before the DHCP client sends the next update. For example, you can specify the IF-MAP server to always delete existing ip-mac and ip-duid information before the next update. For information, see [Deleting Existing Data Before Publishing](#) on page 821.

Following are the tasks to enable DHCP servers in a Grid to function as IF-MAP clients:

1. Enable IF-MAP in the Grid and specify the URL and port of the IF-MAP server, as described in [Configuring a Grid to Support IF-MAP](#) on page 817.
2. Optionally, enable the validation of the IF-MAP server certificate and import the CA certificate, as described in [Validating the IF-MAP Server Certificate](#) on page 818.
3. Enable IF-MAP on each Grid member and specify an authentication method the member uses to connect to the IF-MAP server, as described in [Configuring Members as IF-MAP Clients](#) on page 818.
4. Optionally, override publishing settings at the member, network, or range level, as described in [Overriding IF-MAP Publishing Settings](#) on page 820.

You can also delete DHCP data published by a specific member, or define how the IF-MAP server deletes existing DHCP data before a client publishes an update. For information, see [Deleting Data from the IF-MAP Server](#) on page 821.

Configuring a Grid to Support IF-MAP

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab, and then complete the following:
 - **Enable IF-MAP:** Select this check box to enable the IF-MAP service for the Grid. Note that you must enable the IF-MAP service in order to enable or disable publishing at the Grid, member, network, or range level.
 - **IF-MAP Server URL:** Enter the URL of the IF-MAP server to which the Grid members publish DHCP data. The URL must begin with **http://** or **https://**; for example, **https://<server_ip_addr>/ifmap**.

- **IF-MAP Server Port:** The default HTTP port is 80 and the default HTTPS port is 443. Optionally, you can specify a different port on the IF-MAP server.
 - **Enable IF-MAP publishing:** Select this check box to enable IF-MAP publishing for the Grid. When you select this, IF-MAP publishing is enabled for all members, networks (IPv4 and IPv6), and DHCP ranges (IPv4 only). You can override the Grid property at a specific level to control the ip-mac and ip-duid metadata you want the client to publish for specific leases. For information, see [Overriding IF-MAP Publishing Settings](#) on page 820.
 - **IF-MAP Protocol Version:** Select the IF-MAP protocol version you want the IF-MAP client to use to connect to the IF-MAP server. The default is IF-MAP 2.0.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
 5. You can also configure how the IF-MAP server deletes existing metadata before the IF-MAP client publishes another update. For information, see [Deleting Data from the IF-MAP Server](#) on page 821.

Validating the IF-MAP Server Certificate

You can configure the IF-MAP client to validate the IF-MAP server certificate before the client establishes a connection or performs IF-MAP transactions. To validate an IF-MAP server certificate, you must first import the certificate of the CA that signs the IF-MAP server certificate.

To configure the IF-MAP client to validate the IF-MAP server certificate:

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab and complete the following:
 - **Enable IF-MAP:** Select this check box to enable the IF-MAP service for the Grid.
 - **Enable IF-MAP server certificate validation:** Select this check box to enable the validation of the IF-MAP server certificate, and then click **Import** to import the CA certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**. You can also copy and paste the CA certificate here.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring Members as IF-MAP Clients

To configure a member to be an IF-MAP client, you must first enable IF-MAP on the member and then configure a client authentication method. The IF-MAP client can authenticate itself to the IF-MAP server through user name and password credentials or digital certificate. Note that each member must have unique credentials or certificates. You cannot use the same credentials or certificates on multiple members. The appliance supports only one CA-signed certificate on each member. If you want to use a roll-over certificate, you must replace the existing certificate and restart services on the member.

To enable an appliance to function as an IF-MAP client:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Member DHCP Properties* dialog box, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab and complete the following:
 - **Enable IF-MAP:** Select this check box to enable the IF-MAP service on the member. Note that you must enable the IF-MAP service in order to enable or disable publishing at the network and range levels.
 - **Authentication:** Select one of the following authentication methods:
 - **Certificate:** Select this to use the IF-MAP client certificate for client authentication. You must already have a certificate configured for the member before you can select and save this configuration. For information about creating a client certificate, see [Creating IF-MAP Client Certificates](#) on page 819.
 - **Basic:** Select this to use username and password credentials for IF-MAP client authentication. Complete the following:

- **Username:** Enter the username the member uses to connect to the IF-MAP server. This username must have been configured as a valid username on the IF-MAP server. Each member must have its own username.
- **Password:** Enter the password the member uses to connect to the IF-MAP server.
- **Confirm Password:** Enter the password again.

Note: When you upgrade to a new NIOS release, the basic authentication credentials are retained if IF-MAP was enabled and basic authentication was used before the upgrade.

- **Enable IF-MAP publishing:** Click **Override** to override the Grid setting. Select this check box to enable IF-MAP publishing for all the networks that are served by this member. Ensure that you enable IF-MAP at either the Grid or member level in order to enable IF-MAP publishing for all networks.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Creating IF-MAP Client Certificates

Before you can select “Certificate” as the client authentication method, you must first create a certificate for the specified member.

You can do one of the following to generate an IF-MAP client certificate:

- Generate a self-signed certificate and save it. For information, see [Generating Self-Signed Certificates](#).
- Request a CA (Certificate Authority) signed certificate. When you receive the certificate from the CA, upload it to the member that you configure as an IF-MAP client. For information, see [Generating Certificate Signing Requests](#) on page 820.

Generating Self-Signed Certificates

You can replace the default certificate with a self-signed certificate that you generate. When you generate a self-signed certificate, you can specify the correct hostname and change the public/private key size, enter valid dates and specify additional information specific to the member. If you have multiple members, you can generate a certificate for each appliance with the appropriate hostname.

To generate a self-signed certificate:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click **IF-MAP Client Certificate** -> **Generate Self-signed Certificate** from the Toolbar.
2. In the *Generate Self-Signed Certificate* dialog box, complete the following:
 - **Key Size:** Select either **2048** or **1024** for the length of the public key.
 - **Days Valid:** Specify the validity period of the certificate.
 - **Common Name:** Specify the domain name of the member. You can enter the FQDN (fully qualified domain name) of the appliance.
 - **Organization:** Enter the name of your company.
 - **Organizational Unit:** Enter the name of your department.
 - **Locality:** Enter a location, such as the city or town of your company.
 - **State or Province:** Enter the state or province.
 - **Country Code:** Enter the two-letter code that identifies the country, such as US.
 - **Admin E-mail Address:** Enter the email address of the appliance administrator.
 - **Comment:** Enter information about the certificate.
3. Click **OK**.
4. If the appliance already has an existing client certificate, the new certificate replaces the existing one. In the *Replace IF-MAP Certificate Confirmation* dialog box, click **Yes**.

Generating Certificate Signing Requests

You can generate a CSR (certificate signing request) that you use to obtain a signed certificate from your own trusted CA. Once you receive the signed certificate, you can import it to the member, as described in [Uploading Certificates](#). To generate a CSR:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click **IF-MAP Client Certificate** -> **Create Signing Request** from the Toolbar.
2. In the *Create Certificate Signing Request* dialog box, enter the following:
 - **Key Size:** Select **2048** or **1024** for the length of the public/private key pair.
 - **Common Name:** Specify the domain name of the member. You can enter the FQDN of the appliance.
 - **Organization:** Enter the name of your company.
 - **Organizational Unit:** Enter the name of your department.
 - **Locality:** Enter a location, such as the city or town of your company.
 - **State or Province:** Enter the state or province.
 - **Country Code:** Enter the two-letter code that identifies the country, such as US.
 - **Admin E-mail Address:** Enter the email address of the appliance administrator.
 - **Comment:** Enter information about the certificate.
3. Click **OK**.

Uploading Certificates

When you receive the certificate from the CA, the appliance finds the matching CSR and takes the private key associated with the CSR and associates it with the newly imported certificate. The appliance then automatically deletes the CSR.

To import a certificate:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click **IF-MAP Client Certificate** -> **Upload Certificate** from the Toolbar.
2. Navigate to where the certificate is located and click **Open**.
3. If the appliance already has an existing IF-MAP client certificate, the new certificate replaces the existing one. In the *Replace IF-MAP Certificate Confirmation* dialog box, click **Yes**.

Downloading Certificates

You can download the current certificate or a self-signed certificate.

To download a certificate:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click **IF-MAP Client Certificate** -> **Download Certificate** from the Toolbar.
2. Navigate to where you want to save the certificate, enter the file name, and then click **Save**.

Overriding IF-MAP Publishing Settings

When you enable IF-MAP publishing at the Grid level, all members, networks (IPv4 and IPv6), and DHCP ranges (IPv4 only) in the Grid inherit the same setting. To control which ip-mac and ip-duid metadata is published for specific leases that belong to a specific network or address range, you can override the Grid settings at a specific member, network, or range level. Note that you must first enable the IF-MAP service at the Grid and member levels in order to enable or disable IF-MAP publishing at other levels. For example, if you want the DHCP server to publish IF-MAP data for specific leases in a specific network that is served by a specific member, you must first enable the IF-MAP service at the Grid and member levels, as described in [Configuring a Grid to Support IF-MAP](#) on page 817. Then, you can enable IF-MAP publishing at the range level, as described in this section.

Though you can configure and save the settings of IF-MAP publishing any time at any level, the publishing does not actually happen unless the IF-MAP service is enabled at the Grid or member level. If a network or DHCP range is served by a specific member and you want to enable IF-MAP publishing for the network or range, you must first enable the IF-MAP service for the specified member.

To override IF-MAP publishing settings:

1. **Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
Network: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network* check box, and then click the Edit icon.
DHCP Range: From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network* -> *addr_range* check box, and then click the Edit icon.
2. In the editor, click **Toggle Advanced Mode**, and then click the **IF-MAP** tab.
3. Click **Override** and complete the following:
 - **Enable IF-MAP Publishing:** Select this check box to instruct the DHCP server to publish metadata to the IF-MAP server when the IF-MAP service is enabled for the Grid or member. Clear this check box so the DHCP server does not publish metadata to the server.

Deleting Data from the IF-MAP Server

The appliance allows you to delete IF-MAP data from the IF-MAP server. You can delete all IF-MAP data published by a specific member. You can also define how the IF-MAP server handles the deletion of existing metadata before the IF-MAP client publishes another update.

Deleting all data

You can delete all IF-MAP data published by a specified member. To delete data published by all members in a Grid, you must delete data for each member individually.

To delete IF-MAP data published by a member from the IF-MAP database:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab, and then click **Clear** -> **IF-MAP Data** from the Toolbar.
2. In the *Purge IF-MAP Data* dialog box, click **Select Member** to select a member. If there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select one. Click the member name in the dialog box, and then click **Purge** to delete all the DHCP data published by the Grid member. You can also click **Clear** to clear the displayed member and select a new one.

Deleting Existing Data Before Publishing

You can define how the IF-MAP server deletes existing metadata before an IF-MAP client publishes new data. You can configure the IF-MAP client to instruct the server to always delete existing data, never delete it, or delete the data before a specified time period.

To define how the IF-MAP server deletes DHCP data before the next publish:

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab and complete the following:
 - **Enable IF-MAP:** Select this check box to enable the IF-MAP service.
 - **Delete existing metadata:** You can define how the IF-MAP server deletes the existing metadata before the IF-MAP client publishes new data. Select one of the following:
 - **Always delete:** Select this to always delete existing metadata before the IF-MAP client publishes updates. This is the default.
 - **Do not delete:** Select this to never delete the existing metadata before the IF-MAP client publishes updates.

- **Earlier than:** Select this to delete metadata that was published before a given time before the IF-MAP client publishes updates. When you select this option, enter a time value, and then select a time unit from the drop-down list.

4. Save the configuration and click **Restart** if it appears if it appears at the top of the screen.

STARTING DHCP SERVICES ON A MEMBER

The DHCP service is disabled by default. After you complete the DHCP configuration, you can start DHCP service on a member. To enable the member to provide DHCPv6 service as well, you must start the DHCP service and then enable the DHCPv6 service on the member. In addition, you must specify the DHCP Unique Identifier (DUID) of the member. IPv6 clients use DUIDs to identify the source of the DHCP messages from servers.

To start DHCP service on a member:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Members** → *member* check box.
2. Expand the Toolbar and click **Start**.
3. In the *Start Member DHCP Service* dialog box, click **Yes**.
4. Grid Manager starts DHCP on the selected member.

You can stop DHCP service on a member by selecting the member check box and click **Stop** from the Toolbar. This will stop DHCP service enabled on the LAN port.

To stop DHCP service enabled on the LAN2 port:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Members** → *member* check box.
2. Click the Edit icon.
3. In the *Member DHCP Properties* editor, select the **General Basic** tab.
4. Clear the check box for LAN2 under DHCP interfaces.
5. Save the configuration.

To enable DHCPv6 service on the member:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → *member* check box.
2. In the *Member DHCP Properties* editor, select the **General Basic** tab.
3. In the **IPv6 Properties** section, do the following:
 - **Server DUID:** Enter the DUID of the member.
 - **Enable DHCPv6 Service:** Select this check box.
4. Save the configuration.

VIEWING DHCP MEMBER STATUS

You can view DHCP member status after you configure DHCP properties and start or stop DHCP services on a member.

To view member status:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Members** section.
2. Grid Manager displays the following information:
 - **Name:** The name of the Grid member.
 - **Status:** The status of the DHCP services on the member. This can be one of the following:
 - **Not Running:** DHCP services have not been started on the member.
 - **Running:** The DHCP services are running properly on the member.
 - **Warning:** The member is connecting or synchronizing with its Grid Master.

- **Error:** The member is offline, is not licensed (that is, it does not have a DNSone license with the Grid upgrade that permits Grid membership), is upgrading or downgrading, or is shutting down.

Note: You can mouse over on the informational icon next to the status to view detailed information.

- **Comment:** The information you entered for the member.
- **IPv4 DHCP Utilization:** The percentage of the total IPv4 DHCP utilization of the member. This is the percentage of the total number of DHCP hosts, fixed addresses, reservations, and leases assigned to the member versus the total number of IP addresses (excluding IP addresses in the exclusion range) and all DHCP objects assigned to the member. Note that only enabled objects are included in the calculation. The appliance updates the utilization data every 15 minutes. The appliance displays the utilization data in one of the following colors:
 - Red: The DHCP resources are 100% utilized.
 - Yellow: The utilization percentage is over the effective high watermark threshold.
 - Blue: The utilization percentage is below the effective low watermark threshold.
 - Black: The utilization percentage is at any number other than 100%, or within the effective thresholds.
- **Site:** The site to which the member belongs. This is one of the predefined extensible attributes.

You can select the following additional columns for display:

- **Address:** The IP address of the member.
- **Static Addresses:** The number of static IP addresses.
- **Dynamic Addresses:** The number of dynamically assigned IP addresses.
- **IF-MAP Connection:** The status of the IF-MAP service connection on the member. This can be one of the following.
 - **Stopped:** The DHCP or IF-MAP service on the member is stopped, or the IF-MAP service is not enabled.
 - **Running:** The IF-MAP client is connected to the IF-MAP server and the IF-MAP service is running properly.
 - **Failed:** The IF-MAP client cannot publish data to the IF-MAP server due to some errors.
 - **Warning:** Some non-fatal errors occurred. The IF-MAP client attempts to reconnect to the server.

Note: You can mouse over on the informational icon next to the status to view detailed information, including the status description and the timestamp when the status initially changed.

- **IF-MAP Last Update:** The timestamp the status of the IF-MAP service was last updated. For example, if the IF-MAP connection status is **Running** and this field shows 2011-11-20 12:30:42 EST, it means that an IF-MAP operation, such as a publish, was last completed on November 20, 2011 at 12:30:42 Eastern Standard Time.

To view status information about the IF-MAP connection on an independent appliance, from the **Data Management** tab -> **DHCP** tab, click **System DHCP Properties** from the toolbar. The appliance displays the following:

- **IF-MAP Connection:** The status of the IF-MAP service on the independent appliance. A color icon associated with the connection status appears before the status.
- **IF-MAP Connection Information:** Detailed information about the status. On a Grid member, this information appears when you mouse over on the informational icon.
- **IF-MAP Last Update:** The timestamp when the status of the IF-MAP service last changed.

Note: For more information about these fields, see descriptions about Grid member status in this section.

You can view detailed information about a specific member by clicking the member link. Grid Manager displays the following information about the selected member:

- **Network:** The network assigned to the member.
- **Comment:** The information about the network.

- **IPv4 DHCP Utilization:** The percentage of the DHCP usage of the network. This is the percentage of the total number of fixed addresses, reservations, hosts, and active leases on the network over the total IP addresses in the range, excluding the number of addresses on the network. Note that only enabled objects are included in the calculation.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes.

In the member panel, you can select the following additional fields for display:

- **Disabled:** Indicates whether the member is disabled or not.
- **IPAM Utilization:** When you define a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network.

When you define a network container that contains subnets, this is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. For example, when you define a /16 network and then 64 /24 networks underneath it, the /16 network container is considered 25% utilized even when none of the IP addresses in the /24 networks is in use.

You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.

- Extensible attributes that associate with the network.

You can also sort the data in ascending or descending order by column. For information, see [Customizing Tables](#) on page 60.

Viewing DHCP Configuration Files

You can view the IPv4 and IPv6 DHCP configuration of a selected member. The format of the configuration file depends on the browser you use.

To view the DHCP configuration of a selected member:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box.
2. Expand the Toolbar, select **View DHCP Configuration**, and then select either **IPv4** or **IPv6**. Grid Manager displays the IPv4 or IPv6 DHCP configuration of the selected member in a new browser. You can print and save the file using the corresponding functions in your browser.



Chapter 25 Managing DHCP Templates

This chapter explains how to configure and manage IPv4 and IPv6 DHCP templates. It contains the following sections:

- [About DHCP Templates](#) on page 826
- [About IPv4 DHCP Templates](#) on page 826
 - [About IPv4 Range Templates](#) on page 826
 - [About IPv4 Fixed Address/Reservation Templates](#) on page 828
 - [About IPv4 Network Templates](#) on page 829
 - [Configuration Example: Creating an IPv4 Network Using a Template](#) on page 832
- [About IPv6 DHCP Templates](#) on page 834
 - [About IPv6 Range Templates](#) on page 834
 - [About IPv6 Fixed Address Templates](#) on page 835
 - [About IPv6 Network Templates](#) on page 837
- [Viewing Templates](#) on page 839
- [Deleting Templates](#) on page 839

ABOUT DHCP TEMPLATES

A template contains a set of predefined properties that you use to create IPv4 and IPv6 DHCP objects. It is metadata that you can modify and reuse. Using a template enables you to create objects in a quick and consistent way. You can define the object properties once in a template, and then create multiple objects that inherit their properties from the template. For example, you can create a network template that has a fixed netmask of /24 and extensible attribute “State” set to California. You can then use the template to create networks in California that contain /24 netmasks. You can also modify and delete a template. Note that modifying or deleting a template does not affect existing objects created based on the template. You must be a superuser or have read/write permissions to add, modify, or delete a template. A superuser can set other admin group privileges on templates. For information, see [Administrative Permissions for IPv4 or IPv6 DHCP Templates](#) on page 211. You can also define extensible attributes for these templates when you create them. For information, see [Using Extensible Attributes](#) on page 332.

ABOUT IPv4 DHCP TEMPLATES

You can use templates to create DHCP IPv4 ranges, fixed addresses, reservations, roaming hosts, and networks. You can create the following IPv4 templates:

- A DHCP range template, containing DHCP range settings, such as the total number of IP addresses allocated to a range. You can add a DHCP range template to a network template. For information, see [About IPv4 Range Templates](#) on page 826.
- A fixed address/reservation template, containing information for creating fixed addresses, reservations, or roaming hosts. You can add a fixed address/reservation template to a network template. For information, see [About IPv4 Network Templates](#) on page 829.
- A network template, containing basic network properties for creating networks. It is also a container that holds your DHCP range templates and fixed address/reservation templates. When you create a network using a network template, the network inherits the properties of the range and fixed address/reservation templates. You can create a network in any network view using a network template. For information, see [About IPv4 Network Templates](#) on page 829.

Because you can potentially add DHCP range and fixed address/reservation templates to a network template, create the DHCP range and fixed address/reservation templates before you create a network template. For information, see [Configuration Example: Creating an IPv4 Network Using a Template](#) on page 832.

About IPv4 Range Templates

When you create an IPv4 range template, the start and end address fields are based on the specified offset from the network start address and the number of IP addresses in the range. After you create a DHCP range template, you can configure additional properties such as exclusion ranges and DHCP filters, as described in [Modifying IPv4 Range Templates](#) on page 827. Then when you use the template to create a DHCP range, the range inherits the properties of the template. You can also include a DHCP range template in a network template to automatically create a DHCP range when you use that network template.

Adding IPv4 Range Templates

To create an IPv4 DHCP range template:

1. From the **Data Management** tab, select the **DHCP** tab → **Templates** tab, and then expand the Toolbar and click **Add** → **Templates** → **Range** → **IPv4**.
2. In the *Add IPv4 Range Template* wizard, do the following:
 - **Name:** Enter a name that helps identify the DHCP range template. For example, enter **Region 1 IT** if you want to use this template to create DHCP ranges for the IT department in Region 1.

- **Offset:** An offset in a DHCP range template determines the starting IP address of the range. The appliance adds the offset value you enter here to the start IP address of the network in which you create a DHCP range using this template. That IP address becomes the start IP address of the DHCP range. For example, you specify an offset value of 25 for a 25.0.0.0/8 network using the DHCP range template, the appliance creates a DHCP range with the start IP address of 25.0.0.25 in the network.
 - **Number of Addresses:** Enter the total number of IP addresses to be included in the DHCP range.
 - **Comment:** Enter useful information about the template.
3. Click **Next** and select one of the following to provide DHCP services for the range:
 - **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
 - **Grid Member:** Click **Select** and choose a Grid member from the drop-down list.
 - **Failover Association:** Click **Select** and choose a failover association. Only failover associations that provide DHCP services in the network view of the DHCP range appear in the drop-down list.
 - **Microsoft DHCP Server:** Click **Select** and choose a Microsoft server from the drop-down list. The drop-down list displays only the servers that are associated with the network to which the DHCP range belongs.
 4. Click **Next** to configure or override DHCP options as described in [Defining Basic IPv4 Options](#) on page 801.
 5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 6. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying IPv4 Range Templates

After you use the wizard to create an IPv4 DHCP range template, you can set additional properties for the template. Following are some guidelines:

- In the **DHCP Options** tab of a DHCP range template, the broadcast address is an address offset number rather than a broadcast IP address; network router addresses are offset numbers as well.
An offset in a DHCP range template indicates the starting IP address of the DHCP range object created from the template. For example, you can create a network template called *test_network_template* and a DHCP range template *test_range_template* linked to this network template. If the *test_range_template* has an offset value 10, when you create a 10.0.0.0/8 network using the *test_network_template*, the appliance creates a DHCP range with the starting IP address 10.0.0.10. If you create a 20.0.0.0/8 network using the *test_network_template*, the appliance creates a DHCP range with the starting IP address 20.0.0.10.
- For the exclusion range in the template, the start and end addresses are determined by the number of offsets in the DHCP range template's start address and the number of IP addresses in the exclusion range. For more information about exclusion ranges, see [About DHCP Ranges](#) on page 780.

To modify and set properties for a DHCP range template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Edit icon.
2. The *DHCP Range Template* editor contains the following tabs from which you can modify data:
 - **General:** Modify general information described in [Adding IPv4 Range Templates](#) on page 826.
 - **Member Assignment:** Change the Grid member, failover association, or Microsoft server that provides DHCP services for this template. You can also add or delete a member or failover association. For information, see [Adding IPv4 Address Ranges](#) on page 854.
 - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with this template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

- **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#) on page 160.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the template. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.
 - **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use for dynamic address assignments. Complete the following:
 - **Offset:** An offset for an exclusion range determines the start IP address of the exclusion range. The appliance adds the offset value you enter here to the start IP address of the DHCP range created using this template. That IP address becomes the start IP address of the exclusion range.
 - **Number of Addresses:** Enter the number of IP addresses to be included in the exclusion range.
 - **Comment:** Enter useful information about the exclusion range.
 - **IPv4 Thresholds:** Keep the inherited thresholds settings or override them and enter unique settings for the template. For information, see [Configuring Thresholds for DHCP Ranges](#) on page 807.
 - **IPv4 Filters:** Add DHCP filters to the Class Filter List and Logic Filter List. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.
 4. Save the configuration and click **Restart** if it appears at the top of the screen.

Note: Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

About IPv4 Fixed Address/Reservation Templates

You can use an IPv4 fixed address/reservation template to create fixed addresses, reservations and roaming hosts. When you create an IPv4 fixed address/reservation template, you can specify an offset and number of addresses. This is used when you include the template in a network template. When you include a fixed address/reservation template in a network template, the DHCP server automatically creates reservations based on the offset and number of addresses you specified in the fixed/address reservation template. It does not create fixed addresses.

After you create a fixed address/reservation template using the wizard, you can configure additional properties as described in [Modifying IPv4 Fixed Address/Reservation Templates](#) on page 829. Then when you use the template to create a fixed address, it inherits the properties of the template.

Adding IPv4 Fixed Address/Reservation Templates

To create an IPv4 fixed address/reservation template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** section.
2. Click the Add drop-down list and select **IPv4 Fixed Address/Reservation Template**.
3. In the *Add IPv4 Fixed Address/Reservation Template* wizard, enter the following:
 - **Name:** Enter a name that helps identify the fixed address/reservation template. For example, you can enter **HP Printer** when you create a template that contains settings for assigning fixed addresses or reservations to HP printers.
 - **Comment:** Optionally, enter additional information about the template.

In the **Optional Settings For Range of Objects** section, do the following:

- **Offset:** An offset in a fixed address/reservation template determines the start IP address of the object created from the template. The appliance adds the offset value you enter here to the start IP address of the network in which you create objects using this template. That IP address becomes the start IP address of the object.
- **Number of Addresses:** Enter the number of IP addresses to be used as fixed addresses, reservations, or roaming hosts.

Note: The appliance uses the offset and number of addresses only when this template is used in a network template.

4. Click **Next** to configure or override DHCP options as described in [Defining Basic IPv4 Options](#) on page 801.
5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying IPv4 Fixed Address/Reservation Templates

To modify a fixed address/reservation template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Edit icon.
2. The *Fixed Address/Reservation Template* editor contains the following tabs from which you can modify data:
 - **General:** Modify general information for the template as described in [Adding IPv4 Fixed Address/Reservation Templates](#) on page 828.
 - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#) on page 160.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the template. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

About IPv4 Network Templates

You can create IPv4 network templates to facilitate network configuration. You can use network templates to create networks in any network view. When you create a network template, you do not specify a network address. You enter the network address when you create an actual network from the template. You can specify a netmask or allow the user to define the netmask when they create the actual network.

A network template is useful for setting up a network with fixed addresses and DHCP ranges already defined. You can add DHCP range or fixed address/reservation templates to a network template. Once the fixed address and DHCP range information is set up, the network template contains a range template list and a fixed address/reservation template list.

When you enable support for RIR updates, you can create IPv4 network templates specific for RIR associated networks. For information about RIR updates, see [RIR Registration Updates](#) on page 437.

Adding IPv4 Network Templates

To create a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** section.
2. Click the Add drop-down list and select **IPv4 Network Template**.
3. In the *Add IPv4 Network Template* wizard, do the following:

- **Regional Internet Registry:** This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#) on page 437. Complete the following to create a network template for an RIR IPv4 network container or network:
 - **Internet Registry:** Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
 - **Organization ID:** Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
 - **Registration Status:** The default is **Not Registered**. When using this template to add an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** check box below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**.
 - **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you use this template to add an existing RIR allocated network to NIOS, select **None**. When you use this template to add networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
 - **Do not update registrations:** Select this check box if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
 - **Name:** Enter a name that helps identify the network template. For example, you can enter **Class C** if you want to configure the template for creating Class C networks.
 - **Netmask:** Select one of the following options:
 - **Fixed:** Select this and adjust the netmask slider to a fixed netmask for this network template. When you select this option, users cannot specify another netmask when they use this template to create a network. For example, if you select /24 as the fixed netmask, all networks created using this template have a /24 netmask.
 - **Allow User to Specify Netmask:** Select this to allow users to specify the subnet mask when creating networks using this template.
 - **Comment:** Optionally, enter additional information about the template.
 - **Automatically Create Reverse-Mapping Zone:** This function is enabled if the fixed netmask of the template equals /8, /16, and /24, or if you select the **Allow User to Specify Netmask** option. Select this if you want the appliance to automatically create the corresponding reverse-mapping zone for the networks created using this template. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network level.
4. Click **Next** and do the following to assign either Grid members or Microsoft DHCP servers to this network template. Ensure that you include members or Microsoft servers that are associated with other templates that you plan to add to this network template. You can assign one or multiple members to this template. However, you cannot assign a combination of NIOS Grid members and vNIOS Grid members to the template. You can also assign multiple Microsoft servers to a template, but you cannot assign a mix of Microsoft servers and Grid members to a template.
- click the Add icon and select one of the following options:
 - **Add Infoblox Member:** Select this option to add a Grid member as a DHCP server for the networks created using this template. Select the Grid member from the *Member Selector* dialog box. Keep in mind, DHCP properties for the network are inherited from this member. Networks created using this template can be served by multiple members, but a member can serve networks in one network view only.

or

- **Add Microsoft Server:** Select this option to add a Microsoft server as a DHCP server for the networks created using this template. Select the Microsoft server from the *Microsoft Server Selector* dialog box.
5. Click **Next** and do the following to include IPv4 address range and fixed address/reservation templates in the network template. Note that when you select a fixed address/reservation template, only reservations, not fixed addresses, are created for networks created using this template. You cannot add a fixed address/reservation template that does not contain an offset value or a total number of IP addresses for a range.
 - a. Click the Add icon.
 - b. In the *DHCP Template Selector* dialog box, choose the template that you want to include in this network template. You can choose a DHCP range or fixed address/reservation template. Use SHIFT+click and CTRL+click to select multiple templates.
 - c. Click the Select icon.

You can delete a template from the table by selecting it and clicking the Delete icon.
 6. Click **Next** to configure or override DHCP options as described in [Defining Basic IPv4 Options](#) on page 801.
 7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#) on page 447.
 8. Save the configuration and click **Restart** if it appears at the top of the screen.

Modifying IPv4 Network Templates

To modify a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Edit icon.
2. The *IPv4 Network Template* editor contains the following tabs from which you can modify data:
 - **General:** Modify general information described in [Adding IPv4 Network Templates](#) on page 829.
 - **Member Assignment:** Change the Microsoft servers or Grid members that provide DHCP services for this template. For information, see [Adding IPv4 Networks](#) on page 845.
 - **Templates:** Add or delete DHCP range and fixed address/reservation templates. For information, see [About IPv4 Range Templates](#) on page 826 and [About IPv4 Fixed Address/Reservation Templates](#) on page 828.
 - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **RIR Registration:** Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#) on page 443.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the template. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.
 - **IPv4 Thresholds:** Keep the inherited thresholds settings or override them and enter unique settings for the template. For information, see [Configuring Thresholds for DHCP Ranges](#) on page 807.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

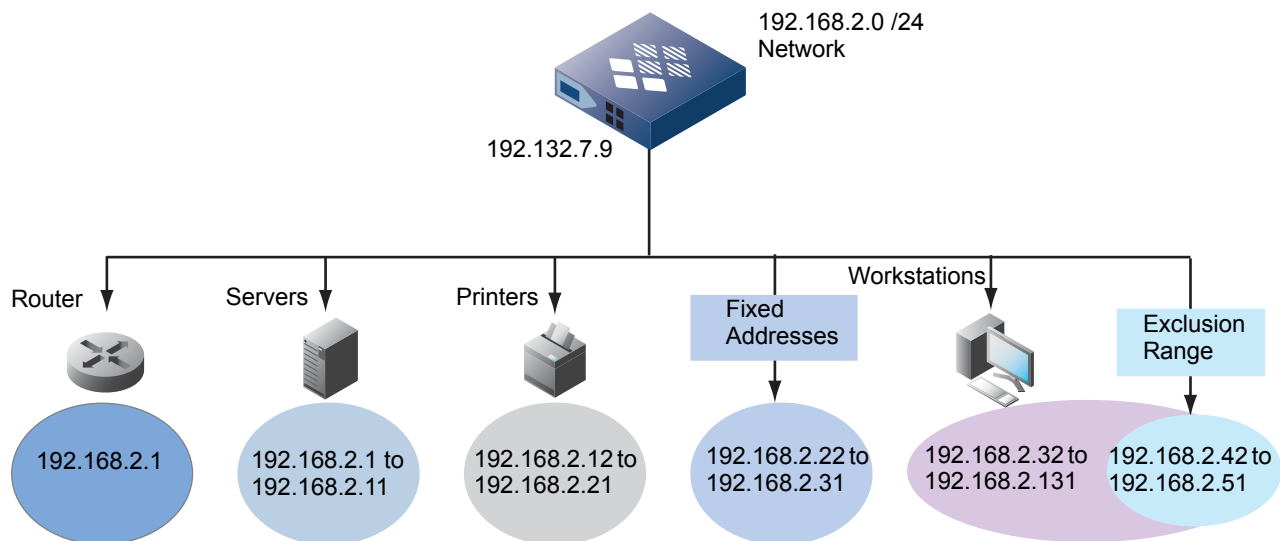
Configuration Example: Creating an IPv4 Network Using a Template

This example describes how to create a /24 network template and how to use the template to create a 192.168.2/24 network with the following configurations:

- First address 192.168.2.1 is reserved for the router
- Next 10 addresses (192.168.2.2 to 192.168.2.11) reserved for servers
- Next 10 addresses (192.168.2.12 to 192.168.2.21) reserved for printers
- Next 10 addresses (192.168.2.22 to 192.168.2.31) assigned as fixed addresses
- 100 addresses (192.168.2.32 to 192.168.2.131) reserved for workstations. The appliance assigns these dynamically.
- 10 addresses (192.168.2.42 to 192.168.2.51) are in an exclusion range. If you assigned static addresses to certain hosts in the middle of an address range template, you can exclude the addresses from the address range template so the appliance does not assign these IP addresses to clients.

[Figure 25.1](#) illustrates the configurations of the 192.168.2/24 network using the network template you create:

Figure 25.1 Creating a Network Using a Template



Use the following steps to create the sample network template (shown in [Figure 25.1](#)).

1. Create the following DHCP range templates. For information, see [Adding IPv4 Range Templates](#) on page 826.
 - Server template with the following values:
 - **Name:** Servers
 - **Offset:** 2
 - **Number of Addresses:** 10
 - **Comment:** Address range 2 to 11 for Servers
 - Printer template with the following values:
 - **Name:** Printers
 - **Offset:** 12
 - **Number of Addresses:** 10
 - **Comment:** Address range 12 to 21 for printers.

- Workstation template with the following values:
 - **Name:** Workstations
 - **Offset:** 32
 - **Number of Addresses:** 100
 - **Comment:** Address range 32 to 131 for DHCP on workstations
 - Exclusion range with the following values. You must modify the *Workstations* template to add the exclusion range. For information, see [Modifying IPv4 Range Templates](#) on page 827.
 - **Name:** Exclusion
 - **Offset:** 42
 - **Number of Addresses:** 10
 - **Comment:** Excluding addresses 42 to 51 from the DHCP range 32 to 131.
2. Create a fixed address/reservation template with the following values. For information, see [Adding IPv4 Fixed Address/Reservation Templates](#) on page 828.
 - **Name:** Router
 - **Comment:** Fixed address template
 - **Offset:** 1
 - **Number of Addresses:** 1
 3. Create a fixed address/reservation template with the following values. For information, see [Adding IPv4 Fixed Address/Reservation Templates](#) on page 828.
 - **Name:** myFixedAddress
 - **Comment:** Fixed address template
 - **Offset:** 22
 - **Number of Addresses:** 10
 4. Create a network template with the following values. For information, see [Adding IPv4 Network Templates](#) on page 829.
 - **Name:** myNetworkTemplate
 - **Netmask:** Select /24 as the fixed subnet mask for the network
 - **Comment:** Network template for /24 network
 - **Automatically create a reverse-mapping zone:** Select this so that the NIOS appliance automatically creates the corresponding reverse-mapping zone for the network.
 5. Add the DHCP range templates *Servers*, *Printers*, and *Workstations* to the network template.
 6. Add the fixed address/reservation template *myFixedAddress* to the network template.
 7. Add a fixed address with the following values:
 8. Create a network using the network template *myNetworkTemplate* with the following values. For information, see [Adding IPv4 Networks](#) on page 845.
 - **Address:** Enter the IP address 192.168.2.0 of the network that you want to create using the template.
 - **Select template:** Select the network template *myNetworkTemplate*.
 9. To verify your configuration, from the **Data Management** tab, select the **DHCP** tab -> **Templates** tab. Select *myNetworkTemplate* and click the Edit icon. In the *Network Template* editor, click the **Templates** tab. The Grid Manager displays the DHCP range templates and fixed address templates.
 10. Click **Restart** to restart services.

ABOUT IPv6 DHCP TEMPLATES

You can use templates to create DHCP IPv6 ranges, fixed addresses, roaming hosts, and networks. You can create the following IPv6 templates:

- A DHCP range template that specifies an offset and the total number of addresses in a range. You can add a DHCP range template to a network template. For more information, see [About IPv6 Range Templates](#) on page 834.
- A fixed address template, containing information for creating fixed addresses and roaming hosts. You can add a fixed address template to a network template. For information, see [About IPv6 Fixed Address Templates](#) on page 835.
- A network template, containing basic network properties for creating networks. It is also a container that holds your DHCP range templates and fixed address/reservation templates. When you create a network using a network template, the network inherits the properties of the range and fixed address/reservation templates. You can create a network in any network view using a network template. For information, see [Adding IPv6 Network Templates](#) on page 837.

Because you can potentially add DHCP range and fixed address/reservation templates to a network template, create the DHCP range and fixed address/reservation templates before you create a network template.

ABOUT IPv6 RANGE TEMPLATES

You can create range templates to specify an offset and the number of addresses allocated to a range. Note that you cannot create templates for prefix-delegated ranges because the start or end prefix can be outside of the subnet address boundary.

After you create a DHCP range template, you can configure additional properties such as exclusion ranges and DHCP properties, as described in [Modifying IPv6 Range Templates](#). Then when you use the template to create a DHCP range, the range inherits the properties of the template. You can also include a DHCP range template in a network template to automatically create a DHCP range when you use that network template.

Adding IPv6 Range Templates

To create an IPv6 range template:

1. From the **Data Management** tab, select the **DHCP** tab → **Templates** tab.
2. Click the Add drop-down menu and select **IPv6 DHCP Range Template**.
3. In the *Add IPv6 Range Template* wizard, complete the following:
 - **Name:** Enter a name that helps identify the IPv6 DHCP range template.
 - **Offset:** An offset in a DHCP range template determines the starting IP address of the range. The appliance adds the offset value you enter here to the start IP address of the network in which you create a DHCP range using this template. That IP address becomes the start IP address of the DHCP range. For example, you specify an offset value of 10 for the 2001:db8:1263:/48 network using the DHCP range template, the appliance creates a range with the start address 2001:db8:1263:0:0:0:0:a.
 - **Number of Addresses:** Enter the total number of IPv6 addresses to be included in the DHCP range.
 - **Comment:** Optionally, enter additional information about the template.
4. Click **Next** and select one of the following to provide DHCP services for the range:
 - **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
 - **Grid Member:** Click **Select** and choose a Grid member from the drop-down list.

5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Save the configuration.

Modifying IPv6 Range Templates

You can modify the properties of a DHCP range template and define an exclusion range. For the exclusion range in the template, the start and end addresses are determined by the number of offsets in the DHCP range template's start address and the number of IP addresses in the exclusion range. For more information about exclusion ranges, see [About DHCP Ranges](#) on page 780.

To modify a DHCP range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Edit icon.
2. The *IPv6 DHCP Range Template* editor contains the following tabs from which you can modify data:
 - **General:** Modify general information as described in [Adding IPv6 Range Templates](#).
 - **Member Assignment:** Change the Grid member that provides DHCP services for ranges created from this template. For information, see [Adding IPv6 Range Templates](#).
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with this template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#) on page 160.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
 - **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use for dynamic address assignments. **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use for dynamic address assignments. Complete the following:
 - **Offset:** An offset for an exclusion range determines the start IP address of the exclusion range. The appliance adds the offset value you enter here to the start IP address of the DHCP range created using this template. That IP address becomes the start IP address of the exclusion range.
 - **Number of Addresses:** Enter the number of IP addresses to be included in the exclusion range.
 - **Comment:** Enter useful information about the exclusion range.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration.

ABOUT IPV6 FIXED ADDRESS TEMPLATES

A fixed address template is useful when you want to create multiple fixed addresses in a network. When you create a fixed address template, you specify the offset value and number of fixed addresses to be created. You can also specify additional properties for the fixed addresses.

Note that you can use the template to create address-based fixed addresses. You cannot specify prefixes in the template because a fixed address could use a prefix that is not part of the subnet to which the fixed address belongs. You can enter prefixes when you create the individual fixed address objects using the template.

Adding IPv6 Fixed Address Templates

To create an IPv6 fixed address template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Click the Add drop-down menu and select **IPv6 Fixed Address Template**.

3. In the *Add IPv6 Fixed Address Template* wizard, enter the following:
 - **Name:** Enter a name that helps identify the IPv6 fixed address template. For example, you can enter **HP Printer** when you create a template that contains settings for assigning fixed addresses or reservations to HP printers.
 - **Comment:** Optionally, enter additional information about the template.

In the **Optional Settings For Range of Objects** section, do the following:

- **Offset:** An offset in a fixed address template determines the IP address of the first fixed address created from the template. The appliance adds the offset value you enter here to the start IP address of the network in which you create objects using this template, and that IP address becomes the IP address of the object. For example, you specify an offset value of 50 for the 2001:db8:1263:/48 network, when you create a fixed address using the fixed address template, the appliance assigns it the address 2001:db8:1263:0:0:0:32.
- **Number of Addresses:** Enter the number of IP addresses to be used as fixed addresses or roaming hosts.

Note: The appliance uses the offset and number of addresses only when this template is used in a network template.

4. Click **Next** to configure or override DHCP options as described in [Configuring DHCP Properties](#) on page 791.
5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Save the configuration.

Modifying IPv6 Fixed Address Templates

To modify a fixed address template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Edit icon.
2. The *IPv6 Fixed Address Template* editor contains the following tabs from which you can modify data:
 - **General:** Modify general information for the template as described in [Adding IPv6 Fixed Address Templates](#) on page 835.
 - **IPv6 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining General IPv6 Properties](#) on page 809.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#) on page 160.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
 - **IPv6 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration.

ABOUT IPV6 NETWORK TEMPLATES

You can create IPv6 network templates to facilitate network configuration. You can use network templates to create networks in any network view. When you create a network template, you do not specify a network address. You enter the network address when you create an actual network from the template. You can specify a netmask or allow the user to define the netmask when they create the actual network.

A network template is useful for setting up a network with fixed addresses and DHCP ranges already defined. You can add DHCP range or fixed address templates to a network template.

When you enable support for RIR updates, you can create IPv6 network templates specific for RIR associated networks. For information about RIR updates, see [RIR Registration Updates](#) on page 437.

Adding IPv6 Network Templates

To create a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Click the Add drop-down menu and select **IPv6 Network Template**.
3. In the *Add IPv6 Network Template* wizard, do the following:
 - **Regional Internet Registry:** This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#) on page 437. Complete the following to create a network template for an RIR IPv6 network container or network:
 - **Internet Registry:** Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
 - **Organization ID:** Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
 - **Registration Status:** The default is **Not Registered**. When using this template to add an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** check box below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**.
 - **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you use this template to add an existing RIR allocated network to NIOS, select **None**. When you use this template to add networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
 - **Do not update registrations:** Select this check box if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
 - **IPv6 Prefix:** If you are adding a template for a previously defined global IPv6 prefix, you can select it from the list.
 - **Name:** Enter a name that helps identify the network template.
 - **Netmask:** Select one of the following options:
 - **Fixed:** Select this and adjust the netmask slider to a fixed netmask for this network template. When you select this option, users cannot specify another netmask when they use this template to create a network. For example, if you select /24 as the fixed netmask, all networks created using this template have a /24 netmask.
 - **Allow User to Specify Netmask:** Select this to allow users to specify the subnet mask when creating networks using this template.

- **Comment:** Enter useful information about the template.
 - **Automatically create a reverse-mapping zone:** This function is enabled if the fixed netmask of the template is a multiple of 4 (4, 8, 24, and so on), or if you select the **Allow User to Specify Netmask** option. Select this if you want the appliance to automatically create the corresponding reverse-mapping zone for the networks created using this template. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network level.
4. Click **Next** to assign Grid members to this network template. Ensure that you include members that are associated with other templates that you plan to add to this network template. You can assign one or multiple members to this template. However, you cannot assign a combination of NIOS Grid members and vNIOS Grid members to the template.
 - Click the Add icon to add a Grid member as a DHCP server for the networks created using this template. Select the Grid member from the *Member Selector* dialog box. Keep in mind, DHCP properties for the network are inherited from this member. Networks created using this template can be served by multiple members, but a member can serve networks in one network view only.
 5. Click **Next**, and then click the Add icon to include DHCP range and fixed address templates in the network template. Choose the template that you want to include in this network template. Use SHIFT+click and CTRL+click to select multiple templates.
You can remove a template from the list by selecting the template and clicking the Delete icon.
 6. Click **Next** to configure or override DHCP options as described in [Defining General IPv6 Properties](#) on page 809.
 7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#) on page 447.
 8. Save the configuration.

Modifying IPv6 Network Templates

To modify and set the properties of a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Edit icon.
2. The *Network Template* editor contains the following tabs from which you can modify data:
 - **General:** Modify general information described in [Adding IPv6 Network Templates](#) on page 837.
 - **Member Assignment:** Change the Grid members that provide DHCP services for networks created from this template. For information, see [Adding IPv6 Networks](#) on page 871.
 - **Templates:** Add or delete DHCP range and fixed address templates. For information, see [Adding IPv6 Range Templates](#) on page 834 and [Adding IPv6 Fixed Address Templates](#) on page 835
 - **IPv6 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining General IPv6 Properties](#) on page 809.
 - **RIR Registration:** Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#) on page 443.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
 - **IPv6 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration.

VIEWING TEMPLATES

To view a list of all IPv4 and IPv6 DHCP templates:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Grid Manager displays the following information:
 - **Name:** The name of the template.
 - **Type:** The template type, such as **IPv4 Network Template** or **IPv6 Network Template**.
 - **Comment:** The information you entered about the template.
 - **Site:** The site to which the template belongs. This is one of the predefined extensible attributes.

You can select predefined and user defined extensible attributes for display.

You can also do the following in this panel:

- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Sort the displayed data in ascending or descending order by column.
- Delete a selected template or multiple templates. For information, see [Deleting Templates](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Select an object and edit its information.
- Print or export the data in the panel.

DELETING TEMPLATES

To delete a template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.



Chapter 26 Managing IPv4 DHCP Data

This chapter explains how to configure and manage IPv4 DHCP data. It contains the following sections:

- [Configuring DHCP for IPv4](#) on page 843
- [About the Next Available Network or IP Address](#) on page 844
- [Configuring IPv4 Networks](#) on page 845
 - [Adding IPv4 Networks](#) on page 845
 - [Viewing Networks](#) on page 848
 - [Modifying IPv4 Networks](#) on page 851
 - [Deleting IPv4 Networks](#) on page 852
- [Configuring IPv4 Shared Networks](#) on page 852
 - [Adding IPv4 Shared Networks](#) on page 852
 - [Viewing Shared Networks](#) on page 853
 - [Modifying IPv4 Shared Networks](#) on page 853
 - [Deleting IPv4 Shared Networks](#) on page 854
- [Configuring IPv4 Address Ranges](#) on page 854
 - [Adding IPv4 Address Ranges](#) on page 854
 - [Modifying IPv4 Address Ranges](#) on page 855
 - [Controlling Lease Assignments](#) on page 856
 - [Deleting IPv4 Address Ranges](#) on page 857
- [Configuring IPv4 Fixed Addresses](#) on page 857
 - [Adding IPv4 Fixed Addresses](#) on page 858
 - [Adding IPv4 Fixed Addresses](#) on page 858
 - [Modifying IPv4 Fixed Addresses](#) on page 859
 - [Deleting Fixed Addresses](#) on page 860
- [Configuring IPv4 Reservations](#) on page 860
 - [Adding IPv4 Reservations](#) on page 861
 - [Modifying Reservations](#) on page 862
- [Viewing IPv4 DHCP Objects](#) on page 862
- [About Roaming Hosts](#) on page 863
 - [Configuring Roaming Hosts](#) on page 863
 - [Enabling Support for Roaming Hosts](#) on page 864
 - [Adding IPv4 Roaming Hosts](#) on page 864
 - [Adding IPv6 Roaming Hosts](#) on page 865

- [*Adding IPv4/IPv6 Roaming Hosts*](#) on page 865
- [*Viewing Roaming Hosts*](#) on page 866
- [*Setting Properties for Roaming Hosts*](#) on page 867
- [*Deleting Roaming Hosts*](#) on page 868

CONFIGURING DHCP FOR IPv4

To configure DHCP service for an IPv4 network and the resources in the network, perform the following tasks:

1. Create a network and assign it to Grid members or Microsoft DHCP servers. For information, see [Adding IPv4 Networks](#) and [Modifying IPv4 Networks](#) on page 851.
2. Configure DHCP properties for the network. You can override properties set at the Grid or member level and enter unique values for the network. For information, see [Configuring General IPv4 DHCP Properties](#) on page 793 and [Configuring DHCP IPv4 and IPv6 Common Properties](#) on page 812.
3. Optionally, assign zones to a network. For information, see [Associating Networks with Zones](#) on page 813.
4. Add a DHCP range to the network and assign it to a member, a failover association, or a Microsoft DHCP server. For information, see [Adding IPv4 Address Ranges](#) on page 854 and [Modifying IPv4 Address Ranges](#) on page 855.
5. Optionally, add exclusions to the DHCP range for addresses that are not used for dynamic allocation. For information, see [Configuring IPv4 Fixed Addresses](#) on page 857.
6. Optionally, configure DHCP properties for the address range. You can override properties set at an upper level and enter unique values for the address range. For information, see [Modifying IPv4 Address Ranges](#) on page 855.
7. Optionally, define filters for precise address assignments and apply them to the DHCP range. For information, see [About IPv4 DHCP Filters](#) on page 892.
8. Optionally, add fixed addresses and reservations to the network and configure DHCP properties for them. For information, see [Configuring IPv4 Fixed Addresses](#) on page 857 and [Configuring IPv4 Reservations](#) on page 860.

ABOUT THE NEXT AVAILABLE NETWORK OR IP ADDRESS

When you create certain objects through Grid Manager, the appliance can obtain the next available IPv4 or IPv6 network from a specific network container. It can also obtain the next available IP address from a specific network or address range. This feature automates the allocation of networks and IP addresses so you can manage your network space more efficiently. You can also use this feature to organize network devices. For example, you can create a reserved range called “Printer Range” to reserve static IP addresses for printers in your network. When you allocate IP addresses for printers, you can have the appliance search for the next available IP address within “Printer Range,” and then allocate the next available address to a new printer.

Obtaining the Next Available Network

When you create a new network, you can request the appliance to look for the next available network address within a specific network container. The next available network address is the first unused network address in the network container for which you have administrative permissions. If you are not within a specific network container or network, Grid Manager displays a selector from which you can select the network for the next available network address. For information about creating IPv4 and IPv6 networks using the next available feature, see [Adding IPv4 Networks](#) on page 845 and [Adding IPv6 Networks](#) on page 871.

For information about how the appliance select the next available network, see [Guidelines for the Next Available Network and IP Address](#).

Obtaining the Next Available IP Address

When you create an object such as a fixed address within a network or an address range, you can request the appliance to search for the next available IP address within the network or address range. The next available IP address is the first unused IP address in the specified network, DHCP range, or reserved range. Though you can request the next available IP if you have Read-only permission for the corresponding network or range, you must have Read/Write permission in order to save the configuration using the next available IP address.

To avoid IP address conflicts, the appliance validates the next available IP address to ensure that it is not used for other objects or associated with another operation, such as a scheduled task or an approval workflow. IP addresses associated with scheduled tasks and approval workflows are considered pending and are not available as the next available IP addresses. The appliance preserves all used and pending IP addresses and returns only the next unused IP address as the next available IP. This ensures that the next available IP address is unique and does not overlap with used or scheduled IP addresses in the system.

When multiple users simultaneously request for the next available IP address, the appliance returns the same unused IP address to all users. However, only the user who first saves the configuration succeeds with the next available IP. When other users try to save their configurations using the same IP address, the appliance displays an error message indicating that the IP address is not available. Users can then choose to request a different IP address with the next available IP address function, enter a new one or cancel the operation.

For information about how the appliance selects the next available IP address, see [Guidelines for the Next Available Network and IP Address](#).

Guidelines for the Next Available Network and IP Address

The appliance follows certain rules when searching for the next available network and IP address in the specified wizard, network container, network, and address range.

In a wizard where you can obtain the next available network or IP address, the following applies:

- In a wizard, if you add a network or IP address and then delete it, the appliance excludes it from the next available. When you try to obtain the next available network or IP address in the same wizard, the appliance does not return the deleted network or IP address until you exit the wizard.

In a network, the appliance searches for the next IP address that meets all of the following criteria:

- It does not match any DNS resource record, such as an A, AAAA, PTR record, or host record, that is associated with an IP address.
- It is not assigned to a DHCP fixed address, reservation, or host address record.
- It is not part of any DNS bulk host record.
- It does not match any unmanaged IP address.
- It is not the network (the first) or broadcast (the last) address in the specified network.
- It is not within any DHCP range in this network.
- It is not within any reserved range in this network.
- It is not within an exclusion range.
- It is not part of a scheduled task or approval task that involves all DNS and DHCP objects.

In a DHCP range, the appliance searches for the next IP address that meets all of the following criteria:

- It is not assigned to a fixed address, reservation, or host record.
- It does not match any unmanaged IP address.
- It is not part of an exclusion range within the DHCP range.
- It is not part of a scheduled task or approval workflow that involves DNS and DHCP objects.
- It does not match any active DHCP lease.

In a reserved range, the appliance searches for the next IP address that meets all of the following criteria:

- It is not assigned to a fixed address, reservation, or host record.
- It does not match any unmanaged IP address.
- It is not part of a scheduled task or approval workflow that involves DNS and DHCP objects.

Note: The appliance does not search for deleted leases in the Recycle Bin.

CONFIGURING IPv4 NETWORKS

When you create an IPv4 network, you can do so from scratch or from a network template that contains predefined properties. When you use a template to create a network, the properties of the template apply to the new network. For information about network templates, see [About IPv4 Network Templates](#) on page 829. You can also create an IPv4 network from the Tasks Dashboard, as described in [The Tasks Dashboard](#) on page 99.

After you create IPv4 networks, you can combine them into shared networks or create ranges and fixed addresses.

Adding IPv4 Networks

When you configure an IPv4 network, you must assign either Grid members or Microsoft servers to the network. A network cannot be served by a mix of Microsoft and Infoblox DHCP servers. Multiple servers can serve a network, but Grid members and Microsoft servers cannot serve the same network.

A Grid member can serve only one network view. Similarly, a Microsoft server can serve only one network view. Therefore when you assign Grid members to networks, you must assign the members to networks in the same network view. For information, see [Configuring DHCP for IPv4](#) on page 843.

If you have enabled support for RIR (Regional Internet Registry) updates and are adding an RIR IPv4 network container or network to NIOS, Grid Manager displays an RIR section in the *Add IPv4 Network* wizard. You must enter RIR related information in this section in order for NIOS to associate the newly added network with an RIR organization. For more information about RIR address allocation and updates, see [RIR Registration Updates](#) on page 437.

To add an IPv4 network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab.
2. In the **Networks** section, select **IPv4 Network** from the Add drop-down menu.
3. In the *Add Network* wizard, select one of the following and click **Next**:
 - **Add Network**: Click this to add a network from scratch.
 - **Add Network using Template**: To use a template, click this, and then click **Select Template** and select a network template. For information, see [About IPv4 Range Templates](#) on page 826. The appliance populates the template properties in the wizard when you click **Next**. You can then edit the pre-populated properties, except for **Netmask**.
4. Complete the following and click **Next**:
 - **Regional Internet Registry**: This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#) on page 437. Complete the following to create an RIR IPv4 network container or network:
 - **Internet Registry**: Select the RIR from the drop-down list. The default is **None**, which means that the network is not associated with an RIR organization. When you select **RIPE**, the appliance displays **Organization ID** field where you can select an RIR organization.
 - **Organization ID**: Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
 - **Registration Status**: The default is **Not Registered**. When adding an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** check box below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**. The updated status and timestamp are displayed in the **Status of last update** field in the *IPv4 Network Container* or *IPv4 Network* editor.
 - **Registration Action**: Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you are adding an existing RIR allocated network to NIOS, select **None**. When you are adding networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
 - **Do not update registrations**: Select this check box if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
 - **Network View**: This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
 - **Netmask**: Enter the netmask or use the netmask slider to select the appropriate number of subnet mask bits for the network. The appliance supports /1 to /32 netmasks. Note that when you use a template that contains a fixed netmask, you cannot adjust the netmask for this network.
 Microsoft servers can serve networks with /1 to /31 netmasks. Infoblox DHCP servers can serve networks with /8 to /32 netmasks.
 Since Infoblox DHCP servers do not support /1 to /7 networks, you can assign these networks to Microsoft DHCP servers only. You can create DHCP ranges and fixed addresses within these subnets.
 - **Networks**: Do one of the following to add new networks:
 Click the Add icon to enter a new network. Grid Manager adds a row to the table. Enter the network address in the **Network** field. Click the Add icon again to add another network.
 or
 Click the Next Available icon to have the appliance search for the next available network. Complete the following in the Next Available Networks section:

- **Create new network(s) under:** Enter the network container in which you want to create the new network. When you enter a network that does not exist, the appliance adds it as a network container. When you enter a network that is part of a parent network, the parent network is converted into a network container if it does not have a member assignment or does not contain address ranges, fixed addresses, reservations, shared networks, and host records that are served by DHCP. When you enter a network that has a lower CIDR than an existing network, the appliance creates the network as a parent network and displays a message indicating that the newly created network overlaps an existing network. You can also click **Select Network** to select a specific network in the *Network Selector* dialog box. For information about how the appliance searches for the next available network, see [Obtaining the Next Available Network](#) on page 844.
- **Number of new networks:** Enter the number of networks you want to add to the selected network container. Note that if there is not enough network space in the selected network to create the number of networks specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing networks in the table and you select one, the number you enter here includes the selected network.
- Click **Add Next** to add the networks. Grid Manager lists the networks in the table. You can click **Cancel** to reset the values.

Note: You must click **Add Next** to add the network container you enter in the Next Available Networks section. If you enter a network in the Next Available Networks section and then use the Add icon to add another network, the appliance does not save the network you enter in the Next Available Networks section until you click **Add Next**.

- **Comment:** Enter useful information about the network, such as the name of the organization it serves.
 - **Automatically Create Reverse-Mapping Zone:** This function is enabled if the netmask of the network equals /8, /16, or /24. Select this to have the appliance automatically create reverse-mapping zones for the network. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network view level.
 - **Disabled:** Select this if you do not want the DHCP server to provide DHCP services for this network at this time. This feature is useful when you are in the process of setting up the DHCP server. Clear this after you have configured the server and are ready to have it serve DHCP for this network.
5. Click **Next** and add a Grid member or Microsoft server as a DHCP server for the network. A network can be served by either Grid members or Microsoft servers, but not both at the same time.
- click the Add icon and select one of the following options:
 - **Add Infoblox Member:** Select this option to add a Grid member as a DHCP server for the network. Select the Grid member from the *Member Selector* dialog box. Keep in mind, DHCP properties for the network are inherited from this member. The network can be served by multiple members, but a member can serve networks in one network view only.
- or
- **Add Microsoft Server:** Select this option to add a Microsoft server as a DHCP server for the network. Select the Microsoft server from the *Microsoft Server Selector* dialog box.
6. Click **Next** to override DHCP properties as described in [About DHCP Properties](#) on page 793. This only applies if you are adding a network that is served by an Infoblox Grid member.
- or
7. (*Applies only with Network Insight*) Click **Next** to initiate or disable discovery of the new network(s). Discovery settings differ based on whether you are defining one network or multiple networks.

- **Configuring one network:** Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member. By default, all Polling Options discovery settings are inherited from the parent network (or Grid, if no parent exists) unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. See the section [Configuring Grid Properties for Discovery](#) on page 529 for a complete description of discovery Polling Options.
 - or
 - **Configuring more than one network:** If the networks are child networks, they automatically inherit the settings of the parent network, including discovery settings and the discovery member. Discovery is disabled for any parent networks. These settings will not appear in the wizard page. For discovery of multiple networks, you can only enable or disable **Immediate Discovery**.
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
- If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#) on page 447. You can preview the information before the appliance submits updates to the RIPE database. To preview registration updates, click **Preview RIR Submissions**. For more information, see [Previewing Registration Updates](#) on page 451.
-
- Note:** You cannot leave an optional RIR attribute value empty. If you do not have a value for an RIR attribute, you must delete it from the table. You can enter up to 256 characters for all RIR attributes.
-
9. Save the configuration.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

After you create a network, you can do the following:

- Use the split network feature to create subnets for the network. For information, see [Splitting IPv4 Networks into Subnets](#) on page 472.
- Use the join networks feature to create a parent network that encompasses multiple subnets into a larger network. For information, see [Joining IPv4 Networks](#) on page 473. You can also create a shared network for subnets that are on the same network segment.

Networks served by Microsoft servers do not support the split and join functions.

Viewing Networks

You can view IPv4 networks from the **IPAM** tab -> Net Map and List panels. The Net Map panel provides a graphical view of your networks and the List panel displays the networks in table format. For more information, see [IPv4 Network Map](#) on page 467 and [Network List](#) on page 470.

You can also view a list of IPv4 and IPv6 networks in the **DHCP** tab -> **Networks** tab -> **Networks** panel. This panel displays all IPv4 and IPv6 networks.

In any of these panels, you can use filters or the **Go to** function to navigate to a specific network. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68. You can add, delete, or edit a network. You can also monitor the DHCP utilization of a selected network.

Depending on where you view your networks from, Grid Manager displays some of the following information by default. You can also select specific information for display.

- **Network:** The network address.
- **Comment:** The information you entered about the network.
- **RIR Organization:** This appears only if support for RIR updates is enabled. This displays the name of the RIR organization to which the network is assigned.

- **RIR Organization ID:** This appears only if support for RIR updates is enabled. This displays the ID of the RIR organization to which the network is assigned.
- **RIR Registration Status:** This appears only if support for RIR update is enabled. This field displays the RIR registration status. This can be **Registered** or **Not Registered**. **Registered** indicates that the network has a corresponding entry in the RIPE database.
- **Last Registration Updated:** Displays the timestamp when the last registration was updated. The displayed timestamp reflects the timestamp used on the Grid Master.
- **Status of Last Registration Update:** Displays the registration status and communication method of the last registration update. The status can be Pending, Sent, Succeeded, or Failed. Each time you send a registration update to create, modify, or delete a network container or network, the updated status will be displayed here. If you have selected not to send registration updates, the previous status is retained.
- **IPAM Utilization:** This information is available for IPv4 networks only. It displays the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network.
- **Site:** The site to which the network belongs. This is one of the predefined extensible attributes.
- **Protocol:** Displays whether the network is an IPv4 or IPv6 network.
- **Disabled:** Indicates if the network is disabled.
- **Leaf Network:** Indicates whether the network is a leaf network or not. A leaf network is a network that does not contain other networks.
- **Discovery Enabled:** (*Applies only with Network Insight*) Indicates whether discovery is allowed on the network container or the network.
- **Managed:** (*Applies only with Network Insight*) Indicates whether the network is set to Managed status under NIOS.
- **First Discovered:** (*Applies only with Network Insight*) The date and timestamp of the first occasion that NIOS discovered the network.
- **Last Discovered:** (*Applies only with Network Insight*) The date and timestamp of the last occasion that NIOS performed discovery on the network.
- **IPv4 DHCP Utilization:** This information is available for IPv4 networks only. It displays the percentage of the total DHCP usage of the IPv4 network. This is the percentage of the total number of DHCP hosts, fixed addresses, reservations, and active leases in the network divided by the total number of IP addresses (excluding IP addresses in the exclusion range) and all DHCP objects in the network. Note that only enabled addresses are included in the calculation. The appliance updates the utilization data approximately every 15 minutes. The utilization data is displayed in one of the following colors:
 - Red: The DHCP resources are 100% utilized.
 - Yellow: The DHCP utilization percentage is over the effective high-water mark threshold.
 - Blue: The DHCP utilization percentage is below the effective low-water mark threshold.
 - Black: The DHCP utilization percentage is at any number other than 100%, or it is not above and below any threshold.
- **Extensible attributes and RIR attributes (Building, Country, Region, State and VLAN):** You can select the extensible attributes such as Building, Country, Region, State, and VLAN for display. When you enable support for RIR updates, you can also select RIR attributes for display. For information about RIR attributes, see [Managing RIR Attributes](#) on page 444.

You can sort the list of networks in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#) on page 60.

You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62. You can edit values of inheritable extensible attributes by double clicking on the respective row. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you

perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329. If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.

Viewing Network Details

You can view detailed information about a specific network by clicking the network link. Grid Manager displays the objects in the network, including DHCP ranges, hosts, fixed addresses and roaming hosts. It displays the following information about the network:

- **IP Address:** The IP address of a DHCP object, such as a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address. For a DHCP range, this field displays the start and end addresses of the range. For a host that has multiple IP addresses, each IP address is displayed separately. Note that the appliance highlights all disabled DHCP objects in gray.
- **Type:** The DHCP object type, such as **DHCP Range** or **Fixed Address**.
- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Comment:** The information you entered for the object.
- **IPv4 DHCP Utilization:** The percentage of the total DHCP usage of a DHCP range. This is the percentage of the total number of fixed addresses, reservations, hosts, and active leases in the DHCP range divided by the total IP addresses in the range, excluding the number of addresses in the exclusion ranges. Note that only enabled objects are included in the calculation.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes.

You can select the following additional columns for display:

- **Static Addresses:** Indicates whether the IP address is a static address.
- **Dynamic Addresses:** Indicates whether the IP address is a dynamically assigned address.
- **Disabled:** Indicates whether the object is disabled.
- **Priority:** Displays the priority of a DHCP range when NAC filters are applied.
- Available extensible attributes.

You can also do the following in this panel:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only. You can edit values of inheritable extensible attributes by double clicking on the respective row. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329. If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.
- Sort the displayed data in ascending or descending order by column.
- Click **Go to IPAM View** to view information about the object in the **IPAM** tab.
- Add new objects, such as DHCP ranges, to the network.
- Delete or schedule the deletion of a selected object or multiple objects.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print or export the data.

Modifying IPv4 Networks

You can modify existing network settings and override the Grid or member DHCP properties, with the exception of the network address and netmask.

To modify an IPv4 network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* check box, and then click the Edit icon.
or
From the **Data Management** tab, select the **IPAM** tab -> *network* check box, and then click the Edit icon.
2. The *IPv4 Network* editor contains the following basic tabs from which you can modify data:
 - **General Basic:** You can modify the following fields:
 - **Comment:** The information you entered for the network.
 - **Disabled:** This field is displayed only if the selected network is a network without a child network under it. You can disable and enable existing networks instead of removing them from the database, if the selected network does not have a child subnet. This feature is especially helpful when you have to move or repair the server for a particular network.
 - **Member Assignment:** Add or delete a Grid member that provides DHCP services for this network. For information, see [Adding IPv4 Networks](#) on page 845.
 - **IPv4 DHCP Options:** Keep the inherited DHCP properties or override them and enter unique settings for the network. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **RIR Registration:** Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#) on page 443.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322. You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329. If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Optionally, click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **General Advanced:** You can associate zones with a network. For information, see [Associating Networks with Zones](#) on page 813.
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the network. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.
 - **IPv4 Thresholds:** Keep the inherited thresholds settings or override them and enter unique settings for the network. For information, see [Configuring Thresholds for DHCP Ranges](#) on page 807.
4. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting IPv4 Networks

When you delete a network, all of its data, including all DHCP records, subnets, and records in its subnets, is deleted from the database. Because of the potentially large loss of data that can occur when you delete a network, the appliance stores the deleted network in the Recycle Bin. You can restore a deleted network from the Recycle Bin, if enabled. You can also disable a network instead of deleting it. For information, see [Modifying IPv4 Networks](#) on page 851.

To delete a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* check box, and then select **Delete** or **Schedule Delete** from the Delete drop-down menu.
2. To delete the network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#) on page 76.

The appliance puts the deleted network in the Recycle Bin, if enabled.

CONFIGURING IPV4 SHARED NETWORKS

You can combine individual contiguous networks into a shared network to allow the DHCP server to assign IP addresses from any subnet (that resides on the same network interface) in the shared network.

Before creating a shared network, you must first create the subnets. For example, you must first create the networks 10.32.1.0 and 10.30.0.0 before designating them to a shared network. For more information, see [About Shared Networks](#) on page 779.

Adding IPv4 Shared Networks

To add a shared network:

1. Select the **Data Management** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network.
3. Select the **DHCP** tab -> **Networks** tab.
4. In the **Shared Networks** section, select **IPv4 Shared Network** from the Add drop-down menu.
5. In the *Add IPv4 Shared Network* wizard, complete the following and click **Next**:
 - **Name:** Enter the name of the shared network.
 - **Comment:** Enter information about the shared network.
 - **Disabled:** Select this if you want to enable the shared network at a later time. You can disable and enable existing networks instead of removing them from the database. This feature is especially helpful when you have to move or repair the server for a particular network.
6. Do the following to add networks:
 - a. Click the Add icon.
 - b. In the *Select Network* dialog box, select the networks that you want to include in the shared network. Ensure that the networks are served by the same Grid members to avoid DHCP inconsistencies.
7. Click **Next** to configure or override DHCP options as described in [Defining Basic IPv4 Options](#) on page 801.
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes for the shared network. For information, see [Using Extensible Attributes](#) on page 332.
9. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Viewing Shared Networks

To view IPv4 and IPv6 shared networks:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks**.
2. Grid Manager displays the following information:
 - **Name:** The name of the shared network.
 - **Protocol:** Displays whether the network is an IPv4 or IPv6 network.
 - **Comment:** The information you entered about the shared network.
 - **IPv4 DHCP Utilization:** The percentage of the DHCP utilization of the networks that belong to the shared network. This is the percentage of the total number of available IP addresses from all the networks that belong to the shared network versus the total number of all IP addresses in all of the networks in the shared network.
 - **Site:** The site to which the shared network belongs. This is one of the predefined extensible attributes.

You can select **Disabled** or available extensible attributes for display. You also can view detailed information about a network in a shared network by clicking the network link.

In this panel, you can use filters or the **Go to** function to navigate to a specific network. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

You can sort the list of networks in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#) on page 60.

You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.

Modifying IPv4 Shared Networks

You can modify existing network settings and override the Grid or member DHCP properties:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared_network* check box, and then click the Edit icon.
2. The *Shared Network* editor contains the following tabs from which you can modify data:
 - **General:** Modify the fields **Name**, **Comments**, and **Disabled** as described in [Adding IPv4 Shared Networks](#) on page 852.
 - **Networks:** Displays the networks that are currently assigned to the shared network. You can add or delete a network. To add a network, click the Add icon. In the *Select Network* dialog box, select the network you want to add. To delete an existing network, select the *network* check box, and then click the Delete icon.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **IPv4 DHCP Options:** Keep the inherited DHCP properties or override them and enter unique settings for the shared network. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the shared network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the shared network. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting IPv4 Shared Networks

Though you can delete the networks in a shared network, a shared network must have at least one network in it.

To delete a shared network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared_network* check box, and then select **Delete** or **Schedule Delete** from the drop-down menu.
2. To delete the shared network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#) on page 76.

The appliance puts the deleted shared network in the Recycle Bin, if enabled.

CONFIGURING IPV4 ADDRESS RANGES

In a network, you define address ranges from which the DHCP server or failover association assigns IP addresses to client requests. When a DHCP client requests an IP address, the appliance allocates an address within a defined DHCP range. The DHCP client can use the assigned IP address until the lease expires.

When you do not assign a DHCP server or failover association to an address range, the range becomes a reserved range. A reserved range contains IP addresses that are reserved for static hosts, not for dynamic assignments. You can allocate the next available IP from a reserved range.

You can also apply filters to DHCP ranges to control how the DHCP server allocates IP addresses. For information about DHCP filters, see [About IPv4 DHCP Filters](#) on page 892.

Adding IPv4 Address Ranges

To add an IPv4 address range:

1. Navigate to the IPv4 network to which you want to add an address range, and then select **Range** from the Add drop down menu.

or

From any panel in the DHCP tab, expand the Toolbar and click **Add** -> **Range** -> **IPv4**.

2. In the *Add IPv4 Range* wizard, select one of the following and click **Next**:
 - **Add Range**: Select this to add an address range from scratch.
- or
- **Add Range Using Template**: Click **Select Template** and select the template that you want to use. Note that when you use a template to create an address range, the configurations of the template apply to the new range. The appliance automatically populates the range properties in the wizard. You can then edit the pre-populated properties.
3. Complete the following:
 - **Network**: Click **Select Network**. Grid Manager displays the network address here if you have only one network configured. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.
 - **Start**: Enter the first available IP address in the range.
 - **End**: Enter the last available IP address in the range.

- **Name:** Optionally, enter a name for the range.
 - **Comment:** Enter additional information about the address range.
 - **Disabled:** Select this if you want to save the configuration for the address range but do not want to activate the address range yet. You can clear this check box when you are ready to allocate addresses from this range.
4. Click **Next** and select one of the following:
 - **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
 - **Grid Member:** Select this if you want a Grid member to serve DHCP for this address range. Select a Grid member from the drop-down list. The drop-down list displays only the Grid members that are associated with the network to which the DHCP range belongs.
 - **Failover Association:** Select this if you want a failover association to serve DHCP for this address range. Click **Select Association**. In the *DHCP Failover Association Selector* dialog box, choose a failover association, and then click the Select icon. The appliance lists failover associations that serve DHCP in the network view of the DHCP range. For information, see [Chapter 28, DHCP Failover](#), on page 883.
 5. Click **Next** to configure or override DHCP options as described in [Defining Basic IPv4 Options](#) on page 801.

Note: Steps 6-7 apply only in deployments using Network Insight discovery features. Otherwise, skip to Step 8.

6. Click **Next** to initiate or disable discovery of the new DHCP range.
 7. Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member. By default, all Polling Options discovery settings are inherited from the parent network unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. See the section [Configuring Grid Properties for Discovery](#) on page 529 for a complete description of discovery Polling Options.
 8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 9. Save the configuration and click **Restart** if it appears at the top of the screen.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

For information on viewing address ranges in a network, see [Viewing IPv4 DHCP Objects](#) on page 862

Modifying IPv4 Address Ranges

You can modify settings for the DHCP range. You can also define an exclusion range to prevent the appliance from assigning the addresses in the exclusion range to clients. IP addresses in an exclusion range are excluded from the pool of IP addresses. For more information, see [About Exclusion Ranges](#) on page 780.

To modify an IPv4 address range:

1. From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** section → *network* → *addr_range* check box, and then click the Edit icon.
2. The *DHCP Range* editor contains the following basic tabs from which you can modify data:
 - **General:** Modify the fields, except the network address, as described in [Adding IPv4 Address Ranges](#) on page 854.
 - **Member Assignment:** Modify the Grid member or failover association that provides DHCP services for the DHCP range as described in [Adding IPv4 Address Ranges](#) on page 854.
 - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the DHCP range. For information, see [Defining Basic IPv4 Options](#) on page 801.

- **Extensible Attributes:** You can add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332. You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329. If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the DHCP range. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the DHCP range. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.
 - **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use to assign to clients. You can use these exclusion addresses as static IP addresses. Enter the start and end addresses of the exclusion range, and optionally, enter information about this exclusion range.
 - **IPv4 Thresholds:** Keep the inherited thresholds settings or override them and enter unique settings for the DHCP range. For information, see [Configuring Thresholds for DHCP Ranges](#) on page 807.
 - **Filters:** You can add or delete DHCP filters to the range. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
 4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Controlling Lease Assignments

You can set parameters to control how the DHCP server responds to lease requests within a specific DHCP range. When you set a DHCP range to deny all leases requests, the appliance does not assign IP addresses within this range to DHCP clients. This is useful when you want DHCP clients with IP addresses within this range to obtain new IP addresses when they renew their leases. When a client with an IP address within this range broadcasts a DHCPREQUEST message for its old IP address, the authoritative DHCP server responds with a DHCPNAK. This causes the client to move to the INIT state and to send a DHCPDISCOVER message for a new IP address.

You can also configure the DHCP server to assign or deny IP addresses within a DHCP range to known and unknown DHCP clients. Known clients include roaming hosts and clients with fixed addresses or DHCP host entries. Unknown clients include clients that are not roaming hosts and clients that do not have fixed addresses or DHCP host entries.

To control how the appliance assigns leases to client requests:

1. **DHCP Range:** From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* → *addr_range* check box, and then click the Edit icon.
2. In the *IPv4 Range* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **IPv4 DHCP Options** tab → **Advanced** tab and complete the following:

Note: The **IPv4 DHCP Options** tab is enabled when you select a **Grid Member** or **IPv4 DHCP Failover Association** in the **Member Assignment** tab.

- **Allow/Deny Clients**
 - **Known Clients:** Select this check box, and then select **Allow** or **Deny** from the drop-down list to assign or deny IP addresses within this range to known DHCP clients. Known DHCP clients include roaming hosts and clients with fixed addresses or DHCP host entries. Note that the appliance cannot deny an IP address to a fixed address within this range. You must disable the fixed address if you do not want it to obtain an IP address here.
 - **Unknown Clients:** Select this check box, and then select **Allow** or **Deny** from the drop-down list to assign or deny IP addresses within this range to unknown DHCP clients. Unknown DHCP clients include clients that are not roaming hosts and clients that do not have fixed addresses or DHCP host entries.
 - **Deny Leases:** Select **Deny all lease requests for this range** to deny all lease requests from DHCP clients.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Deleting IPv4 Address Ranges

To delete a DHCP range:

- From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *addr_range* check box, and then click the Delete icon.

CONFIGURING IPV4 FIXED ADDRESSES

A fixed address represents a persistent link between an IP address and one of the following:

- MAC address
- Client identifier
- Circuit ID or remote ID in the DHCP relay agent option (option 82)

You can create fixed addresses as described in [Adding IPv4 Fixed Addresses](#) or from the Tasks Dashboard. For information about the Tasks Dashboard, see [The Tasks Dashboard](#) on page 99. You can also create a fixed address when you create a host record or when you convert an active, dynamically leased address to a fixed address. For more information, see [Adding Host Records](#) on page 462 and [Converting DHCP Leases](#) on page 488.

When you create a fixed address, you must define a host identifier that the DHCP server uses to match the DHCP client. Every time the DHCP client with the matching identifier requests an IP address, the DHCP server assigns it the same address.

When a DHCP client sends a DHCPDISCOVER, it can include the MAC address or a unique client identifier as option 61 in the DHCP section of the packet. Using a client identifier is especially useful for virtualized server processes that might be moved to different hardware platforms. For information about option 61, refer to *RFC2132, DHCP Options and BOOTP Vendor Extensions*. You can select either the MAC address or client identifier as the host identifier in a fixed address. The DHCP server matches the option 61 value in the client request using either the MAC address or client identifier, depending on your configuration. When a DHCP client renews an IP address using a matching MAC address or client identifier, the DHCP server tracks the allocation of IP addresses and reserves the same IP address for the client.

When you enter a MAC address, you can use one of the following formats:

- aa:bb:cc:dd:ee:ff — Six groups of two hexadecimal digits separated by colons (:))
- aa-bb-cc-dd-ee-ff — Six groups of two hexadecimal digits separated by hyphens (-)
- aabb.ccdd.eeff — Three groups of four hexadecimal digits separated by periods (.)
- aabbcc-ddeeff — Two groups of six hexadecimal digits separated by a hyphen (-)
- aabbccddeeff — One group of 12 hexadecimal digits without any separator

After you save the entry, the appliance displays the MAC address in the AA:BB:CC:DD:EE:FF format.

When a DHCP client requests an IP address through a DHCP relay agent, the agent adds either the circuit ID or remote ID, or both, to the DHCP relay agent information option (option 82). For information, see [About the DHCP Relay Agent Option \(Option 82\)](#) on page 806. When you select the DHCP relay agent option (circuit ID or remote ID) as the host identifier in a fixed address, the DHCP server matches the DHCP client request using either the circuit ID or the remote ID, depending on your configuration. When a DHCP client renews an IP address using a matching relay agent ID, the DHCP server tracks the allocation of IP addresses and reserves the same IP address for the client. Note that leases are not renewed at the standard renewal time (T1) when option 82 information is not available as a unicast renewal. Instead, leases are renewed at the rebinding time (T2) when renewals are sent as broadcasts to the relay agents and contain option 82 information. For information about how to configure the lease time, see [Configuring General IPv4 DHCP Properties](#) on page 793.

Adding IPv4 Fixed Addresses

To add an IPv4 fixed address:

1. Navigate to the network to which you want to add a fixed address, and then select **Fixed Address** from the Add drop-down menu.

or

From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> Fixed Address -> IPv4**.

2. In the *Add IPv4 Fixed Address* wizard, select one of the following and click **Next**:

- **Add Fixed Address**

or

- **Add Fixed Address using Template**

Click **Select Template** and select the template that you want to use. Note that when you use a template to create a fixed address, the configurations of the template apply to the new address. The appliance automatically populates the fixed address properties in the wizard. You can then edit the pre-populated properties.

3. Complete the following:

- **Network:** Click **Select Network**. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.
- **IP Address:** Enter the IPv4 address for the fixed address, or click **Next Available IP** to obtain the next available IP address. For information about the next available IP address, see [About the Next Available Network or IP Address](#) on page 844.

Note: When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

- If the network of the IP address is served by a Grid member, Grid Manager displays the **Assign IP Address by** section. Select one of the following to match your criteria:
 - **MAC Address:** Select this to assign a fixed address to a host with the MAC address that you specify here. Enter the MAC address in the field. For MAC address format, see [Configuring IPv4 Fixed Addresses](#) on page 857.
 - **DHCP Client Identifier:** Select this to assign a fixed address to a host with the DHCP client identifier that you specify here. In the field, enter the client identifier of the host to which you want the DHCP server to assign this IP address. The client identifier must be unique within the network.
 - **Match null (\0) at beginning of DHCP client identifier:** This is enabled when you select **DHCP client identifier**. Select this when a DHCP client sends a \000 prefixed to the DHCP client identifier. \0 is the null character. Some DHCP clients (for example, Microsoft) send the client identifier in a \000foo format (with the null character prefix instead of just foo). The client identifier for the requesting host and the client identifier stored in the appliance must match.

- **DHCP Relay Agent:** Select this to assign a fixed address to a host with the circuit ID or remote ID you specify here. From the drop-down list, select **Circuit ID** or **Remote ID**, and then enter the ID in the field. For information about circuit IDs and remote IDs, see [About the DHCP Relay Agent Option \(Option 82\)](#) on page 806. You can enter the ID in hexadecimal format, such as ex:aa, ab, 1f:cd, or ef:23:56, or in string format, such as abcd or aa:gg. The appliance matches the value you enter here with the value sent by the DHCP client in counted octet sequence format. For information about how to use hexadecimal values, see [DHCP Option Data Types](#) on page 800. The ID is case sensitive and can contain up to 230 characters.

Note: You cannot use the same circuit ID or remote ID for different fixed addresses if the addresses are in the same network or the same shared network.

- **Name:** Enter a name for the fixed address. This field is required if the network is served by a Microsoft server. For information, see [Adding Fixed Addresses/Microsoft Reservations](#) on page 993.
 - **Comment:** Optionally, enter additional information about the fixed address.
 - **Disabled:** Select this if you do not want the DHCP server to allocate this IP address at this time.
4. Click **Next** to configure or override DHCP options as described in [About IPv4 DHCP Options](#) on page 800.
 5. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new fixed IP address.
 - Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the fixed IP address. If you choose **Enable Immediate Discovery**, discovery will execute on the host after you save your settings. You may also choose to leave both options disabled.
 - By default, the new host inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** checkbox and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials. For descriptions of SNMP credentials for discovery, see the section [Configuring Grid SNMPv1/v2 Properties](#) on page 530 and [Configuring Grid SNMPv3 Properties](#) on page 530.
 6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 7. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

For information on viewing fixed addresses and other DHCP objects, see [Viewing IPv4 DHCP Objects](#) on page 862.

Modifying IPv4 Fixed Addresses

To modify the settings of a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed_address* check box, and then click the Edit icon.
2. The *Fixed Address* editor contains the following basic tabs from which you can modify data:
 - **General:** You can modify the fields, except the network address, as described in [Adding IPv4 Fixed Addresses](#) on page 858.
 - **IPv4 DHCP Options:** You can keep the inherited DHCP options or override them and enter unique settings for the fixed address. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **Discovered Data:** Displays the discovered data of the fixed address. For information, see [Viewing Discovered Data](#) on page 510.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332. You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing.

- **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **IPv4 DDNS:** You can keep the inherited DDNS settings or override them and enter unique settings for the fixed address. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** You can keep the inherited BOOTP properties or override them and enter unique settings for the fixed address. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
 4. Save the configuration and click **Restart** if it appears at the top of the screen.
 - or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting Fixed Addresses

To delete a fixed address within the DHCP range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed_address* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, do the following:
 - **Delete associated leases with the fixed address (selected fixed IP address):** When you clear this check box and click **Yes**, the appliance changes the status of the associated leases from **Static** to **Active**. When you select this check box and click **Yes**, the appliance deletes all the leases associated with the fixed address.

Note: NIOS removes all the static leases associated with a fixed address when you delete a fixed address out of the DHCP range, regardless of the selection of the **Delete associated leases with the fixed address (selected fixed IP address)** check box in the *Delete Confirmation* dialog box.

CONFIGURING IPv4 RESERVATIONS

You can create a reservation as a static IP address for future use. A reservation is a pre-provisioned fixed address that is associated with a MAC address of 00:00:00:00:00:00. Since 00:00:00:00:00:00 is not a real MAC address, no client can receive this IP address from the address pool. You can reserve this static IP address and assign it to a client in the future.

To create a reservation, you can do one of the following:

- Add a reservation. For information, see [Adding IPv4 Reservations](#).
- Convert a fixed address or a dynamic address with an active lease to a reservation. For information, see [Converting Objects Associated with IP Addresses](#) on page 487.
- Define a fixed address with an IP address. For information, see [Adding IPv4 Fixed Addresses](#) on page 858.

Adding IPv4 Reservations

To create a reservation:

1. Navigate to the network to which you want to add a reservation, and then select **Reservation** from the Add drop down menu.
or
From any panel in the DHCP tab, expand the Toolbar and click **Add -> IPv4 Reservation**.
2. In the *Add Reservation* wizard, select one of the following and click **Next**:
 - **Add Reservation**
or
 - **Add Reservation using Template**
Click **Select Template** and select the template that you want to use. Note that when you use a template to create a reservation, the configurations of the template apply to the new address. The appliance automatically populates the reservation properties in the wizard. You can then edit the pre-populated properties.
3. Complete the following:
 - **Network**: The displayed network address can either be the last selected network or the network from which you are adding the DHCP range. If no network address is displayed or if you want to specify a different network, click **Select Network**. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.
 - **IP Address**: Enter the IP address that you want to reserve for manual assignment, or click **Next Available IP** to obtain the next available IP address. For information about obtaining the next available IP address, see [Adding IPv4 Fixed Addresses](#) on page 858.

Note: When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

 - **Name**: Optionally, enter a name for the reservation.
 - **Comment**: Optionally, enter additional information about the reservation.
 - **Disabled**: Select this if you do not want the DHCP server to use this reservation at this time.
4. Click **Next** to configure or override DHCP options as described in [Defining Basic IPv4 Options](#) on page 801.
5. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new IPv4 reservation.
 - Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the IP reservation. If you choose **Enable Immediate Discovery**, discovery will execute on the host after you save your settings. You may also choose to leave both options disabled.
 - By default, the new reservation inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** check box and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials. See the section [Configuring Grid SNMPv1/v2 Properties](#) on page 530 and [Configuring Grid SNMPv3 Properties](#) on page 530 for a complete description of SNMP credentials for discovery.
6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
7. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Modifying Reservations

To modify a reservation:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *reservation* check box, and then click the Edit icon.
2. The *Reservation Address* editor contains the following tabs from which you can modify data:
 - **General:** Modify the fields, except the network address, as described in [Adding IPv4 Reservations](#) on page 861.
 - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the reservation. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **Discovered Data:** Displays the discovered data of the reservation. For information, see [Viewing Discovered Data](#) on page 510.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a reservation. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332. You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) on page 329. If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the reservation. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.
 - **IPv4 BOOTP/PXE:** You can keep the inherited BOOTP properties or override them and enter unique settings for the reservation. For information, see [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

VIEWING IPv4 DHCP OBJECTS

To view the address ranges, fixed addresses and reservations in a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range*.
2. Grid Manager displays the following information:
 - **IP Address:** The IP address of the object in the DHCP range. For exclusion ranges, this displays the start and end IP addresses. For host records with multiple IP addresses, each IP address is displayed separately. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
 - **Type:** The object type, such as **Fixed Address**.

- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.
- **Comment:** The information you entered for the object.
- **Site:** The site to which the object belongs. This is one of the predefined extensible attributes. You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing.

You can select **Disabled** or available extensible attributes for display.

You can also do the following:

- Sort the data in ascending or descending order by column.
- Create a bookmark for the range.
- Delete or schedule the deletion of a selected object or multiple objects in the range.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Select an object and view detailed information.
- Print or export the data.

ABOUT ROAMING HOSTS

A roaming host is a host with a dynamically assigned IP address and a specific set of properties and DHCP options. When you create a roaming host for a network device, the device can receive any dynamically assigned address from the network to which it belongs. You can create roaming hosts for devices, such as laptop computers and mobile phones, that require different IP addresses each time they are moved from one network to another and require a unique set of DHCP options.

You can configure IPv4 addresses, IPv6 addresses, or IPv4 and IPv6 addresses for roaming hosts that require both types of addresses. When you configure IPv4 addresses for a roaming host, you must specify the host MAC address or a DHCP client identifier that the appliance uses to match the host, and specify DHCP options for the host. The appliance assigns an IP address from the DHCP range associated with the network from which the address request originates. You can configure an IPv6 prefix or address for a DHCP client. When you do, you must specify the DUID of the host so the appliance can use the DUID to match the host.

A roaming host also receives DHCP options from the Grid, member, network, or shared network with which it associates.

When you configure a roaming host, you must configure it in a specific network view. If you have multiple network views, you must specify the network view to which the requesting hosts belong so the appliance can assign addresses to the hosts from the networks within the same network view.

After you enable support for roaming hosts at the Grid level, you can add a roaming host that supports IPv4, IPv6, or both protocols. You can also convert an IPv4 roaming host to an IPv6 roaming host and vice versa, or convert an IPv4 or IPv6 roaming host to one that supports both IPv4 and IPv6.

Configuring Roaming Hosts

To configure a roaming host, perform the following tasks:

1. Enable support for roaming hosts at the Grid level. For information, see, [Enabling Support for Roaming Hosts](#).
2. Add a roaming host.
 - To add an IPv4 roaming host, see [Adding IPv4 Roaming Hosts](#)

- To add an IPv6 roaming host, see [Adding IPv6 Roaming Hosts](#)
- To add a dual stack roaming host, see [Adding IPv4/IPv6 Roaming Hosts](#)
- Optionally, configure DHCP properties for the roaming host. You can override properties set for the upper levels and enter unique values for the roaming hosts. For information, see [Defining Basic IPv4 Options](#) on page 801

You can do the following after you configure roaming hosts:

- View the configured roaming hosts. For information, see [Viewing Roaming Hosts](#)
- Modify existing roaming hosts. For information, see [Setting Properties for Roaming Hosts](#) on page 867.
- Delete roaming hosts that are not currently in use. For information, see [Deleting Roaming Hosts](#) on page 868.

Enabling Support for Roaming Hosts

You must first enable support for roaming hosts before adding them. After you enable this feature, you can disable it only after you delete all the existing roaming hosts.

To enable support for roaming hosts:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Expand the Toolbar and click **Grid DHCP Configuration**.
3. In the *General Advanced* tab, select **Enable support for roaming host**.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Adding IPv4 Roaming Hosts

To add an IPv4 roaming host:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Select a network view from the drop-down list.
3. Expand the Toolbar and click **Add -> Roaming Host -> IPv4**.
4. In the *Add Roaming Host* wizard, select one of the following and click **Next**:
 - **Add Roaming Host**
 - or
 - **Add Roaming Host using Template**

Click **Select Template** to create a roaming host using a fixed address/reservation template. In the *DHCP Template Selector* dialog box, select the template that you want to use. Note that when you use a template to create a roaming host, the configurations of the template apply to the new host. The appliance automatically populates the host properties in the wizard. You can then edit the pre-populated properties.
5. Complete the following:
 - **Name:** Enter the name of the roaming host. The name must be unique for each roaming host in a given network view.
 - **Assign IPv4 Address by:** Select one of the following criteria on which the appliance matches when assigning an IP address to the host.
 - **MAC Address:** Select this to assign a dynamic IP address to a host, provided that the MAC address of the requesting host matches the MAC address that you specify here.
 - **DHCP Client Identifier:** Select this to assign a dynamic IP address to a host with the same DHCP client identifier that you specify here. When you select this, the **Match null (\0) at beginning of DHCP client identifier** check box is displayed. Select this when a DHCP client sends a \000 prefixed to the DHCP client identifier. \0 is the null character. Some DHCP clients (for example, Microsoft) send the client identifier in a \000foo format (with the null character prefix instead of just foo). The client identifier for the requesting host and the client identifier stored in the appliance must match.
 - **Comment:** Enter useful information about the roaming host.

- **Disabled:** Select this if you do not want the DHCP server to use this roaming host definition. When you disable a roaming host, the host gets an IP address without the defined DHCP options.
- 6. Click **Next** to configure the IDHCP options for the roaming host, as described in [Defining Basic IPv4 Options](#) on page 801.
- 7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
- 8. Save the configuration and click **Restart** if it appears at the top of the screen.
- or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Adding IPv6 Roaming Hosts

To add an IPv6 roaming host:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Select a network view from the drop-down list.
3. Expand the Toolbar and click **Add -> Roaming Host -> IPv6**.
4. In the *Add Roaming Host* wizard, select one of the following and click **Next**:
 - **Add IPv6 Roaming Host**
 - or
 - **Add Roaming Host Using IPv6 Template**

Click **Select IPv6 Template** to create a roaming host using an IPv6 fixed address template. In the *DHCP Template Selector* dialog box, select the template that you want to use. Note that when you use a template to create a roaming host, the configurations of the template apply to the new host. The appliance automatically populates the host properties in the wizard. You can then edit the pre-populated properties.
5. Complete the following:
 - **Name:** Enter the name of the roaming host. The name must be unique for each roaming host in a given network view.
 - **DUID:** Enter the DHCP unique identifier of the host.
 - **Comment:** Optionally, enter additional information about the roaming host.
 - **Disabled:** Select this if you do not want the DHCP server to use this roaming host definition. When you disable a roaming host, the host gets an IP address without the defined DHCP options.
6. Click **Next** to configure the DHCP options for the roaming host, as described in [Defining General IPv6 Properties](#) on page 809.
7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
8. Save the configuration and click **Restart** if it appears at the top of the screen.
- or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Adding IPv4/IPv6 Roaming Hosts

To add an IPv4/IPv6 roaming host:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Select a network view from the drop-down list.
3. Expand the Toolbar and click **Add -> Roaming Host -> Both**.

4. In the *Add Roaming Host* wizard, select one of the following and click **Next**:

- **Add Roaming Host**

or

- **Add Roaming Host using Both IPv4 and IPv6 Templates**

When you use both templates to create a roaming host, the appliance applies the IPv4 template and then the IPv6 template. Therefore, the comments and extensible attributes from the IPv6 template override those from the IPv4 template.

5. Complete the following:

- **Name:** Enter the name of the roaming host. The name must be unique for each roaming host in a given network view.
- **Assign IP Address by:** Select one of the following criteria on which the appliance matches when assigning an IP address to the host.
 - **MAC Address:** Select this to assign a dynamic IP address to a host, provided that the MAC address of the requesting host matches the MAC address that you specify here.
 - **DHCP Client Identifier:** Select this to assign a dynamic IP address to a host with the same DHCP client identifier that you specify here. When you select this, the **Match null (\0) at beginning of DHCP client identifier** check box is displayed. Select this when a DHCP client sends a \000 prefixed to the DHCP client identifier. \0 is the null character. Some DHCP clients (for example, Microsoft) send the client identifier in a \000foo format (with the null character prefix instead of just foo). The client identifier for the requesting host and the client identifier stored in the appliance must match.
- **DUID:** Specify the DHCP unique identifier of the host.
- **Comment:** If both IPv4 and IPv6 templates were used to create the host, this field displays the comment from the IPv6 template. You can change or add information.
- **Disabled:** Select this if you do not want the DHCP server to use this roaming host definition. When you disable a roaming host, the host gets an IP address without the defined DHCP options.

6. Click **Next** to configure the IPv4 DHCP options for the roaming host, as described in [Defining Basic IPv4 Options](#) on page 801.
 7. Click **Next** to configure IPv6 properties described in [Defining General IPv6 Properties](#) on page 809.
 8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. If both IPv4 and IPv6 templates were used to create the host, this panel displays the attributes from the IPv6 template. You can change or add information. For information, see [Using Extensible Attributes](#) on page 332.
 9. Save the configuration and click **Restart** if it appears at the top of the screen.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Viewing Roaming Hosts

To view a list of roaming hosts in a specific network view:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts**.
2. From the Network View drop-down list, select the network view to which the roaming hosts belong.
3. The Grid Manager displays the following for each roaming host:
 - **Name:** The name of the roaming host.
 - **Address:** The IP address of the roaming host.
 - **Comment:** The information that you entered for the roaming host.
 - **Site:** The site to which the template belongs. This is one of the predefined extensible attributes.

You can select **Disabled** and available extensible attributes for display.

You can also do the following:

- Sort the displayed data in ascending or descending order by column.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.

Setting Properties for Roaming Hosts

You can modify an existing roaming host to add, modify or delete IPv4 or IPv6 addresses, and to set IPv4 and IPv6 DHCP properties.

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** section -> *roaming_host* check box, and then click the Edit icon.
 2. The *Roaming Host* editor contains the following tabs from which you can modify data:
 - **General:** Edit the fields as described in [Adding IPv4 Roaming Hosts](#) on page 864, except for the **Templates** field.
 - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the roaming host. For information, see [Defining Basic IPv4 Options](#) on page 801.
 - **IPv6 DHCP Options:** Keep the inherited IPv6 DHCP properties or override them. For more information, see [Defining General IPv6 Properties](#) on page 809.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a roaming host. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
 3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **IPv4 DDNS:** Click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. You can specify the following:
 - **DDNS Domain Name:** Specify the domain name that the appliance uses to update DNS.
 - **DDNS Hostname:** Select the **Replace the host name dynamically provided by the client/member with the roaming host name** check box to use the name of the roaming host record as the name of the client for DDNS updates.

For information about DDNS, see [Chapter 20, Configuring DDNS Updates from DHCP](#), on page 689.
 - **IPv4 BOOTP/PXE:** Keep the inherited PXE and BOOTP properties or override them and enter unique settings for the roaming host. For information, see [Configuring DHCP for IPv4](#) on page 843.
 - **IPv6 DDNS:** Click **Override** and select **Enable DDNS Updates** for the DDNS settings you configure in this tab to take effect. You can specify the following:
 - **DDNS Domain Name:** Specify the domain name that the appliance uses to update DNS.
 - **DDNS Hostname:** Select the **Replace the host name dynamically provided by the client/member with the roaming host name** check box to use the name of the roaming host record as the name of the client for DDNS updates.

For information about DDNS, see [Chapter 20, Configuring DDNS Updates from DHCP](#), on page 689.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
 4. Save the configuration and click **Restart** if it appears at the top of the screen.
- You can also click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting Roaming Hosts

To delete a roaming host:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming_host* check box, and then select **Delete** or **Schedule Delete** from the drop-down menu.
2. To delete the roaming host now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#) on page 76. The Grid Manager puts the deleted roaming host in the Recycle Bin, if enabled.



Chapter 27 Managing IPv6 DHCP Data

This chapter explains how to configure and manage IPv6 DHCP data. It contains the following sections:

- [*Configuring IPv6 Networks*](#) on page 870
- [*Defining Global IPv6 Prefixes*](#) on page 870
- [*Managing IPv6 Networks*](#) on page 871
 - [*Adding IPv6 Networks*](#) on page 871
 - [*Modifying IPv6 Networks*](#) on page 874
 - [*Deleting IPv6 Networks*](#) on page 874
- [*About IPv6 Shared Networks*](#) on page 875
 - [*Adding IPv6 Shared Networks*](#) on page 875
 - [*Modifying IPv6 Shared Networks*](#) on page 875
 - [*Deleting IPv6 Shared Networks*](#) on page 876
- [*Configuring IPv6 Address Ranges*](#) on page 876
 - [*Adding IPv6 Address Ranges*](#) on page 876
 - [*Setting the Priority of IPv6 Address Ranges*](#) on page 877
 - [*Modifying IPv6 Address Ranges*](#) on page 878
 - [*Deleting IPv6 Address Ranges*](#) on page 878
- [*Configuring IPv6 Fixed Addresses*](#) on page 878
 - [*Adding IPv6 Fixed Addresses*](#) on page 878
 - [*Modifying IPv6 Fixed Addresses*](#) on page 880
 - [*Deleting IPv6 Fixed Addresses*](#) on page 880
- [*Viewing IPv6 DHCP Objects*](#) on page 880

CONFIGURING IPV6 NETWORKS

To configure DHCP services for an IPv6 network and the resources in the network, perform the following tasks:

1. To facilitate network creation, you can specify the IPv6 global prefixes for the Grid. For more information, see [Defining Global IPv6 Prefixes](#).
2. Create a network and assign it to Grid members. For information, see [Managing IPv6 Networks](#) on page 871 and [About IPv6 Shared Networks](#) on page 875.
3. Optionally, configure DHCP properties for the network. You can override properties set at the Grid or member level and enter unique values for the network and fixed addresses. For information, see [Configuring DHCPv6 Properties](#) on page 809 and [Configuring DHCP IPv4 and IPv6 Common Properties](#) on page 812.
4. Optionally, assign zones to a network. For information, see [Associating Networks with Zones](#) on page 813.
5. Add a DHCP range to the network and assign it to a member. For information, see [Configuring IPv6 Address Ranges](#) on page 876.
6. Optionally, add exclusions to the DHCP range for addresses that are not used for dynamic allocation. For information, see [Modifying IPv6 Address Ranges](#) on page 878.
7. Optionally, configure DHCP properties for the address range. You can override properties set at an upper level and enter unique values for the address range. For information, see [Modifying IPv6 Address Ranges](#) on page 878.
8. Optionally, add fixed addresses to the network and configure DHCP properties for them. For information, see [Configuring IPv6 Fixed Addresses](#) on page 878.
9. Start the DHCP service and the IPv6 DHCP service. For more information, see [Starting DHCP Services on a Member](#) on page 822.

DEFINING GLOBAL IPV6 PREFIXES

To simplify network creation, you can define IPv6 prefixes that are used for networks served by the Grid members. If your organization is assigned IPv6 prefixes, you can enter them globally at the Grid level, and then just select the appropriate IPv6 prefix when you define the networks.

To add global IPv6 prefixes:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Expand the Toolbar and click **Grid DHCP Properties**.
3. In the *Grid DHCP Properties* editor, select the **IPv6 Global Prefixes** tab.
4. Click the Add icon and enter a name for the prefix and the prefix. Select the **Default** check box if you'd like to specify a default IPv6 prefix for the Grid.
5. Save the configuration.

MANAGING IPV6 NETWORKS

You can create an IPv6 network from scratch or create a network template and then use that template to create one or more networks. Using a network template facilitates creating multiple IPv6 networks with similar properties. You can also create an IPv6 network from the Tasks Dashboard. For information about the Tasks Dashboard, see [The Tasks Dashboard](#) on page 99.

An IPv6 network inherits its DHCP options & DDNS settings from its shared network, if it is part of a shared network, or from the member to which it is assigned.

If you have enabled support for RIR (Regional Internet Registry) updates and are adding an RIR IPv6 network container or network to NIOS, Grid Manager displays an RIR section in the *Add IPv6 Network* wizard. You must enter RIR related information in this section in order for NIOS to associate the newly added network with an RIR organization. For more information about RIR address allocation and updates, see [RIR Registration Updates](#) on page 437.

Adding IPv6 Networks

To add an IPv6 network:

1. Select the **Data Management** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network.
3. Select the **DHCP** tab -> **Networks** tab.
4. In the **Networks** section, click the Add drop-down list and select **IPv6 Network**.
5. In the *Add IPv6 Network* wizard, select one of the following and click **Next**:
 - **Add IPv6 Network**: Click this to add an IPv6 network from scratch.
 - **Add IPv6 Network using Template**: To use a template, click this, and then click **Select Template** and select an IPv6 network template. For information about network templates, see [About IPv6 Network Templates](#) on page 837. When you use a template to create a network, the configurations of the template apply to the new network. The appliance populates the template properties in the wizard when you click **Next**. You can then edit the pre-populated properties. If the template specified a fixed netmask, you cannot edit the netmask.
6. Complete the following and click **Next**:
 - **Regional Internet Registry**: This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#) on page 437. Complete the following to create an RIR IPv6 network container or network:
 - **Internet Registry**: Select the RIR from the drop-down list. The default is **None**, which means that the network is not associated with an RIR organization. When you select **RIPE**, the appliance displays **Organization ID** field where you can select an RIR organization.
 - **Organization ID**: Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
 - **Registration Status**: The default is **Not Registered**. When adding an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** check box below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**. The updated status and timestamp are displayed in the **Status of last update** field in the *IPv6 Network Container* or *IPv6 Network* editor.

- **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you are adding an existing RIR allocated network to NIOS, select **None**. When you are adding networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
- **Do not update registrations:** Select this check box if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
- **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
- **Netmask:** Use the netmask slider to select the appropriate number of subnet mask bits for the network.
- **Networks:** Do one of the following to add new networks:

Click the Add icon to enter a new network. If you are adding a network for a previously defined global IPv6 prefix, you can select the prefix from the **IPv6 Prefix** drop-down list. The default is **None**, which means that you are not creating an IPv6 network for a previously defined subnet route. If you have defined a global prefix at the Grid level, the default is the global prefix value. Click **Add** and Grid Manager adds a row to the table. Enter the network address in the **Network** field. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2001:0db8:0000:0000:0000:0000:0102:0304 can be shortened to 2001:db8::0102:0304. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Click **Add** again to add another network. You can also select a network and click the **Delete** icon to delete it.

or

Click the Next Available icon to have the appliance search for the next available network. Complete the following in the Next Available Networks section:

- **Create new network(s) under:** Enter the network container in which you want to create the new network. When you enter a network that does not exist, the appliance adds it as a network container. When you enter a network that is part of a parent network, the parent network is converted into a network container if it does not have a member assignment or does not contain fixed addresses and host records that are served by DHCP. You can also click **Select Network** to select a specific network in the *Network Selector* dialog box. For information about how the appliance searches for the next available network, see [Obtaining the Next Available Network](#) on page 844.
- **Number of new networks:** Enter the number of networks you want to add to the selected network container. Note that if there is not enough network space in the selected network to create the number of networks specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing networks in the table and you select one, the number you enter here includes the selected network.
- Click **Add Next** to add the networks. Grid Manager lists the networks in the table. You can click **Cancel** to reset the values.

Note: You must click **Add Next** to add the network container you enter in the Next Available Networks section. If you enter a network in the Next Available Networks section and then use the Add icon to add another network, the appliance does not save the network you enter in the Next Available Networks section until you click **Add Next**.

- **Comment:** Enter additional information about the network, such as the name of the organization it serves.

- **Automatically create reverse-mapping zone:** This function is enabled if the netmask of the network is a multiple of four, such as 4, 8, 12 or 16. Select this to have the appliance automatically create reverse-mapping zones for the network. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network view level.
 - **Disabled:** Select this if you do not want the DHCP server to provide DHCP services for this network at this time. This feature is useful when you are in the process of setting up the DHCP server. Clear this after you have configured the server and are ready to have it serve DHCP for this network.
7. Click **Next** and add one or more Grid members as DHCP servers for the network.
 - click the Add icon and select a Grid member from the *Member Selector* dialog box. Keep in mind, some DHCP properties for the network are inherited from this member. The network can be served by multiple members, but a member can serve networks in one network view only.
 8. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new network(s). Discovery settings differ based on whether you are defining one network or multiple networks.
 - **Configuring one network:** Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member. By default, all Polling Options discovery settings are inherited from the parent network (or Grid, if no parent exists) unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered.
or
 - **Configuring more than one network:** If the networks are child networks, they automatically inherit the settings of the parent network, including discovery settings and the discovery member. These settings will not appear in the wizard page. For discovery of multiple networks, you can only enable or disable **Immediate Discovery**. Click **Next** to override the DHCP properties described in [Defining General IPv6 Properties](#) on page 809.
 9. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#) on page 447. You can preview the information before the appliance submits updates to the RIPE database. To preview registration updates, click **Preview RIR Submissions**. For more information, see [Previewing Registration Updates](#) on page 451.

Note: You cannot leave an optional RIR attribute value empty. If you do not have a value for an RIR attribute, you must delete it from the table. You can enter up to 256 characters for all RIR attributes.

10. Save the configuration.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

After you create a network, you can do the following:

- Add it to a shared network. For more information, see [Adding IPv6 Shared Networks](#) on page 875.
- Use the split network feature to create subnets for the network. For information, see [Splitting IPv6 Networks into Subnets](#) on page 485.
- Use the join networks feature to create a parent network that encompasses multiple subnets into a larger network. For information, see [Joining IPv6 Networks](#) on page 485. You can also create a shared network for subnets that are on the same network segment.
- View a list of networks. For more information, see [Viewing Networks](#) on page 848.

Modifying IPv6 Networks

You can modify existing network settings and override the Grid or member DHCP properties, with the exception of the network address and netmask.

To modify an IPv6 network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
2. The *IPv6 Network* editor contains the following basic tabs from which you can modify data:
 - **Genera Basic:** You can modify the following fields:
 - **Comment:** The information you entered for the network.
 - **Disabled:** This field is displayed only if the selected network is a network without a child network under it. You can disable and enable existing networks instead of removing them from the database, if the selected network does not have a child subnet. This feature is especially helpful when you have to move or repair the server for a particular network.
 - **Member Assignment:** Add or delete a Grid member that provides DHCP services for this network.
 - **IPv6 DHCP Options:** Keep the inherited DHCP properties or override them and enter unique settings for the network. For information, see [Defining General IPv6 Properties](#) on page 809.
 - **RIR Registration:** Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#) on page 443.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **General Advanced:** You can associate zones with a network. For information, see [Associating Networks with Zones](#) on page 813.
 - **IPv6 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting IPv6 Networks

When you delete a network, all of its data, including all DHCP records, subnets, and records in its subnets, is deleted from the database. Because of the potentially large loss of data that can occur when you delete a network, the appliance stores the deleted network in the Recycle Bin. You can restore a deleted network from the Recycle Bin, if enabled. You can also disable a network instead of deleting it. For information, see [Modifying IPv6 Networks](#) on page 874.

To delete a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network* check box, and then select **Delete** or **Schedule Delete** from the Delete drop-down menu.
 2. To delete the network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#) on page 76.
- The appliance puts the deleted network in the Recycle Bin, if enabled.

ABOUT IPV6 SHARED NETWORKS

You can combine two or more contiguous IPv6 networks into a shared network. When you do, the DHCP server allocates IP addresses from both subnets. To create a shared network, create the individual subnets, and then create the shared network and add the subnets to it. For more information about shared networks, see [About Shared Networks](#) on page 779.

Adding IPv6 Shared Networks

To add an IPv6 shared network:

1. Select the **Data Management** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network.
3. Select the **DHCP** tab -> **Networks** tab.
4. In the **Shared Networks** section, select **IPv6 Shared Network** from the Add drop-down menu.
5. In the *Add IPv6 Shared Network* wizard, do the following:
 - **Name:** Enter the name of the shared network.
 - **Comment:** Enter information about the shared network.
 - **Disabled:** Select this if you want to enable the shared network at a later time. You can disable and enable existing networks instead of removing them from the database. This feature is especially helpful when you have to move or repair the server for a particular network.
6. Click **Next** and do the following to add networks:
 - a. Click the Add icon.
 - b. In the *Network Selector*, select the networks that you want to include in the shared network. Ensure that the networks are served by the same Grid members to avoid DHCP inconsistencies.
7. Click **Next** to configure DHCP properties described in [Defining General IPv6 Properties](#) on page 809.
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes for the shared network. For information, see [Using Extensible Attributes](#) on page 332.
9. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

For information on viewing shared networks, see [Viewing Shared Networks](#) on page 853.

Modifying IPv6 Shared Networks

To modify a shared network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared_network* check box, and then click the Edit icon.
2. The *IPv6 Shared Network* editor contains the following tabs from which you can modify data:
 - **General:** Modify the fields **Name**, **Comments**, and **Disabled** as described in [Adding IPv6 Shared Networks](#) on page 875.
 - **Networks:** Displays the networks that are currently assigned to the shared network. You can add or delete a network. To add a network, click the Add icon. To delete a network, select the *network* check box, and then click the Delete icon.
 - **IPv6 DHCP Options:** Keep the inherited DHCP properties or override them and enter unique settings for the shared network. For information, see [Defining General IPv6 Properties](#) on page 809.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

- **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
- 3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **IPv6 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the shared network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) on page 695.

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
- 4. Save the configuration and click **Restart** if it appears at the top of the screen.
 - or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting IPv6 Shared Networks

Though you can delete the networks in a shared network, a shared network must have at least one network in it.

To delete a shared network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared_network* check box, and then select **Delete** or **Schedule Delete** from the drop-down menu.
2. To delete the shared network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#) on page 76.

The appliance puts the deleted shared network in the Recycle Bin, if enabled.

CONFIGURING IPV6 ADDRESS RANGES

You can configure IPv6 ranges that are used to delegate IPv6 prefixes only, to assign IPv6 addresses only, or to delegate IPv6 prefixes and assign IP v6 addresses. When you define a DHCP range to delegate prefixes, the prefixes can be outside of the network where they are being defined. IPv6 ranges inherit their properties from their network, so each range in a subnet provides the same set of options to their DHCP clients.

Note that when an Infoblox DHCP server grants IPv4 leases, it starts from the last IP address in the range to the first. When the server grants IPv6 leases, it uses an algorithm based on the DUID of the client.

Adding IPv6 Address Ranges

To add a an IPv6 address range:

1. Navigate to the IPv6 network to which you want to add an address range, and then select **Range** from the Add drop down menu.
 - or
 - From any panel in the DHCP tab, expand the Toolbar and click **Add** -> **Range** -> **IPv6**.
2. In the *Add IPv6 Range* wizard, select one of the following and click **Next**:
 - **Add IPv6 Range:** Select this to add an address range from scratch.
 - or
 - **Add IPv6 Range Using Template**

Click **Select Template** and select the template that you want to use. Note that when you use a template to create a DHCP range, the configurations of the template apply to the new range. The appliance automatically populates the address range properties in the wizard. You can then edit the pre-populated properties. For more information, see [About IPv6 Range Templates](#) on page 834.

3. Complete the following:

- **Network:** Click **Select Network**. Grid Manager displays the network address here if you have only one network configured. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.

Specify one of the following:

- **Address:** Select this if the address range is used to allocate IPv6 addresses only to DHCP clients, and then enter the start and end addresses in the range.
- **Prefix Delegated:** Select this if the DHCP server uses this address range to delegate IPv6 prefixes only to DHCP clients. Enter the start and end prefixes, and the prefix length.
- **Both:** Select this if the DHCP server delegates IPv6 prefixes and allocates IPv6 addresses from this range. Enter the start and end addresses in the range, and the start and end prefixes, and the prefix length.

Complete the following:

- **Name:** Enter a name for the address range.
- **Comment:** Enter additional information about the address range.
- **Disabled:** Select this if you want to save the configuration for the address range but do not want to activate the address range yet. You can clear this check box when you are ready to allocate addresses from this range.

4. Click **Next** and select one of the following to provide DHCP services for the DHCP range:

- **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
- **Grid Member:** Select this if you want a Grid member to serve DHCP for this DHCP range. Select a Grid member from the drop-down list. The drop-down list displays only the Grid members that are associated with the network to which the DHCP range belongs.

5. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new DHCP range.

- **Configuring one network:** Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member. By default, all Polling Options discovery settings are inherited from the parent network unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. See the section [Configuring Grid Properties for Discovery](#) on page 529 for a complete description of discovery Polling Options.

6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.7. Save the configuration and click **Restart** if it appears at the top of the screen.

or

- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Setting the Priority of IPv6 Address Ranges

The DHCP server allocates IP addresses from the configured DHCP ranges according to the order in which the ranges are listed. By default, ranges are listed according to their start addresses. You can move the ranges up and down in the list to change their order. For information about viewing DHCP ranges and other objects in a network, see [Viewing IPv6 DHCP Objects](#) on page 880.

To change the order of DHCP ranges in a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network*.
2. Expand the Toolbar and click **Order DHCP Ranges**.
3. In the *Order DHCP Ranges* dialog box, click the up and down arrows to move ranges up or down on the list. The Priority value changes accordingly. Click **OK** to save the configuration.

Modifying IPv6 Address Ranges

To modify an IPv6 address range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.
2. The *DHCP Range* editor contains the following basic tabs from which you can modify data:
 - **General:** Modify the fields, except the network address, as described in [Adding IPv6 Address Ranges](#) on page 876.
 - **Member Assignment:** Modify the Grid member that provides DHCP services for the DHCP range as described in [Adding IPv6 Address Ranges](#) on page 876.
 - **IPv6 DHCP Options:** Keep active leases in a deleted DHCP range. For more information, see [Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges](#) on page 814.
 - **Extensible Attributes:** You can add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
 - **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use to assign to clients. You can use these exclusion addresses as static IP addresses. For more information, see [About Exclusion Ranges](#) on page 780.

Note that Grid Manager displays both the basic and advanced mode tabs the next time you log in to the GUI.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting IPv6 Address Ranges

To delete an IPv6 address range:

- From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Delete icon.

CONFIGURING IPV6 FIXED ADDRESSES

You can configure IPv6 fixed addresses with either an IPv6 address or prefix. You can assign prefix-based fixed addresses to routers so they can advertise the prefixes associated with a link. DHCP hosts, in turn, use these prefixes to generate IP addresses using the stateless autoconfiguration mechanism defined in *RFC 2462, IPv6 Stateless Autoconfiguration*. You can also create IPv6 fixed addresses from the Tasks Dashboard. For information about the Tasks Dashboard, see [The Tasks Dashboard](#) on page 99.

Note that dynamic DNS updates for IPv6 fixed addresses are not supported.

Adding IPv6 Fixed Addresses

To add an IPv6 fixed address:

1. Navigate to the network to which you want to add a fixed address, and then select **Fixed Address** from the Add drop down menu.

or

From any panel in the DHCP tab, expand the Toolbar and click **Add -> Fixed Address -> IPv6**.

2. In the *Add Fixed Address* wizard, select one of the following and click **Next**:

- **Add IPv6 Fixed Address**

or

- **Add IPv6 Fixed Address Using Template**

Click **Select Template** and select the template that you want to use. Note that when you use a template to create a fixed address, the configurations of the template apply to the new address. The appliance automatically populates the fixed address properties in the wizard. You can then edit the pre-populated properties.

3. In this panel, the displayed network address can either be the last selected network or the network from which you are adding the fixed address. If no network address is displayed or if you want to specify a different network, click **Select Network**. When there are multiple networks, Grid Manager displays the *Select Network* dialog box.

Specify one of the following:

- Select **Address** to assign an IPv6 address to a fixed address. You can either enter an IPv6 address or select **Next Available IP** to obtain the next available IP address.

Note: When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

- Select **Prefix Delegated** to assign an IPv6 prefix. Enter the prefix and prefix length.
- Select **Both** to assign an IPv6 prefix and address. Enter the IPv6 address, prefix, and prefix length.

Complete the following:

- **DUID:** Specify the DHCP Unique Identifier (DUID) of the DHCP client assigned to this fixed address.
- **Name:** Enter a name for the fixed address.
- **Comment:** Optionally, enter additional information.
- **Disabled:** Select this if you do not want the DHCP server to allocate this IP address at this time.

4. Click **Next** to configure or override DHCP options as described in [Defining General IPv6 Properties](#) on page 809.

5. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new fixed IP address.

- Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the fixed IP address. If you choose **Enable Immediate Discovery**, discovery will execute on the host after you save your settings. You may also choose to leave both options disabled.
- By default, the new host inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** checkbox and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials. See the section [Configuring Grid SNMPv1/v2 Properties](#) on page 530 for a complete description of SNMP credentials for discovery.

6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

7. Save the configuration and click **Restart** if it appears at the top of the screen.

or

- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

For information on viewing IPv6 fixed addresses in a network, see [Viewing IPv6 DHCP Objects](#) on page 880.

Modifying IPv6 Fixed Addresses

To modify a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed_address* check box, and then click the Edit icon.
2. The *Fixed Address* editor contains the following basic tabs from which you can modify data:
 - **General:** You can modify all the fields you filled out in the first step of the wizard described in [Adding IPv6 Fixed Addresses](#) on page 878.
 - **IPv6 DHCP Options:** You can keep the inherited DHCP options or override them and enter unique settings for the fixed address. For information, see [Defining General IPv6 Properties](#) on page 809.
 - **Discovered Data:** You can view discovered data of this address, if any, in this tab. For information, see [Viewing Discovered Data](#) on page 510.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting IPv6 Fixed Addresses

To delete a fixed address, from the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed_address* check box, and then click the Delete icon.

VIEWING IPV6 DHCP OBJECTS

You can view the DHCP objects in an IPv6 network by navigating to the **DHCP** tab -> **Networks** tab -> **Networks** panel, and then clicking the network link. This panel displays the following information about DHCP objects in the selected IPv6 network:

- **IP Address:** The IPv6 address of a DHCP object, such as a DHCP range, fixed address, or host configured for DHCP, or roaming host with an allocated IP address. For a DHCP range, this field displays the start and end addresses of the range. For a host that has multiple IP addresses, each IP address is displayed separately. Note that the appliance highlights all disabled DHCP objects in gray.
- **Type:** The DHCP object type, such as **IPv6 DHCP Range** or **IPv6 Fixed Address**.
- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Comment:** The information you entered for the object.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes.

You can select the following additional columns for display:

- **Priority:** Displays the priority of the DHCP range.
- **Disabled:** Indicates whether the network is disabled.

You can also do the following in this panel:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only.
- Sort the data in ascending or descending order by column.

- Create a bookmark for the object.
- Click **Go to IPAM View** to view information about the object in the **IPAM** tab.
- Delete or schedule the deletion of a selected object or multiple objects.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print or export the data.



Chapter 28 DHCP Failover

This chapter explains how to configure DHCP failover associations. It contains the following sections:

- [*About DHCP Failover*](#) on page 884
 - [*Failover Association Operations*](#) on page 884
- [*Configuring Failover Associations*](#) on page 885
 - [*Adding Failover Associations*](#) on page 886
- [*Managing Failover Associations*](#) on page 887
 - [*Modifying Failover Associations*](#) on page 887
 - [*Monitoring Failover Associations*](#) on page 888
 - [*Deleting Failover Associations*](#) on page 889
 - [*Setting a Peer in the Partner-Down State*](#) on page 889
 - [*Performing a Force Recovery*](#) on page 890

ABOUT DHCP FAILOVER

You can create a failover association between two DHCP servers (a primary server and a secondary) and assign the failover association to serve an IPv4 DHCP range. When you set up a failover association, you greatly reduce DHCP service downtime if one of your DHCP servers is out of service. You can better manage IP address requests by making two servers available for DHCP services. You can also configure one of the servers to assume full DHCP services when you know the other server may go out of service for a period of time.

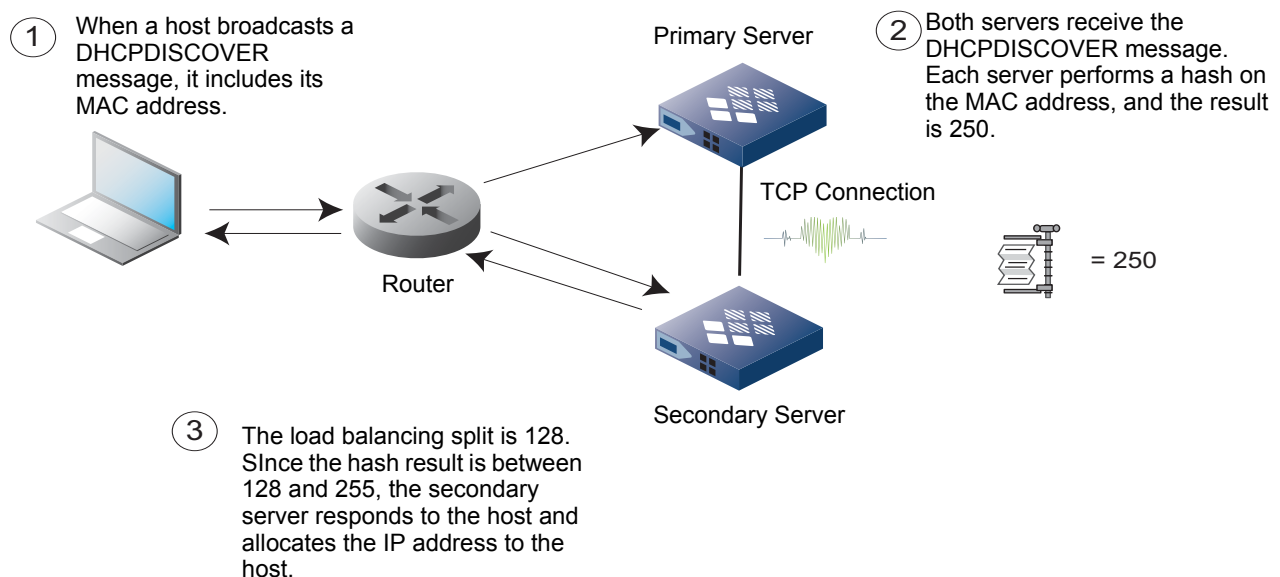
You can configure two NIOS appliances, or one appliance and one external server, to form a failover association. The pairing of a primary and secondary server is called a peer association. The failover peers establish a TCP connection for their communication. They share a pool of IP addresses that they allocate to hosts on their networks based on load balancing. Load balancing is a technique to split the address allocation workload evenly across the two DHCP servers. You can assign a DHCP failover association to serve DHCP ranges in a network. A DHCP failover association can serve DHCP ranges that belong to one network view only. It cannot serve ranges in different network views.

Failover Association Operations

When a host broadcasts a DHCPDISCOVER message, it includes its MAC address. Both the primary and secondary peers receive this message. To determine which server should allocate an IP address to the host, they each extract the MAC address from the DHCPDISCOVER message and perform a hash operation. Each server then compares the result of its hash operation with the configured load balancing split. The split is set to 50% by default to ensure an even split between the two servers. When the split is 50%, the primary server allocates the IP address if the hash result is between 1 and 127, and the secondary server allocates the IP address if the hash result is between 128 and 255. As a server allocates an IP address, it updates its peer so their databases remain synchronized.

As shown in [Figure 28.1](#), when a host broadcasts a DHCPDISCOVER message, both the primary and secondary servers receive the message. They perform a hash operation on the MAC address in the DHCPDISCOVER message, and the result is 250. Since the load balancing split is 50% and the hash result is 250, the secondary server responds to the host with a DHCPOFFER message. The secondary peer allocates an IP address from its assigned pool of IP addresses. It then sends a lease update message to the primary server so that the primary server knows how the address is assigned and can properly take over if the secondary server fails.

Figure 28.1 Load Balancing and IP Addresses Allocation



CONFIGURING FAILOVER ASSOCIATIONS

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar. Click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, select the **General Advanced** tab to complete the following:
 - **Failover Port:** You can modify the port number that members use for failover associations. You can use any available port from 519 to 647. The default is 647 for a new installation and 519 for an upgrade.

The following are tasks and guidelines for configuring a DHCP failover:

1. Identify the primary and secondary DHCP servers and ensure that the appliances are set up correctly for the failover association, using the following guidelines:
 - Configure a failover association using two NIOS appliances, or a NIOS appliance and an ISC DHCP compliant server.
 - One of the DHCP servers must be an independent appliance or in an Infoblox Grid.
 - The DHCP servers do not have to be in the same geographic location.
 - The clocks on both servers must be synchronized. This happens automatically when both servers are on the same Grid.
 - Both servers must use the same version of the DHCP configuration file. This happens automatically when both servers are on the same Grid.
 - If you use firewalls on your networks, ensure that the firewalls allow TCP port 519 between the servers, and that TCP port 7911 is open for partner down operations.
 - Each pair of DHCP servers can participate in only one failover association. An appliance can participate in more than one failover association, as long as it is with a different peer.

Configure the same DHCP properties on the primary and secondary servers, as described in [Configuring General IPv4 DHCP Properties](#) on page 793.

- Both the primary and secondary servers must have the same operational parameters, and they must be able to receive DHCPDISCOVER messages that hosts broadcast on the networks.
- If you change any of the DHCP failover parameters for a peer association definition, you must make the same changes on both the primary and secondary servers.

Note: If both the primary and secondary servers are in a Grid, you configure the properties on the failover association and the configuration applies to both servers.

2. Create a failover association and configure load balancing between the servers. For information, see [Adding Failover Associations](#) on page 886.
 - Ensure that you use the same failover association name on both the primary and secondary servers.
 - The appliance assigns default values to the failover timers and triggers. In general, these default values serve the purpose of a failover. Do not change these values unless you understand the ramification of the changes. For example, when one of the peers in a failover association fails, the other peer goes into a COMMUNICATIONS-INTERRUPTED state, and the lease time changes to the MCLT (Maximum Client Lead Time). You should consider how the MCLT affects the lease time when a failover occurs if you want to change this value.
3. Assign the failover association to the DHCP ranges in the same network view. Failover associations can serve only IPv4 DHCP ranges. For information, see [Configuring IPv4 Address Ranges](#) on page 854.
 - If you configure a shared network, and the subnets in the shared network contain ranges served by a DHCP failover association, both the primary and secondary DHCP server must have the same shared networks defined, containing the same networks and DHCP ranges.

Note: If you have multiple networks that are in a shared network and you plan to use a DHCP failover, you must use the same failover association and specify the same peers on all the networks in the shared network.

4. Enable DHCP on the primary and secondary servers AFTER you complete all the configurations. For information, see [Managing Failover Associations](#) on page 887.

Note: When you set up a failover association for the first time, ensure that both servers are up and running and their databases are synchronized before they can start assigning IP addresses.

When you configure a failover association, the appliance assigns default values for timers and triggers, such as the MCLT and the maximum number of “unacked” packets. A failover may occur when some of the timers expire or when a failover peer goes out of service. When a failover occurs, the functional peer takes over and assigns IP addresses with the lease time set to the MCLT. When the server that is offline comes back online, it synchronizes its database with its peer before it starts allocating IP addresses.

Adding Failover Associations

To add a DHCP failover association, perform the following procedures on both the primary and secondary servers:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **IPv4 Failover Associations** section, and then click the Add icon.
or
Expand the Toolbar and click **Add** -> **IPv4 Failover Association**.
2. In the *Add Failover Association* wizard, complete the following:
 - **Name:** Enter a unique name for the failover association. The failover association name is case sensitive. Enter the same name on both the primary and secondary servers. The appliance validates the names on both servers. The names must be exactly the same. If they do not match, the failover association goes into disconnect mode.
 - **DHCP Failover Primary:** Select one of the following. The default is **Grid Member**.
 - **Grid Member:** Click **Select member**. In the *Select Member* dialog box, select the primary server and click the Select icon.
 - **External Server IP Address:** Select this to use an external ISC DHCP compliant server as the primary server. Enter the IP address of the primary server in the field.
 - **DHCP Failover Secondary:** Select one of the of following. The default is **Grid Member**.
 - **Grid Member:** Click **Select member**. In the *Select Member* dialog box, select the secondary server and click the Select icon.
 - **External Server IP Address:** Select this to use an external ISC DHCP compliant server as the secondary server. Enter the IP address of the secondary server in the field.

Note: You cannot select **External Server IP Address** for both the primary and secondary servers. One of the servers must be an independent appliance or in an Infoblox Grid.

- **Comment:** Enter useful information about the failover association.
3. Click **Next** and do the following to control the IP address allocation between the peers and how they switch from one to the other based on the configuration:
 - **Load Balancing Data:** Adjust the slider to determine which server should handle more IP address requests. The default is 50%. When you adjust the slider, a tooltip window displays the percentage of available IP addresses that each server can allocate. When you move the slider all the way to the left, the primary server responds to all IP address requests, and the secondary server does not respond to any. The opposite applies when you move the slider all the way to the right. Infoblox recommends that you use the default (50/50) to enable the primary and secondary servers to respond to IP address requests on an equal basis.

- **Lease Deletion:** Select the following to override settings at the Grid and member levels.
 - **Keep leases from deleted ranges until one week after expiration:** When you select this and delete a DHCP range with active leases, the appliance stores these leases up to one week after they expire. When you add a new DHCP range that includes the IP addresses of these active leases, the appliance automatically restores the leases.
- 4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
- 5. Save the configuration and click **Restart** if it appears at the top of the screen.

MANAGING FAILOVER ASSOCIATIONS

After you establish a failover association, you can monitor its status periodically to ensure that it is functioning properly. You can also delete a failover association when it is not assigned to any DHCP range.

See the following sections on how to manage failover associations:

- [Modifying Failover Associations](#)
- [Monitoring Failover Associations](#) on page 888
- [Deleting Failover Associations](#) on page 889

Under special circumstances, you can manually adjust the configuration of a failover association. For example, when you know in advance that a peer will be out of service for an extended period of time, you can manually set the functional peer in a PARTNER-DOWN mode. This allows the functional partner to assume all leases and be able to allocate addresses to client requests in full capacity. In addition, when you suspect the databases in a failover association are not synchronized, you can consider doing a force recovery (after you consult with Infoblox Technical Support or your Infoblox representative) so the secondary server can completely rebuild its lease table with updates from the primary server.

See the following sections on how to set a peer to the partner-down mode and perform a force recovery:

- [Setting a Peer in the Partner-Down State](#) on page 889
- [Performing a Force Recovery](#) on page 890

Modifying Failover Associations

To modify a failover association:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Failover Associations** -> *failover_association* check box, and then click the Edit icon.
2. The *DHCP Failover Association* editor contains the following tabs from which you can modify data:
 - **General:** In the **Basic** tab, modify the fields as described in [Adding Failover Associations](#) on page 886. In the **Advanced** tab, complete the following to modify the port number you use for the failover association:
 - **Failover Port:** Click **Override** to enter a port number for the failover association. You can use any available port from 1 to 63999. The default is 647 for a new installation and 519 for an upgrade.
 - **Triggers:** Before editing the triggers and timers, ensure that you understand the ramification of the changes. Improper configuration of the triggers can cause the failover association to fail. For information about the fields in the **Basic** tab, see [Adding Failover Associations](#) on page 886. The following are the triggers in the **Advanced** tab:
 - **Max Response Delay Before Failover(s):** Specifies how much time (in seconds) can transpire before a failover occurs when a failover peer does not receive any communication from its peer. This number should be small enough that the transient network failure does not leave the servers out of communication for a long time, but big enough that the servers are not constantly connecting and disconnecting. The default is 60 seconds.

- **Max Number of Unacked Updates:** Specifies the number of “unacked” packets the server can send before a failover occurs. The default is 10 messages.
- **Max Client Lead Time(s):** Specifies the length of time that a failover peer can renew a lease without contacting its peer. The larger the number, the longer it takes for the peer to recover IP addresses after moving to the PARTNER-DOWN state. The smaller the number, the more load your servers experience when they are not communicating. The default is 3600 seconds.
- **Max Load Balancing Delay(s):** Specifies the cutoff after load balancing is disabled. The cutoff is based on the number of seconds since a client sent its first DHCPDISCOVER message. For instance, if one of the failover peers gets into a state where it is busy responding to failover messages but is not responding to other client requests, the other peer responds to the client requests when the clients retry. This does not cause a failover. The default is three seconds.
- **Extensible Attributes:** Add and delete extensible attributes that are associated with a failover association. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

Monitoring Failover Associations

After you configure a failover association, the peers establish a TCP connection for communication. In a normal operational state, they send keepalive messages and database updates every time they grant a lease. However, there are times when the failover association experiences problems and goes into a state other than NORMAL. You can monitor the overall state of a failover association and the individual status of the peers to verify that the servers are operating and communicating properly.

Both peers in a failover association maintain the same DHCP fingerprinting state (enabled or disabled) even when one of the peers fails or becomes operational again. Note that both peers must be in the same Grid for the fingerprinting state to stay the same. For information about DHCP fingerprinting, see [DHCP Fingerprint Detection](#) on page 1031.

In this panel, you can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.

To monitor the failover association status:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **IPv4 Failover Associations** section. Grid Manager displays the list of failover associations and their overall status.
2. To view detailed information about a failover association, select the *failover_association* check box, and then click the Show Status icon.
3. In the *Failover Association Status* dialog box, Grid Manager displays the overall status of the failover association and the status of both the primary and secondary servers.

The failover association can be in one of the following states:

- **OK** (green): The failover association is functioning properly.
- **DEGRADED** (yellow): The failover association is degraded when one of the peers is giving out limited addresses.
- **FAILURE** (red): The failover association is not functioning, may be because it is not completely configured. The peers are not assigning IP addresses.

For each peer, Grid Manager displays the hostname or IP address, the status, and event date. The peer can be in one of the following states:

- **STARTUP:** The server is starting up.
- **NORMAL:** The server is in a normal operational state in which it responds to its load balancing subset of DHCP clients.
- **PAUSED:** This state allows a peer to inform the other peer that it is going out of service for a short period of time so the other peer can immediately transition to the COMMUNICATIONS-INTERRUPTED state and start providing DHCP service to DHCP clients.

- COMMUNICATIONS-INTERRUPTED: The servers are not communicating with each other. Both servers provide DHCP service to DHCP clients from which they receive DHCP requests.
- PARTNER-DOWN: The server assumes control of the DHCP service because its peer is out of service.
- RECOVER: The server is starting up and trying to get a complete update from its peer and discovers that its peer is in the PARTNER-DOWN state.
- RECOVER-WAIT: The server has got a complete update from its peer and is waiting for MCLT period to pass before transitioning to the RECOVER-DONE state.
- RECOVER-DONE: The server completed an update from its peer.
- POTENTIAL-CONFLICT: The peers are not synchronized due to an administrative error or an incorrect state transition. Check the failover configuration and correct the error.
- CONFLICT-DONE: This is a temporary state that the primary server enters after it received updates from the secondary server when it was in the POTENTIAL-CONFLICT state.
- RESOLUTION-INTERRUPTED: The server responds to DHCP clients in a limited way when it is in this state.
- UNKNOWN: The DHCP server is in an unknown state. The failover association is not functioning properly, may be because it is configured improperly. For example, failover association is not assigned to any DHCP range.
- SHUTDOWN: This state allows a peer to inform the other peer that it is going out of service for a long period of time so the other peer can immediately transition to the PARTNER-DOWN state and completely assume control of the DHCP service.

Deleting Failover Associations

You cannot delete a failover association if it is currently assigned to a DHCP range. If you want to delete a failover association, ensure that it is not assigned to any DHCP range.

To delete a failover association:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Failover Associations** → *failover_association* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.
The appliance puts the failover association in the recycle bin, if enabled.

Setting a Peer in the Partner-Down State

If one of the peers in a failover association is out of service for an extended period of time, you should consider putting the functional peer in the PARTNER-DOWN state. When you place the functional peer in the PARTNER-DOWN state, it assumes full DHCP services for the networks. Since the functional server may not receive all the updates from its peer, it extends all the leases on the MCLT. Once the following conditions are met, the functional peer provides DHCP services autonomously:

- It has reclaimed all the leases that belonged to its peer.
- The MCLT has passed.

When the peer that is offline comes back online, it synchronizes with the functional peer and reestablishes the communication before it provides DHCP services to the clients.

WARNING: BEFORE YOU PUT A PEER IN THE PARTNER-DOWN STATE, ENSURE THAT THE OTHER PEER IS INDEED OUT OF SERVICE. IF BOTH THE PRIMARY AND SECONDARY SERVERS ARE OPERATIONAL WHEN YOU PLACE ONE OF THEM IN THE PARTNER-DOWN MODE, BOTH SERVERS MAY STOP ISSUING LEASES FOR A MINIMUM OF TIME DEFINED IN THE MCLT.

To set a peer in the PARTNER-DOWN state:

1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Failover Associations** → *failover_association* check box.

2. Expand the Toolbar and click **Set Partner Down**.
3. In the *Set Failover Association Partner Down* dialog box, select one of the following:
 - **Primary**: Select this if the secondary server is out of service.
 - **Secondary**: Select this if the primary server is out of service.
4. Click **OK**.

Performing a Force Recovery

When the primary and secondary peers are not synchronized, you can perform a force recovery to set the primary server in the PARTNER-DOWN state while putting the secondary server in the RECOVER state. During a force recovery, all leases in the databases are resynchronized. When you perform a force recovery, the secondary server does not serve any DHCP leases for a minimum of the MCLT while it resynchronizes with the primary server. Before you perform a force recovery, consult with Infoblox Technical Support or your Infoblox representative to ensure that the force recovery is appropriate for the situation.

To perform a force recovery:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab-> **Failover Associations** -> *failover_association* check box.
2. Expand the Toolbar and click **Force Recovery State**.
3. In the *Force Secondary Peer Recovery State* dialog box, click **OK**.
The appliance synchronizes the databases on the primary and secondary servers.



Chapter 29 Configuring IPv4 DHCP Filters

This chapter explains how to configure IPv4 DHCP filters. It contains the following sections:

- [About IPv4 DHCP Filters](#) on page 892
 - [IP Address Allocation](#) on page 892
 - [IP Address Allocation Using Filters](#) on page 895
- [About MAC Address Filters](#) on page 897
 - [Defining MAC Address Filters](#) on page 898
 - [Adding MAC Address Filter Items](#) on page 899
- [About Relay Agent Filters](#) on page 899
 - [Defining Relay Agent Filters](#) on page 900
- [About Option Filters](#) on page 901
 - [Defining Option Filters](#) on page 902
 - [Configuring User Class Filters](#) on page 904
 - [Configuration Example: Using Option Filters](#) on page 905
- [About DHCP Fingerprint Filters](#) on page 906
 - [Defining DHCP Fingerprint Filters](#) on page 907
- [Applying Filters to DHCP Address Ranges](#) on page 907
 - [Adding Filters to the Class Filter List](#) on page 907
 - [Adding Filters to the Logic Filter List](#) on page 908
 - [Configuration Example: Using the Class and Logic Filter Lists](#) on page 909
- [Managing DHCP Filters](#) on page 912
 - [Modifying DHCP Filters](#) on page 912
 - [Viewing DHCP Filters](#) on page 914
 - [Deleting Filters](#) on page 914

ABOUT IPv4 DHCP FILTERS

To control how the appliance allocates IPv4 addresses, you can define DHCP filters and apply them to address ranges and range templates. Depending on your configuration, DHCP filters screen requesting clients by matching MAC addresses, relay agent identifiers, DHCP options, or DHCP fingerprints you define in the filters. If you configure DHCP servers in the Grid to send authentication requests to a RADIUS authentication server group, you can also filter requests by matching the authentication results. (For information about this feature, see [Chapter 30, Authenticated DHCP](#), on page 915.)

When you define DHCP filters, you classify DHCP clients based on the information provided by the clients or by the RADIUS server. When you apply filters to an address range, the appliance responds to your address requests based on your configuration. The appliance also decides which DHCP options to return to the matching clients based on how you apply the filters. For more information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

You can use filters to control address allocation based on your network requirements. For example, you can use DHCP filters to screen unmanaged hosts on a network by denying their address and option requests. If you have multiple DHCP address ranges on the same network and you want to assign IP addresses from specific address ranges to specific clients, you can use filters to screen the address assignments. For information, see [IP Address Allocation](#).

The appliance supports the following filters:

- **MAC address filters:** Use MAC addresses as matching criteria for granting or denying address requests. For information, see [About MAC Address Filters](#) on page 897.
- **Relay agent filters:** Identify remote hosts by matching the relay agent identifiers in the DHCPDISCOVER messages. For information, see [About Relay Agent Filters](#) on page 899.
- **Option filters:** Classify hosts by matching the DHCP options and values sent by the requesting hosts. For information, see [About Option Filters](#) on page 901.
- **DHCP fingerprint filter:** Identify remote clients by matching the option number sequence or vendor ID sent in option 55 and 60 of the DHCP request against the DHCP fingerprints cached on the system. For information about DHCP fingerprint filters, see [About DHCP Fingerprint Filters](#) on page 906. For information about DHCP fingerprint detection, see [DHCP Fingerprint Detection](#) on page 1031.
- **NAC filters:** Use authentication results from a RADIUS authentication server group as matching criteria for granting or denying address requests. For information, see [Chapter 30, Authenticated DHCP](#), on page 915.

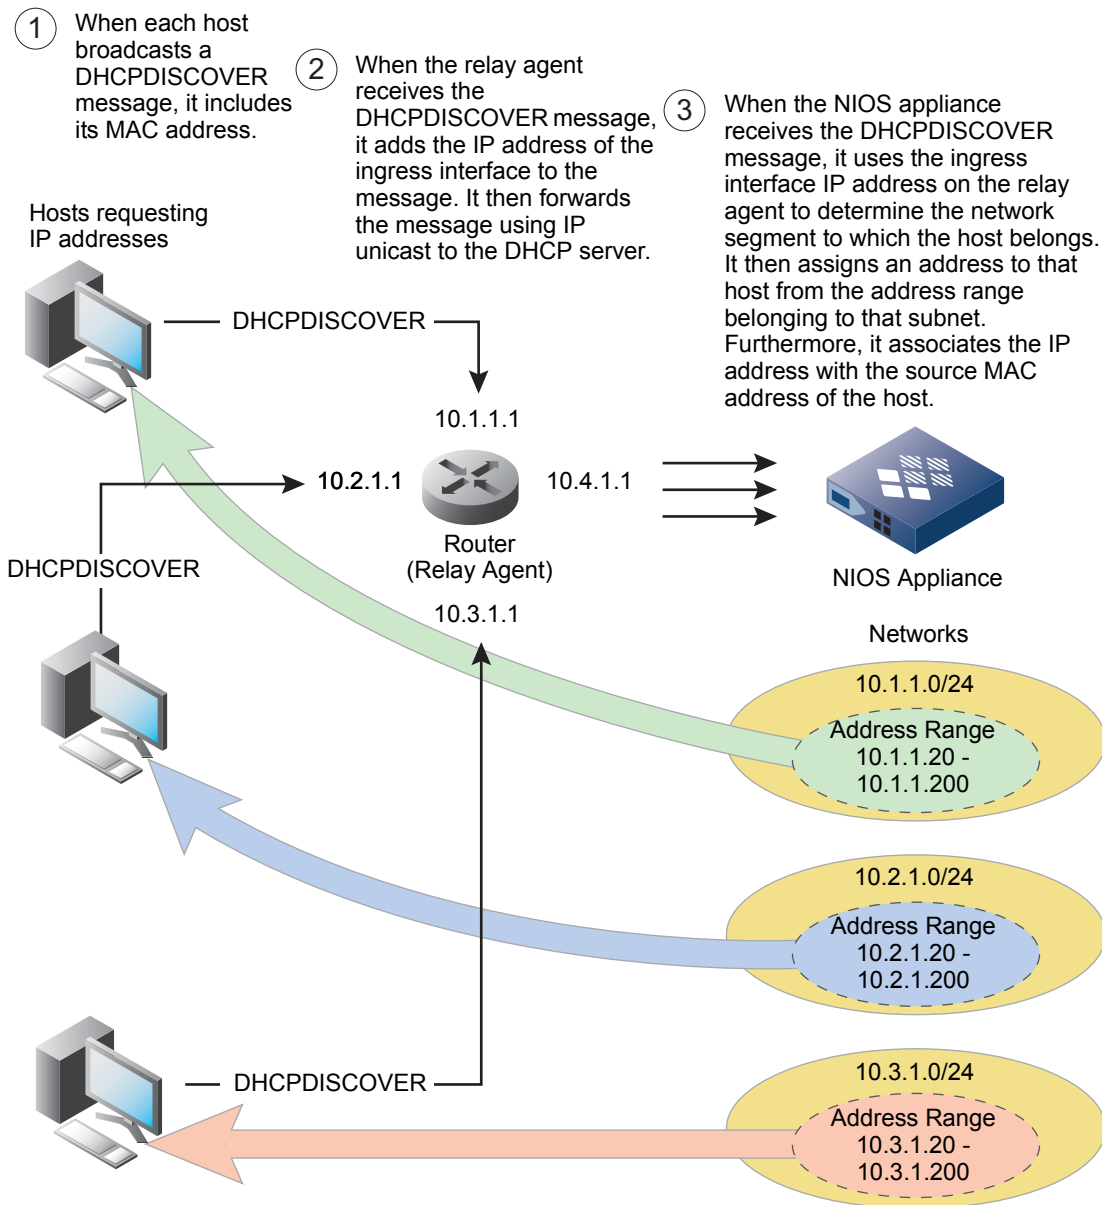
You can use MAC, option, and NAC filters to define DHCP options that matching clients can receive. Depending on how you apply a filter to an address range, all DHCP clients with matching criteria can receive all or some of the DHCP options defined in the filter. DHCP options defined for a matching filter supersede those defined at the Grid, member, network, and DHCP range levels. Options defined for a filter that is in the Class Filter List of an address range supersede those defined in the Logic Filter List. For more information about how the appliance returns options and how to apply DHCP filters, see [Applying Filters to DHCP Address Ranges](#) on page 907.

IP Address Allocation

When a DHCP client requests an IP address, the NIOS appliance draws an address from an address range associated with the network segment for that client. Because you define that range, you can thereby control the IP address (within the defined range) and the associated TCP/IP settings that the client receives.

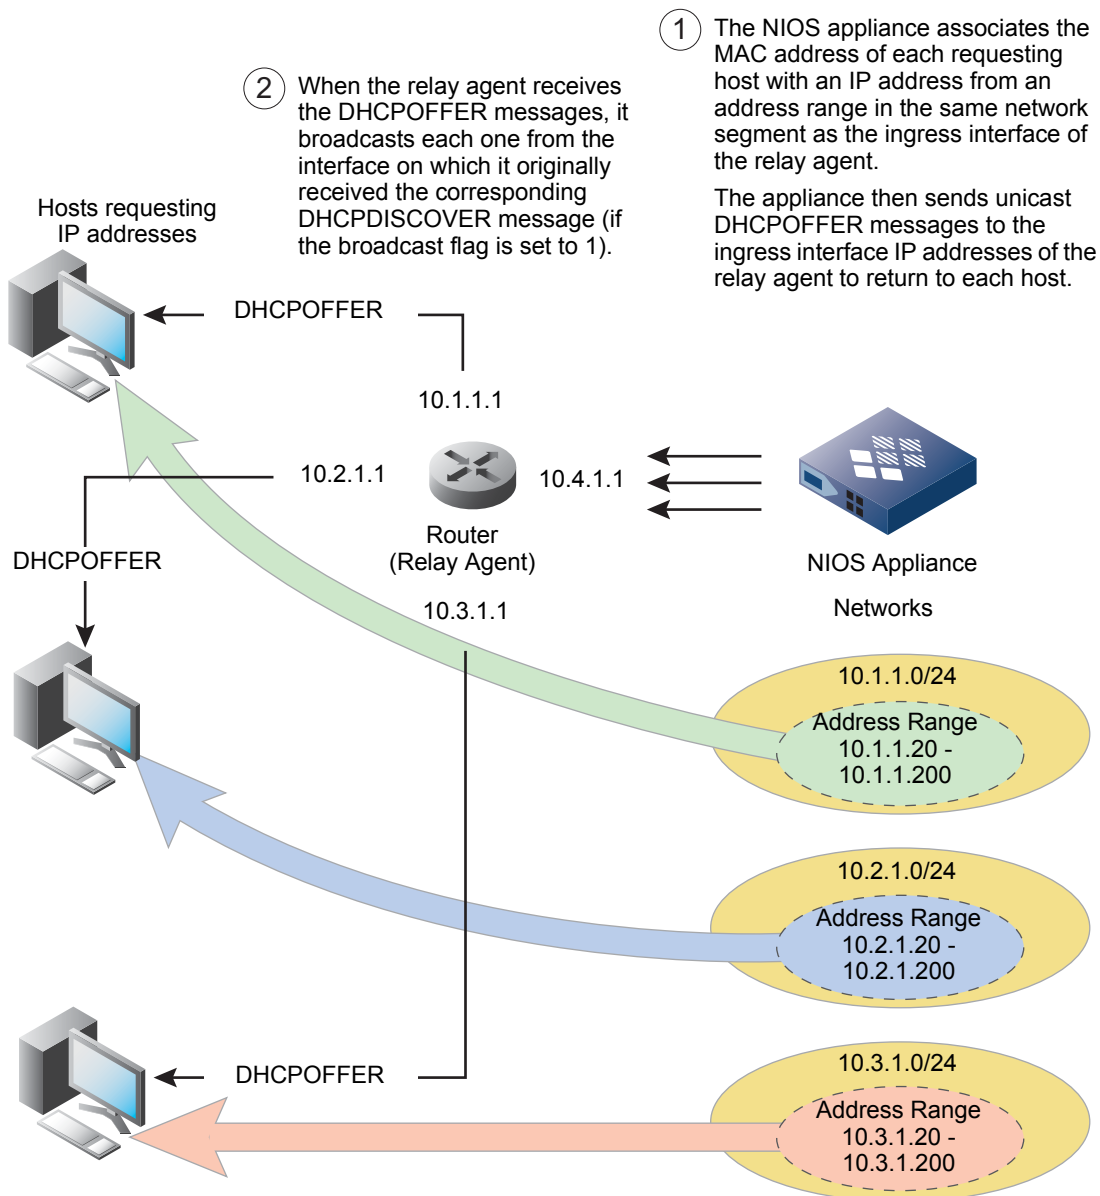
In [Figure 29.1](#), three hosts—each in a different subnet—request an IP address. Each one broadcasts a DHCPDISCOVER message, which includes its MAC address. When the router, which also functions as a DHCP relay agent, receives the message, it adds the IP address of the interface on which the message arrives and forwards the message to the DHCP server—or servers—previously configured on the router. When the NIOS appliance receives the message, it uses the ingress interface IP address of the router to determine the network segment to which the host belongs and associates the MAC address of the requesting host with an IP address from an address range for that network.

Figure 29.1 Requesting Addresses – DHCPDISCOVER Messages



The NIOS appliance replies to DHCPREQUEST messages by sending DHCPOFFER messages through the relay agent to the requesting hosts, as shown in [Figure 29.2](#) on page 894.

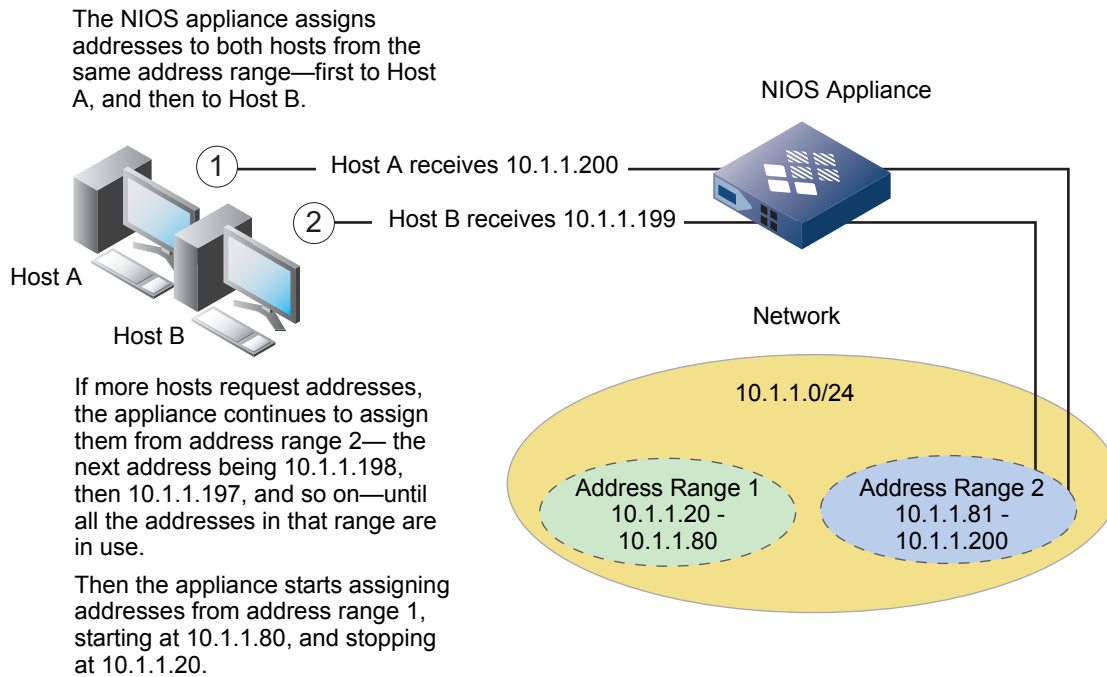
Figure 29.2 Requesting Addresses – DHCP OFFER Messages



The addressing scheme depicted in [Figure 29.1](#) on page 893 and [Figure 29.2](#) is fairly simple: each network has a single address range. Consequently, address assignments are fairly straightforward. However, if you have multiple address ranges in the same network and you want to assign addresses from specific address ranges to specific hosts, you must screen the address assignments through the use of filters. If you do not apply a filter, the NIOS appliance assigns addresses from the highest address range to the lowest range and within each range from the highest address to the lowest address. That is, the appliance chooses the range with the highest addresses first (that is, closest to 255) and begins assigning addresses exclusively from that range, starting with the highest address and finishing with the lowest (closest to 0). When all the addresses from that range are in use, it then begins assigning addresses from the next highest range, and so on, finishing with the range with the lowest addresses. This is shown in [Figure 29.3](#) on page 895.

Note: After the DHCP server runs for a while, it assigns leases based on when it last used addresses, and not just on their positions in the range.

Figure 29.3 Multiple Address Ranges without Filters



IP Address Allocation Using Filters

To control the assignment of addresses from specific address ranges to specific hosts, the NIOS appliance provides the following filters:

- A MAC address filter to which you add MAC addresses as filter criteria. For information, see [About MAC Address Filters](#) on page 897.
- A relay agent filter with configured circuit ID and remote ID as specified by the relay agent (DHCP option 82). For information, see [About Relay Agent Filters](#) on page 899.
- An option filter in which you specify DHCP options and matching values. For information, see [About Option Filters](#) on page 901.
- A NAC filter in which you specify authentication results from a RADIUS authentication server group as filter criteria. For information, see [About NAC Filters](#) on page 941.

When the appliance receives an address request, it checks if the request matches a filter. If it does not, the appliance assigns an address from the address range with the highest available IP address. If the request matches at least one class filter for a range, the appliance applies the following rules:

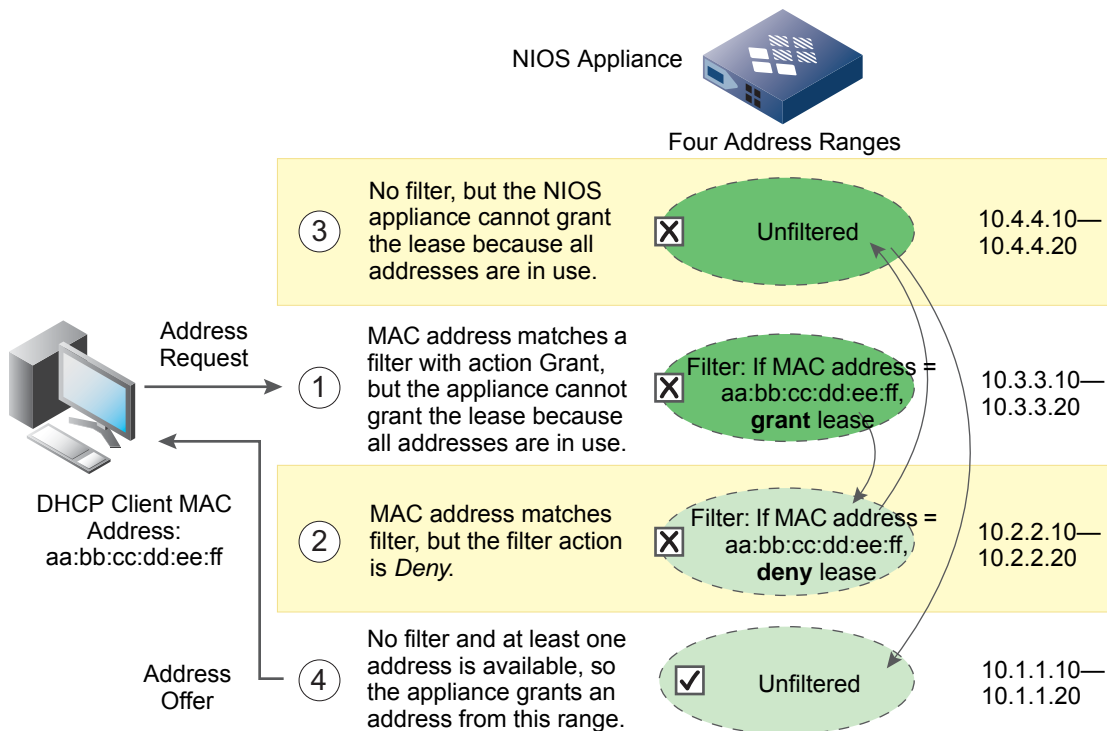
- If there are grant address filters applied to that range, the request must match one of the class filters or the appliance does not grant an address from that range.
- If there are deny address filters applied to that range, the request must not match any of the filters. If the request matches a deny filter, the appliance does not grant an address from that range.
- If an address range has a combination of grant and deny filters, the request must:
 - Match a grant filter
 - Not match a deny filter

Two rules govern the behavior of the appliance in relation to DHCP filters:

1. Depending on your filter configuration, the appliance checks if any data in an address request (such as the MAC address of the client, DHCP options 77 and 82, etc.) matches any filters applied to an address range.
2. The appliance checks for available addresses in the address ranges containing the highest addresses first. (“Highest” means closest to 255.255.255.255, and “lowest” means closest to 0.0.0.0.)

These two rules can work in coordination. For example, when the appliance receives an address request, it first checks if the request matches any filter. If it matches more than one filter assigned to different address ranges, the appliance first applies the filter that belongs to the range with the highest IP addresses. If that address does not grant an address lease (because the filter action is Deny or all address leases in that range are already in use), the appliance then applies the matching filter for the range with the next higher set of IP addresses. If the appliance still has not granted a lease from the address ranges whose filters match data in the request and there are unfiltered address ranges, the appliance attempts to assign an address from one of these ranges, again beginning with the range having the highest IP addresses. [Figure 29.4](#) presents an example illustrating the sequence in which the appliance assigns addresses when a request matches a MAC address filter. For information about MAC address filters, see [About MAC Address Filters](#) on page 897.

Figure 29.4 DHCP Address Assignment with Multiple Filters



The following explains how the NIOS appliance applies filters to DHCP address requests:

If	then
the appliance receives a request that matches a filter for one address range,	it applies the action specified in the filter for that address range. If it does not assign an address from that range (the action is <i>deny</i> or the action is <i>grant</i> but all addresses in that range are in use), the appliance then checks if it can assign an address from an unfiltered address range (if there are any), starting with the range with the highest addresses first, as shown in Figure 29.3 on page 895.
the same filter applies to multiple address ranges and the appliance receives an address request matching that filter,	it checks the address range with the highest IP addresses matching that filter. If the appliance does not assign an address from that range, it checks the filtered address range with the next highest IP addresses, and so on. If it still has not assigned an address, the appliance starts checking unfiltered address ranges (if there are any), again beginning with the range with the highest address first.
multiple filters for the same address range conflict with each other (one filter grants a lease and another denies it) and a requesting client matches both filters,	the filter denying the lease takes precedence. For example, if a requesting client matches both a MAC address filter (granting a lease) and a user class filter (denying a lease) for the same address range, the appliance denies the lease. When faced with a choice to either allow or deny a lease based on equal but contradictory filters, the appliance takes the more secure stance of denying it.

ABOUT MAC ADDRESS FILTERS

The appliance can filter an address request by the MAC address of a requesting host. Depending on how you apply the MAC filter, the appliance can grant or deny the address request if the requesting host matches the filter criteria. You can also define DHCP options that you want to return to the matching client if the options are so configured. The client can also request specific options to be returned through DHCP option 55. The appliance returns DHCP options to matching clients based on how you apply the filters. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

You can configure a MAC address filter or specific MAC addresses within a filter to expire after a certain amount of time has passed. Filter expiration is useful in situations where you want to keep filters running against updated MAC addresses. The permission to use the MAC addresses assigned to an IP address may become invalid after a certain period of time. For example, you can use a MAC address filter to restrict the right to use MAC addresses assigned to IP addresses for visiting guests or temporary workers. You can avoid removing invalid addresses from address filters manually by configuring the appliance to expire filters or to expire specific addresses within filters.

To apply a MAC address filter to an address range:

1. Define a MAC address filter. For information, see [Defining MAC Address Filters](#).
2. Add a MAC address to the filter. For information, see [Adding MAC Address Filter Items](#) on page 899.
3. Apply the filter to a DHCP address range or range template, and specify that if the MAC address of a requesting host matches the filter definition, the appliance either grants or denies the address assignment. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

Defining MAC Address Filters

To define a MAC address filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 MAC Address Filter**.
or
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 MAC Address Filter**.
2. In the *Add IPv4 MAC Filter* wizard, complete the following:
 - **Name:** Enter a meaningful name for the filter. For example, if you want to filter address requests by department, you can name one filter “Marketing”, another “Finance”, and so on. The name must be unique within a specific network. If you want to specify option settings in the filter, the filter name must be unique among all MAC filters.
 - **Comment:** Enter useful information about the filter.
3. Click **Next** and complete the following to define the DHCP options to return to the matching client:
 - **Option Space:** Select an option space from the drop-down list. This field is displayed only when you have custom option spaces. The appliance uses the **DHCP** option space as the default.
 - **Lease Time:** Enter the value of the lease time in the field and select the time unit from the drop-down list. The lease time applies to hosts that meet the filter criteria.

Options to Merge with Object Options

Click the Add icon. Grid Manager adds a new row to the table with the default **DHCP** option space and option name displayed. Complete the following:

- **Option Space:** Click the down arrow and select an option space from the drop-down list. The selected option space contains the corresponding DHCP options.
- **Option Name:** Click the down arrow and from the drop-down list, select the DHCP option you want to return to the requesting host.
- **Value:** Enter the value that you want the filter to return for the selected DHCP option. For example, enter the value 255.255.255.0 for the subnet-mask option.

To add more options to the filter, click the Add icon and repeat the steps.

4. Click **Next** and complete the following to configure the expiration setting:
 - **Default MAC Address Expiration**
Select one of the following to configure the expiration setting for the filter:
 - **Never Expires:** Select this if you want the MAC address filter to never expire. This is selected by default.
 - **Automatically Expires in:** Select this if you want the filter to expire after a specific time frame. You can specify the time in seconds, minutes, hours, or days.
The filter expiration time you configure here affects how long the DHCP server grants a lease to a client. It has an upper limit of 15 minutes on the lease time you configure for the Grid. For example, if both the filter expiration time and the lease time are less than 15 minutes, the appliance uses the lease time. If both the filter expiration time and lease time are greater than 15 minutes, the appliance uses the filter expiration time. If the filter expiration time is less than 15 minutes and the lease time is greater than 15 minutes, the DHCP server grants a lease for 15 minutes. If the filter expiration time is greater than 15 minutes and the lease time is less than 15 minutes, the appliance uses the lease time.
 - **Enforce Expiration Times :** Select this to enable the expiration setting.
 - **Enabled:** The filter is enabled by default. Clear the check box to disable this filter.
5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

Adding MAC Address Filter Items

To add a MAC address to a MAC address filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 MAC Address Filter Item**.
or
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 MAC Address Filter Item**.
2. In the *Add IPv4 MAC Address Filter Item* wizard, complete the following:
 - **MAC Address Filter:** Click **Select Filter**. In the *DHCP Filter Selector* dialog box, select the MAC address filter to which you want to add a MAC address, and then click the Select icon. If you are adding a MAC address to a filter that you have selected in the *Filters* panel, Grid Manager displays the selected filter in this field.
 - **MAC Address:** Enter the MAC address in one of the following formats: aa:bb:cc:dd:ee:ff, aa-bb-cc-dd-ee-ff, aabb.ccdd.eeff, aabbcc-ddeeff, and aabbccddeeff. The appliance displays the address in the AA:BB:CC:DD:EE:FF format. You can also enter a vendor prefix in the three hexadecimal format using the same separators supported in the MAC address format. For example, you can enter aa.bb.cc as the vendor prefix. The appliance displays AA:BB:CC.
 - **Comment:** Enter useful information about the filter item.
 - **Expiration Time**
MAC addresses in a filter stay valid until you explicitly configure them to expire. You can enable expiration for specific MAC addresses in the filter. Select one of the following:
 - **Never Expires:** Select this if you want the MAC address to never expire. This is selected by default.
 - **Expires on:** Select this and specify the **Date** and **Time** for the expiration. The fields display the current date and time. If you have already configured an expiration time for the filter, the appliance displays the time here by adding the filter expiration time to the current time. For example, if the expiration time for the filter is two days and the current date is June 6, 2009, the appliance displays June 8, 2009 in the **Date** field.
3. Click **Next** and select one of the following to configure user registration (optional):
 - **Register as User:** Select this and enter a username in the field.
 - **Register as Guest:** Select this and enter the first name, middle name, last name, email address, and phone number of the guest user.

The appliance displays the information you enter here in the lease viewers.
4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

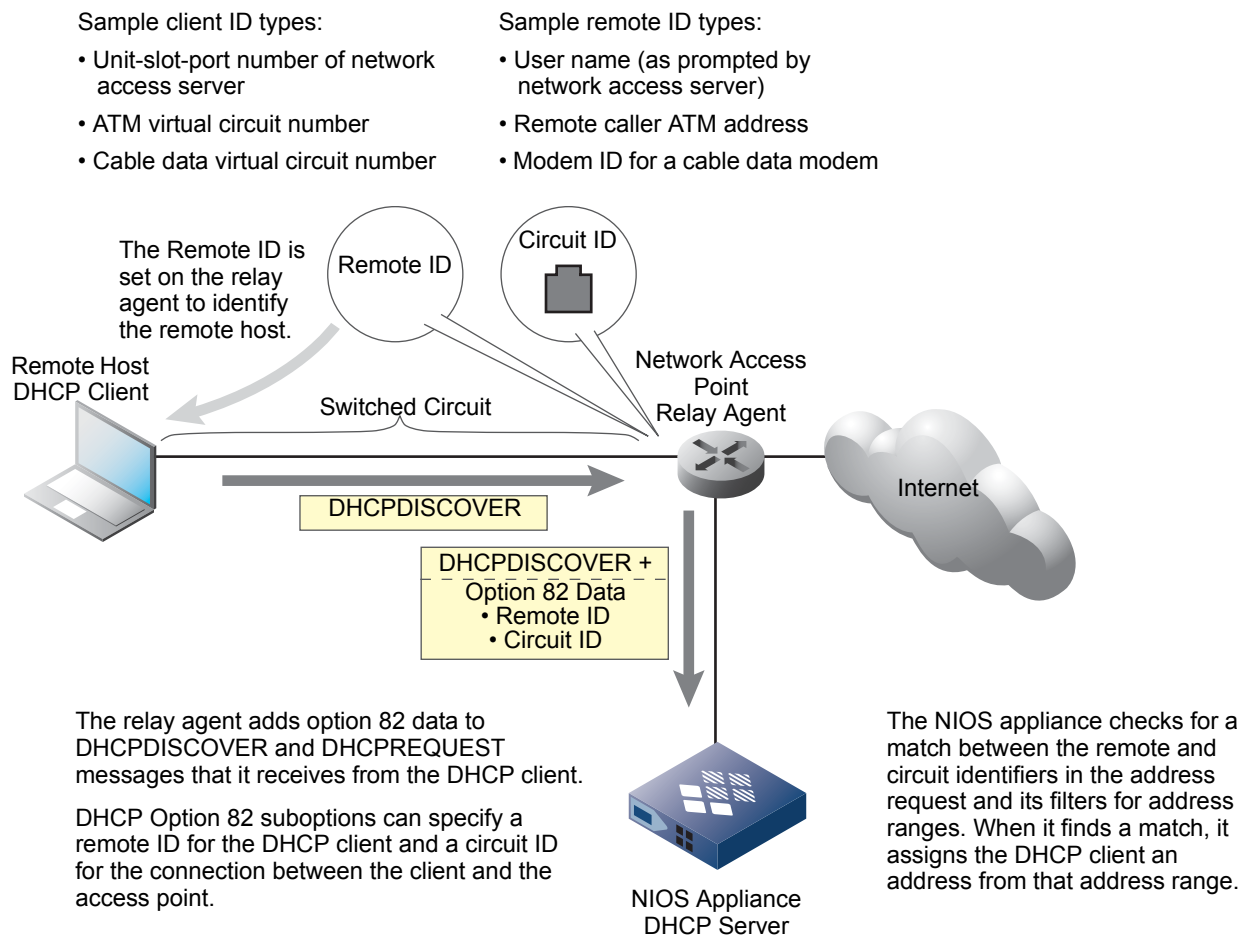
After you define a MAC address filter and add MAC addresses to it, you can assign the filter to a DHCP range. The appliance filters IP address requests based on the filter criteria. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

ABOUT RELAY AGENT FILTERS

The NIOS appliance can filter an address request by the circuit ID and remote ID of a requesting host. The filter instructs the appliance either to grant or deny an address request if the requesting host matches the filter. For information about the DHCP relay agent option, see [About the DHCP Relay Agent Option \(Option 82\)](#) on page 806.

Option 82 assists the agent in forwarding address assignments across the proper circuit. When a relay agent receives a DHCPDISCOVER message, it can add one or two agent IDs (circuit ID and remote ID) in the DHCP option 82 suboption fields to the message, as illustrated in [Figure 29.5](#). If the agent ID strings match those defined in a relay agent filter applied to a DHCP address range, the appliance either assigns addresses from that range or denies the request based on the configured parameters.

Figure 29.5 Relay Agent Filtering



To apply a relay agent filter to an address range:

1. Define a relay agent filter. For information, see [Defining Relay Agent Filters](#).
2. Apply the filter to a DHCP address range or range template, and specify that if the circuit ID or remote ID of a requesting host matches the filter definition, the appliance either grants or denies the address assignment. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.
3. Define the access privileges of limited-access admin group for relay agent filters. For information, see [Managing Administrators](#) on page 149.

Defining Relay Agent Filters

To define a relay agent filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 Relay Agent Filter**.
or
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 Relay Agent Filter**.
2. In the *Add IPv4 Relay Agent Filter* wizard, complete the following:
 - **Name:** Enter a meaningful name for the filter. For example, you can enter the IP address or the name of the router acting as the relay agent. The name must be unique within a specific network.
 - **Comment:** Enter useful information about the filter.

3. Click **Next** to define the relay agent ID type. If you apply both ID types, the relay agent must provide both identifiers when submitting a DHCP address request.
 - Select one of the following for both **Circuit ID** and **Remote ID**:
 - **Any**: Select this and the filter matches any of the circuit identifiers for remote hosts. You cannot select this for both circuit ID and remote ID at the same time.
 - **Not Set**: Select this and no circuit identifier is set for remote hosts.
 - **Matches Values**: Select this and enter the circuit ID or remote ID. You can enter the ID in hexadecimal format, such as 1f:cd, ab, or ef:23:56, or in string format, such as abcd or aa:gg. The appliance matches the value you enter here with the value sent by the DHCP client in counted octet sequence format.
4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

After you define a relay agent filter, you can assign it to a DHCP range. The appliance responds to address requests based on the filter criteria. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

ABOUT OPTION FILTERS

You can use option filters to classify DHCP clients and decide which DHCP options each group of clients can receive. By default, regardless of the networks in which the DHCP clients reside and whether an option filter is applied to a DHCP range or range template, all DHCP clients that match the filter criteria receive the DHCP options and values you define in the filter. You can change this configuration so the appliance does not use the filter to classify DHCP clients. For information about how to configure this, see [Defining Option Filters](#) on page 902.

You can add DHCP options and the Hardware Operator option to an option filter. (For information about the Hardware Operator option, see [DHCP Hardware Operator](#).) Depending on whether the options you add to the filter are also defined at the Grid, member, network, and DHCP range levels, and whether you add the filter to the Class Filter List or Logic Filter List of a range or range template, the appliance either appends them to the existing options or overwrites the option values before returning them to the matching clients. For more information about how the appliance returns DHCP options, see [Adding Filters to the Logic Filter List](#) on page 908.

The appliance can filter an address request by the options (such as root-server-ip-address or user-class) of the requesting host. Depending on how you apply an option filter, the appliance can grant or deny an address request if the requesting host matches the filter criteria. You can also create complex match rules that use the AND and OR logic to further define the filter criteria. When you select match rules in Grid Manager, you can preview the rules before committing them to the filter. Grid Manager provides an expression builder that automatically builds the rules after you define them. For information, see [Defining Option Filters](#) on page 902.

To define an option filter and apply it to an address range:

1. Define an option filter based on either the predefined or custom DHCP options. For information, see [Defining Option Filters](#).
2. Apply the filter to a DHCP address range or range template in the Class Filter List or Logic Filter List. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

After you define an option space and add options to it, you can set up option filters and define option values. For example, to handle two different client classes, you can define two option filters (vendor-class_1 and vendor-class_2) and send different option values to different clients based on the vendor-class-identifier options that you obtain from the clients.

DHCP Hardware Operator

You can define the Hardware Operator option and add it as a match rule to an option filter. This option enables the appliance to match the hardware type and MAC address of the DHCP client, which it derives from the hardware type, hlen (hardware length) and chaddr (client hardware address) fields of the client's DHCP Discover and Renew packets.

To add Hardware Operator to an option filter, fill in the fields as follows:

- In the first drop-down list, select **Hardware Operator**. Note that because it is not a DHCP Option, it does not have an actual option number.
- In the second drop-down list, select one of the following operators: **equals**, **does not equal**, **substring equals** and **substring does not equal**.
If the operator is **substring equals** or **substring does not equal**, specify the offset and length.
- In the text field, enter the string that represents the hardware type and MAC address to match. For example, the htype value is 1 for the Ethernet hardware type. The hardware types (hrd) are defined at <http://www.iana.org/assignments/arp-parameters/arp-parameters.xml#hardware-type-rules>.

This filter rule assumes that the values exist in the DHCP packets.

The following table provides examples of rules that include the Hardware Operator option. The entry in the first drop-down list for all rules is **Hardware Operator**.

Table 29.1 Hardware Operator Sample Rules

Rule Description	Second Drop-Down List (operator)	Text Field (string)	Offset	Length
Match a hardware type and MAC address.	equals	01:00:C0:B0:AA:BB:CC		
Match hardware type only.	substring equals	01	0	1
Match the vendor MAC prefix (first three bytes of MAC address).	substring equals	00:C0:B0	1	3

Defining Option Filters

To define an option filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 Option Filter**.
or
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 Option Filter**.
2. In the *Add IPv4 Option Filter* wizard, complete the following:
 - **Name:** Enter a meaningful name for the option filter. For example, you can enter Linux if you plan to use this option filter to screen Linux systems. The name must be unique within a specific network. If you want to specify option settings in the filter, the filter name must be unique among all option filters.
 - **Comment:** Enter useful information about the filter.
 - **Apply this filter as a global DHCP class:** This check box is selected by default. When you select this check box, the appliance defines a global class statement in the dhcpd configuration file for members that have DHCP enabled, regardless of whether the filter is applied to a DHCP range or range template. All DHCP clients that belong to this class receive the DHCP options and values you define in the filter. When you clear this check box, you cannot apply this filter to the Class Filter List of a range or range template. You cannot clear this check box if the filter is currently applied to a range or range template. The appliance displays an error message when you try to save this configuration.
3. Click **Next** and complete the following to add match rules:
 - In the first drop-down list, select a DHCP option. For example, select **user-class (77)** for a specific user class, such as mobile users.
 - In the second drop-down list, select an operator.
If you select **equals** or **does not equal**, enter the value of the selected option you want the filter to match in the field.

If your operator and match value include a substring of an option value, enter the offset and length of the substring based on the following definitions:

- **Offset:** Enter the number of characters at which the match value substring starts in the option data. Enter 0 to start at the beginning of the option data, enter 1 for the second position, and so on. For example, when you enter 2 and have a match value of MSFT, the appliance matches the value MSFT starting at the third character of the option data.
- **Length:** Enter the length of the match value. For example, if the match value is MSFT, the length is 4.

You can do the following and repeat the filter selection steps to add another rule:

- Click **+** to add another rule at the same level.
- Click **|&-** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level and above the first rule.
- Click **->|** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level.

After you add all the match rules, you can click **Preview** to view the rules that are written to the dhcpd configuration file or click **Reset** to remove the previously configured rules and start again. For information about how to use match rules, see [Using Match Rules in Option Filters](#) on page 903.

4. Click **Next** and complete the following to define which DHCP options to return to the matching client:
 - **Option Space:** Select an option space from the drop-down list. This field is not displayed if you do not have custom option spaces. The appliance uses the **DHCP** option space as the default.
 - **Lease Time:** Enter the value of the lease time in the field and select the time unit from the drop-down list. The lease time applies to hosts that meet the filter criteria.

Options to Merge with Object Options

Click the Add icon. Grid Manager adds a new row to the table with the default **DHCP** option space and option name displayed. Complete the following:

- **Option Space:** Click the down arrow and select an option space from the drop-down list. The selected option space contains the corresponding DHCP options that you can use as filter criteria.
- **Option Name:** Click the down arrow and from the drop-down list, select the DHCP option you want to use as filter criteria.
- **Value:** Enter the match value that you want the filter to use for the selected DHCP option.

To add more options to the filter, click the Add icon and repeat the steps.

5. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

Using Match Rules in Option Filters

Each match rule you define in an option filter further defines the filter criteria of a matching client. You can add multiple match rules to an option filter. The appliance writes these rules to the dhcpd configuration file. You can also create complex match rules that use the AND and OR logic to further define the filter criteria. After you define the match rules, you can preview the rules before committing them to the filters.

For example, you can define the following rules in an option filter:

DHCP option = vendor-class-identifier

Substring offset = 0 (the match value starts at the beginning of the option data received from the client)

Substring length = 4 (the length of the match value MSFT)

Match value = MSFT

The appliance generates the following rules in the dhcpd configuration file:


```

class "microsoft-other" {
  match if substring (option vendor-class-identifier,
0, 4) = "MSFT";
  vendor-option-space MSFT;
}

```

Substring offset

Match value

Length of the match value

DHCP option

You can also define more complex rules using the AND and OR logic as follows:

DHCP option = vendor-class-identifier

Match value = infoblox2000a

OR

DHCP option = vendor-encapsulated-options

Substring offset = 0 (the match value starts at the first character of the option data received from the client)

Substring length = 8 (the length of the match value infoblox)

Match value = infoblox

AND

DHCP option = vendor-encapsulated-options

Substring offset = 10 (the match value starts at the ninth character of the option data received from the client)

Substring length = 5, the length of the match value 2000a

Match value = 2000a

The appliance generates the following rules in the dhcpd configuration file:

```

class "infoblox" {
  match if (option vendor-class-identifier=infoblox2000a:) or
((substring(option vendor-encapsulated-options,0,8)="infoblox") and
(substring(option vendor-encapsulated-options,10,5)="2000a"));
  vendor-option-space DHCP
}

```

Configuring User Class Filters

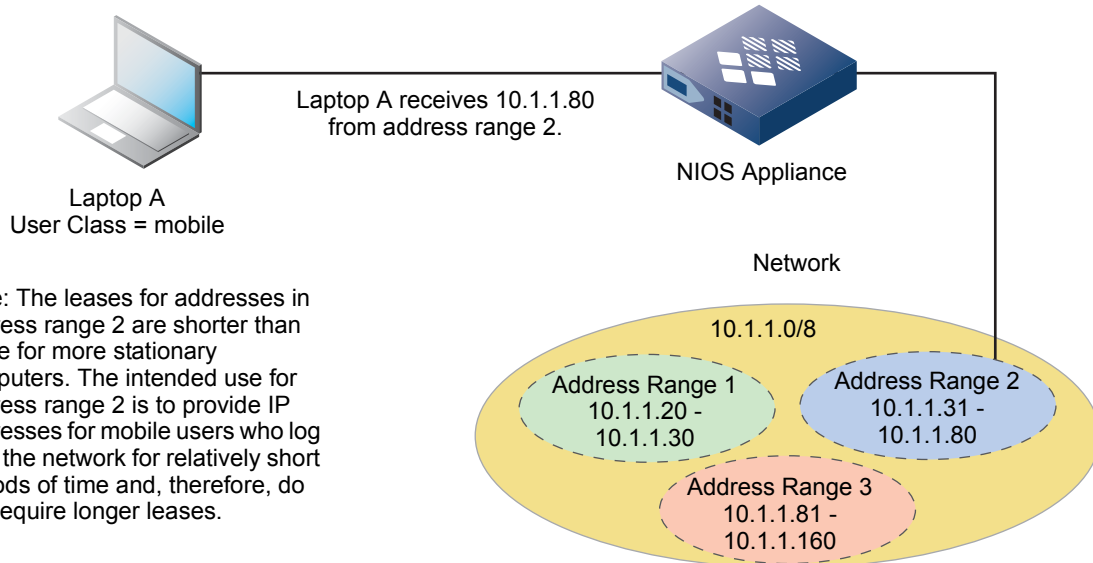
The NIOS appliance can filter DHCP address requests by user class filters. A user class indicates a category of user, application, or device of which the DHCP client is a member. User class identifiers are configured on DHCP clients and are sent during a DHCP address request operation. The client includes the user class identifier in DHCP option 77 when sending DHCPDISCOVER and DHCPREQUEST messages.

By using user class identifiers, a DHCP server can screen address requests and assign addresses from select address ranges based on the different user class identifiers it receives. For example, if you assign a user class filter named mobile to a range of addresses from 10.1.1.31–10.1.1.80, the appliance selects an address from that range if it receives an address request that includes the user class name mobile and there are still addresses available in that range. You might want mobile users to receive these addresses because you have given them shorter lease times than other, more stationary DHCP clients. See [Figure 29.6](#).

Figure 29.6 Applying User Class Filtering

The user class for laptop A is *mobile*. When it sends DHCPDISCOVER and DHCPREQUEST messages, it includes its user class in the DHCP option 77 field.

The NIOS appliance has a filter that screens address requests by user class. If the user class for a DHCP client is *mobile*, the appliance assigns it an address from address range 2.



Note: The leases for addresses in address range 2 are shorter than those for more stationary computers. The intended use for address range 2 is to provide IP addresses for mobile users who log in to the network for relatively short periods of time and, therefore, do not require longer leases.

If the NIOS appliance receives address requests with the user class *mobile* and there are no available addresses in address range 2 but there are available addresses in ranges 1 and 3, the appliance begins assigning addresses from address range 3 (because its addresses are higher than those in range 1). Then, if all addresses in range 3 are in use, the appliance begins assigning addresses from address range 1. If you want the appliance to assign addresses to mobile users (that is, those identified with the user class *mobile*) exclusively from address range 2, then you must apply user class filters for “mobile” to address ranges 1 and 3 that deny lease requests matching that user class.

Configuration Example: Using Option Filters

The following example shows you how to create an option space, add custom options to it, create an option filter, and a match rule to filter the options so that the NIOS appliance can filter an address request by the vendor options of the requesting hosts. It can grant or deny an address request if the requesting host matches the filter.

1. Add an option space called MSFT, and then add the following options to it. For information, see [Applying DHCP Options](#) on page 804.

Option name	Code	Type
root-mount-options	1	Text
root-server-ip-address	2	IP address
root-server-host-name	3	Text
root-server-path-name	4	Text
swap-server-ip-address	5	IP address
swap-file-path-name	6	Text
boot-file-path-name	7	Text
posix-timezone-string	8	String

Option name	Code	Type
boot-read-size	9	16-Bit unsigned integer

- From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab and click the Add icon.
- In the *Add IPv4 Filter* wizard, enter the filter name **i86pc**, and then select **Options** as the filter type.
- Select **MSFT** as the option space, select an option, specify a value for it, and then add it to the i86pc option filter. You can select multiple options. Add the following options to the i86pc option filter:

Option name	Code	Type
root-server-ip-address	2	IP address
root-server-host-name	3	Text
root-server-path-name	4	Text
boot-file-path-name	7	Text

- From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab -> *filter_name*, and then click the Add icon.
- In the *Add IPv4 Match Rule* wizard, select **i86pc** as the option filter, select **vendor-class-identifier (60)** as the matching option, and then enter **MSFT** as the matching value.
- Add a DHCP range to the network. For information, see [Configuring IPv4 Address Ranges](#) on page 854.
- Apply the i86pc option filter to the DHCP address range. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.
- Click **Restart** to restart services.

ABOUT DHCP FINGERPRINT FILTERS

The appliance can filter an address request by the DHCP fingerprint of a requesting client. Depending on how you apply DHCP fingerprint filters, the appliance can grant or deny the address request if the requesting client matches the filter criteria. Note that only superusers can add, modify, and delete DHCP fingerprint filters. Limited-access users cannot perform any DHCP fingerprint filter related tasks, though with the correct permissions they can apply DHCP fingerprint filters to DHCP ranges and range templates. For information about how to apply filters to DHCP ranges, see [Applying Filters to DHCP Address Ranges](#) on page 907.

You can define a DHCP fingerprint filter by selecting one or multiple DHCP fingerprints from the existing list of DHCP fingerprints, and then assign a grant or deny permission to the filter. You can then apply the filter to a DHCP address range, if DHCP fingerprint detection is enabled. For information about how to enable DHCP fingerprint detection, see [Enabling and Disabling DHCP Fingerprint Detection](#) on page 1034.

Note that once you apply a DHCP fingerprint filter to an address range, you cannot disable DHCP fingerprint detection or disable individual DHCP fingerprints that have been included in the filter. You must first delete or disable the DHCP fingerprint filter that you have applied to the address range before you can disable any fingerprint related tasks. For information about how to delete a DHCP fingerprint filter, see [Deleting Filters](#) on page 914.

On lease renewals, requesting clients must send the same DHCP fingerprint information in order for the appliance to properly grant or deny leases based on the configured DHCP fingerprint filters. For example, if a client sends option 55 in the original request but does not send the same information in the renewal request, and you have configured a DHCP fingerprint filter to grant a lease to this client, the appliance may not be able to properly grant a lease to this client.

To apply a DHCP fingerprint filter to an address range:

- Define a DHCP fingerprint filter. For information, see [Defining DHCP Fingerprint Filters](#).

2. Apply the filter to a DHCP address range or range template, and specify that if the DHCP fingerprint of a requesting host matches the filter definition, the appliance either grants or denies the address assignment. For information, see [Applying Filters to DHCP Address Ranges](#) on page 907.

Defining DHCP Fingerprint Filters

To define a DHCP fingerprint filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 Fingerprint Filter**.
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 Fingerprint Filter**.
2. In the *Add IPv4 Fingerprint Filter* wizard, complete the following:
 - **Name:** Enter a meaningful name for the filter. For example, if you want to filter address requests by a specific device class, you can name one filter “Gaming Console,” another “Android Phones,” and so on. The filter name must be unique among all DHCP fingerprint filters.
 - **Comment:** Enter useful information about the filter.
3. Click **Next** and then click the Add icon in the Select Fingerprints table. In the *Fingerprint Selector* dialog box, select the DHCP fingerprint you want to include in this filter. Click Add icon to select another DHCP fingerprint.

Note: When you select **No Match**, the appliance applies the filter to all requesting clients that do not send option 55 and option 60 or to clients that send values in option 55 and 60 that do not match any existing DHCP fingerprints.

4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
5. Save the configuration.

APPLYING FILTERS TO DHCP ADDRESS RANGES

To further control how the appliance allocates IPv4 addresses to DHCP client requests, you can apply DHCP filters to an address range or range template so the appliance can determine the following:

- The class statements
- The address ranges from which it assigns leases
- When to grant or deny leases to the matching clients
- Which DHCP options to return to the matching clients

Adding Filters to the Class Filter List

You can apply any DHCP filter to the Class Filter List of a DHCP range or range template. The appliance uses the matching rules of these filters to select the address range from which it assigns a lease. You can define permissions for these filters to instruct the appliance whether to grant or deny a lease to the matching client. When you add a filter with a grant permission, the client must match the filter criteria to receive a lease. When you define a filter with a deny permission, clients that do not match the filter criteria still receive leases. Only the client that matches the filter criteria is denied a lease.

Filters in the Class Filter List correspond to the class statement generated in the `dhcpd` configuration file, which is a classification of the client packet. All DHCP clients that match the option filter and relay agent filter criteria become members of the same class and are eligible to receive DHCP options for that class, regardless of the networks in which the clients reside. However, a client can only become a member of the MAC or NAC filter class when it is granted a lease from the DHCP range based on the filter criteria. Whether a client receives specific options and option values depends on the hierarchy of the options and how you apply the filters. For information about how the appliance returns DHCP options, see [Adding Filters to the Logic Filter List](#).

Adding Filters to the Logic Filter List

The filters you add to the Logic Filter List correspond to the match rules that are written to the `dhcpcd` configuration file. The appliance uses these filters to identify DHCP options and values to return to the matching clients. You can apply option, MAC, and NAC filters to the Logic Filter List. Note that a DHCP client is eligible to receive DHCP options defined in a filter if it matches the filter criteria. Whether the client receives specific options and their corresponding values depends on the hierarchy of the options and the list of options requested by the client through DHCP option 55. You can configure the appliance to ignore the option list requested by a matching client and return all the options that the client is eligible to receive. For information about how to ignore the option list requested by a client, see [Configuring General IPv4 DHCP Properties](#) on page 793.

The appliance decides which options and values to return to a client based on the following:

- If you have different DHCP options defined in a range and any DHCP filters in the Class Filter and Logic Filter lists, and these options do not overlap, the appliance merges and returns all options to the matching client. For example, a DHCP client obtains a lease from a DHCP address range (R) through an option filter in the Class Filter List (CF), which contains an option statement (O1) with a value of (S1). The appliance then matches a filter in the Logic Filter List (LF) that contains an option statement (O2) with a value of (S2). In this case, option statements O1 and O2 and their values S1 and S2 are merged and returned to the matching client.
- If there are overlapping DHCP options in a range and any DHCP filters in the Class Filter and Logic Filter lists, the values defined in the Class Filter List filters take precedence over those defined in the range and filters in the Logic Filter List. The appliance returns the option value defined in the class filters to the matching client. For example, a DHCP client obtains a lease from a DHCP address range (R) through an option filter in the Class Filter List (CF), which contains an option statement (O1) with a value of (S1). The appliance then matches a filter in the Logic Filter List (LF) that contains the same option statement (O1) with a value of (S2). In this case, the option value S1 defined in the option filter in the Class Filter List takes precedence and is returned to the DHCP client.
- When you apply option, MAC, and NAC filters to the Logic Filter List, the appliance translates their match rules into a DHCP if/elseif/else statement using the match rules of the first filter on the list as the “if” expression in the statement. Match rules in subsequent filters are translated into the “elseif” statements, and the last filter that does not contain any match rules is translated into the “else” statement. Note that a filter without any match rules can only be added as the last filter in the Logic Filter List.

For more information about how the appliance grants and denies leases to requesting clients and determines which DHCP options to return to the matching clients, see [Configuration Example: Using the Class and Logic Filter Lists](#) on page 909.

To apply filters to a DHCP address range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.
2. In the *DHCP Range* editor, click **Toggle Advanced Mode**, and then select the **IPv4 Filters** tab.
3. **Class Filter List:** Click the Add icon to add a filter to identify the class of a matching client, and to grant or deny a lease to a client. For more information, see [Adding Filters to the Logic Filter List](#) on page 908.

If you have only one configured DHCP filter, the appliance displays the filter in the table. Otherwise, in the *DHCP Filter Selector* dialog box, click the filter you want to add. Use SHIFT+click and CTRL+click to select multiple filters.

For each filter you add, click the **Action** column and select one of the following from the drop-down list:

— **Grant lease:**

For MAC address filters: Select this to assign an IP address from the address range to a requesting host whose MAC address matches the MAC address in the filter.

For relay agent filters: Select this to assign an IP address from the address range when one or both of the relay agent identifiers of the requesting host match the filter criteria.

For option filters: Select this to assign an IP address from the address range to a requesting host whose DHCP options match the DHCP options and match rules defined in the filter.

For NAC filters: Select this to assign an IP address from the address range to a requesting host based on the authentication results from a RADIUS authentication server group.

For DHCP fingerprint filters: Select this to grant a lease from the address range to a requesting host based whose DHCP fingerprint matches the DHCP fingerprint in the filter.

— **Deny lease:**

For MAC address filters: Select this to deny an address request from a host whose MAC address matches a MAC address in the filter.

For relay agent filters: Select this to deny an address request when one or both relay agent identifiers match the filter criteria in the filter.

For option filters: Select this to deny an address request from a host whose DHCP options match the options and match rules in the filter.

For NAC filters: Select this to deny an address request from a host based on the authentication results from a RADIUS authentication server group.

For DHCP fingerprint filters: Select this to deny a lease request when the DHCP fingerprint of the requesting host matches the DHCP fingerprint in the filter.

4. **Logic Filter List:** Click the Add icon to add a filter to match a client based on the match rules defined in the filter. The appliance uses filters in both the Class Filter and Logic Filter lists to determine the DHCP options and values it returns to the matching clients. For more information, see [Adding Filters to the Logic Filter List](#) on page 908. If you have only one configured DHCP filter, the appliance displays the filter in the table. Otherwise, in the *DHCP Filter Selector* dialog box, click the filter you want to add. Use SHIFT+click and CTRL+click to select multiple filters.

Note: You can only add a filter that does not contain any match rules as the last filter in the Logic Filter List.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuration Example: Using the Class and Logic Filter Lists

The following example shows you how to define DHCP filters and apply them to the class and logic filter lists. It also shows you the DHCP configuration file that is generated based on the configuration.

In this example, you first define a MAC filter, two option filters (one without match rules), and a NAC filter, and then apply the MAC filter to the Class Filter List and the other filters to the Logic Filter List of the address range 10.34.34.6 - 10.34.34.55.

1. Configure and save a MAC filter as follows. For more information, see [Defining MAC Address Filters](#) on page 898.
 - a. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 MAC Address Filter**.
 - b. In the *Add IPv4 MAC Filter* wizard, complete the following:
 - **Name:** Enter **MAC1**.
 - c. Click **Next** and complete the following to define the DHCP options to return to the matching client:
 - **Lease Time:** Enter **1234** and select **seconds** from the drop-down list.

Options to Merge with Object Options: Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:

 - **Option Name:** Click the down arrow and select **log-server (7)** from the drop-down list.
 - **Value:** Enter **10.34.34.3** as the value for the log-server option that is sent to the client in the OFFER/ACK message.
 - d. Save the configuration.
2. Add a MAC address filter item as follows. For more information, see [Adding MAC Address Filter Items](#) on page 899.
 - a. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 MAC Address Filter Item**.
 - b. In the *Add IPv4 MAC Address Filter Item* wizard, complete the following:

- **MAC Address Filter:** Click **Select Filter**. In the *DHCP Filter Selector* dialog box, click **MAC1**.
 - **MAC Address:** Enter **AB:DE:CC:DD:EE:01** as the MAC address.
- c. Save the configuration.
- 3. Configure and save an option filter with match rules as follows. For more information, see [Defining Option Filters](#) on page 902.
 - a. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 Option Filter**.
 - b. In the *Add IPv4 Option Filter* wizard, complete the following:
 - **Name:** Enter **option1**.
 - c. Click **Next** and complete the following to add match rules:
 - In the first drop-down list, select **vendor-class-identifier**.
 - In the second drop-down list, select **substring equals**, and then enter the following:
 - **Offset:** Enter 0 to match the value starting at the first character of the option data.
 - **Length:** Enter 4.
 - Enter **MSFT** as the matching value.

Click **Preview** and the appliance displays the expression: `(vendor-class-identifier , 0 , 4 = "MSFT")`.
 - d. Click **Next** and complete the following to define the DHCP options to return to the matching client:

Options to Merge with Object Options: Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:

 - **Option Name:** Click the down arrow and from the drop-down list, select **time-server(4)**.
 - **Value:** Enter **10.34.34.2** as the value for the time-server option that is sent to the client in the OFFER/ACK message.
 - e. Save the configuration.
- 4. Configure and save another option filter without match rules as follows:
 - a. In the *Add IPv4 Option Filter* wizard, complete the following:
 - **Name:** Enter **option2**.
 - b. Click **Next**. Do not define any match rules.
 - c. Click **Next** again and complete the following to define the DHCP options to return to the matching client:

Options to Merge with Object Options: Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:

 - **Option Name:** Click the down arrow and from the drop-down list, select **domain-name(6)**.
 - **Value:** Enter **www.infoblox.com**.
 - d. Save the configuration.
- 5. Configure and save a NAC filter as follows. For more information, see [Defining NAC Filters](#) on page 942.
 - a. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 NAC Filter**.
 - b. In the *Add Filter Wizard*, complete the following and click **Next**:
 - **Name:** Enter **NAC1**.
 - c. Create a rule as follows:
 - In the first drop-down list, select **Compliance State**.
 - In the second drop-down list, select **equals**.
 - In the third drop-down list, select **Compliant**.

Click **Preview** and the appliance displays the expression: `(Sophos.ComplianceState = "Compliant")`.
 - d. Click **Next** and complete the following to define DHCP options:
 - **Lease Time:** Enter 1000 and select **seconds** from the drop-down list.

Options to Merge with Object Options: Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:

- **Option Name:** Click the down arrow and from the drop-down list, select **cookies-servers(8)**.
- **Value:** Enter **10.34.34.5**.

e. Save the configuration.

6. Apply the filters to the address range as follows. For more information, see [Applying Filters to DHCP Address Ranges](#) on page 907.
 - a. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> **10.34.34.6-10.34.34.55** check box, and then click the Edit icon.
 - b. In the *DHCP Range* editor, click **Toggle Advanced Mode**.
 - c. Click the **IPv4 Filters** tab and complete the following:

Class Filter List: Click the Add icon and add **MAC1** as a class filter. Click the **Action** column and select **Grant lease** from the drop-down list.

Logic Filter List: Click the Add icon and add **Option1**, **NAC1**, and **Option2** respectively as logic filters.
 - d. Save the configuration.

The appliance generates the following information in the DHCP configuration file based on the filter configuration in this example:

```
# MAC filter "MAC1"
class "MAC1" {
    default-lease-time 1234;
    min-lease-time 1234;
    max-lease-time 1234;
option log-servers 10.34.34.3;
}
# NAC filter "NAC1"
{option sophos.compliance
state="compliant"
}
subnet 10.34.34.0 netmask 255.255.255.0 {
    pool {
        infoblox-range 10.34.34.6 10.34.34.55;
        range 10.34.34.6 10.34.34.55;
        option routers 10.34.34.1;
        # INFOBLOXMACFILTERDEBUGINFO: allow members of "MAC1";

        if (substring(option vendor-class-identifier,0,4)="MSFT") {
            # Option filter "Option1"
            option time-servers 10.34.34.2;
        }
        elsif (option Sophos.ComplianceState="Compliant") {
            # NAC filter "NAC1"
            default-lease-time 1000;
            min-lease-time 1000;
            max-lease-time 1000;
            option cookie-servers 10.34.34.5;
        }
        else {
```

```
# Option filter "Option2"
default-lease-time 2500;
min-lease-time 2500;
max-lease-time 2500;
option domain-name "www.infoblox.com"; }
}
```

Depending on client requests and the matching criteria, the following scenarios can happen in this example:

If the requesting client matches the MAC1 and Option1 filters, the appliance returns the following:

- Lease time = 1234 seconds (from the MAC filter)
- Returned options:
 - Router(3) with a value of 10.34.34.1 (from the address range)
 - Log-server(7) with a value of 10.34.34.3 (from the MAC filter MAC1)
 - Time-server(4) with a value of 10.34.34.2 (from the option filter Option1)

If the requesting client matches the MAC1 and NAC1 filters, the appliance returns the following:

- Lease time = 1234 seconds (from the MAC filter MAC1)
- Returned options:
 - Router(3) with a value of 10.34.34.1 (from the address range)
 - Log-server(7) with a value of 10.34.34.3 (from the MAC filter MAC1)
 - Cookie-server(8) with a value of 10.34.34.5 (from the NAC filter NAC1)

If the client matches the MAC1 filter, but not the Option1 or NAC1 filters, the appliance returns the following:

- Lease time = 1234 seconds (from the MAC filter)
- Returned options:
 - Router(3) with a value of 10.34.34.1 (from the address range)
 - Log-server(7) with a value of 10.34.34.3 (from the MAC filter MAC1)
 - Domain-name(6) with a value of www.infoblox.com (from the option filter Option2)

If the requesting client does not match the MAC1 filter, no lease is granted.

MANAGING DHCP FILTERS

You can do the following to manage DHCP filters:

- Modify filter settings. For information, see [Modifying DHCP Filters](#).
- View a complete list of filters, MAC address items, and match rules. For information, see [Viewing DHCP Filters](#) on page 914.
- Delete filters that are not in use. For information, see [Deleting Filters](#) on page 914.

Modifying DHCP Filters

To modify a filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab -> *filter_name* check box, and then click the Edit icon.
2. For a MAC address filter, the *DHCP MAC Filter* editor provides the following tabs from which you can modify information:
 - **General:** Modify the fields as described in [Defining MAC Address Filters](#) on page 898.
 - **DHCP Options:** Add or delete DHCP options. For information, see [Defining MAC Address Filters](#) on page 898.

- **Extensible Attributes:** Add or delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
- **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#) on page 160.

For a relay agent filter, the *Relay Agent Filter* editor provides the following tabs from which you can modify information:

- **General:** Modify the fields as described in [Defining Relay Agent Filters](#) on page 900.
- **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

For an option filter, the *Option Filter* editor contains the following tabs from which you can modify information:

- **General:** Modify the fields as described in [Defining Option Filters](#) on page 902.
- **Rules:** Modify the match rules as described in [Defining Option Filters](#) on page 902.
- **DHCP Options:** Modify option spaces and DHCP options in the **Basic** tab as described in [Defining Option Filters](#) on page 902. You must define the **PXE Lease Time** in the **Advanced** tab.
- **BOOTP:** Modify BOOTP settings as described in [Configuring IPv4 BOOTP and PXE Properties](#) on page 798.
- **Extensible Attributes:** Add or delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

For a DHCP fingerprint filter, the *Add IPv4 Fingerprint Filter* editor provides the following tabs from which you can modify information:

- **General:** Modify general information, such as the name and device class, as described in [Defining DHCP Fingerprint Filters](#) on page 907.
- **Fingerprints:** Add or delete DHCP fingerprints as described in [Defining DHCP Fingerprint Filters](#) on page 907.
- **Extensible Attributes:** Add and delete extensible attributes that are associated with the DHCP fingerprint filter. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

For a NAC filter, the *NAC Filter* editor contains the following tabs from which you can modifying information:

- **General:** Modify the name and comment.
- **Rules:** Modify the rules as described in [Defining NAC Filters](#) on page 942.
- **DHCP Options:** Add or delete DHCP options. For information, see [Defining NAC Filters](#) on page 942.
- **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can modify the MAC address filter items and match rules for corresponding MAC address filters and option filters. For information, see [Modifying MAC Address Filter Items](#) on page 913 and [Viewing DHCP Filters](#) on page 914.

Modifying MAC Address Filter Items

To modify a MAC address filter item:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab -> *filter_name* -> *mac_filter* check box, and then click the Edit icon.
2. The *MAC Address Filter Item* editor contains the following tabs from which you can edit data:
 - **General:** Modify the fields as described in [Adding MAC Address Filter Items](#) on page 899.
 - **Registration:** Modify registration settings as described in [Adding MAC Address Filter Items](#) on page 899.

- **Extensible Attributes:** Add or delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Viewing DHCP Filters

To view DHCP filters:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab.
2. Grid Manager displays the following for each filter:
 - **Name:** The name of the filter.
 - **Filter Type:** The filter type.
 - **Comment:** The information about the filter.
 - **Site:** The location to which the filter belongs. This is one of the predefined extensible attributes.

Viewing MAC Address Filter Items

To view a list of MAC addresses in a specific MAC address filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab -> *filter_name*.
2. Grid Manager displays the following:
 - **MAC Address:** The DHCP fingerprint information of client's device. This field is displayed only when users use captive portal for authentication. MAC address assigned to the filter.
 - **Username:** Grid Manager displays the username to which the MAC address belongs in the lease viewers.
 - **Comment:** The information you entered about the filter item.
 - **Expiration Time:** The expiration time you configured for the MAC address.
 - **Fingerprint:** The DHCP fingerprint information of client's device. This field is updated when users use Captive Portal for authentication.
 - **Site:** The location to which the filter belongs. This is one of the predefined extensible attributes.

Deleting Filters

You can delete a filter that is not currently assigned to a DHCP range. You can also remove a filter from a DHCP range, and then delete the filter if it is not in use.

To delete a filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab -> *filter_name*, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

The appliance puts the deleted filters in the Recycle Bin, if enabled. You can later restore the filter if needed.

To schedule this task, click the Delete icon -> **Schedule Delete**. In the *Schedule Deletion* dialog box, click **Delete Later**, and then specify a date, time, and time zone.



Chapter 30 Authenticated DHCP

This chapter includes the following sections:

- [About Authenticated DHCP](#) on page 917
 - [DHCP Authentication Process](#) on page 917
- [Configuring DHCP Authentication](#) on page 921
- [About Authentication Server Groups](#) on page 921
 - [Configuring a RADIUS Authentication Server Group](#) on page 921
 - [Configuring an Active Directory Authentication Server Group](#) on page 923
- [About the Captive Portal](#) on page 923
 - [Configuring Captive Portal Properties](#) on page 924
 - [Customizing the Captive Portal Interface](#) on page 925
 - [Managing Captive Portal Certificates](#) on page 926
 - [Starting the Captive Portal Service](#) on page 927
- [Defining the IPv4 Network and DHCP Ranges](#) on page 928
- [Defining MAC Address Filters](#) on page 929
- [Using the Captive Portal Wizard](#) on page 929
- [Adding and Modifying the Filters and Associations](#) on page 930
- [Monitoring DHCP Authentication](#) on page 931
 - [Viewing DHCP Ranges and Filters](#) on page 931
- [Configuration Example: Configuring Authenticated DHCP](#) on page 931
- [NAC Integration](#) on page 937
- [Configuring NAC with RADIUS Servers](#) on page 938
- [About Authentication Servers](#) on page 938
 - [Adding a Server Group](#) on page 938
 - [Associating a Server Group with a Member](#) on page 940
 - [Managing Server Groups](#) on page 940
 - [Clearing the Authentication Cache](#) on page 940
- [Configuring DHCP Ranges](#) on page 940
 - [Listing DHCP Ranges](#) on page 940

- [About NAC Filters](#) on page 941
 - [Defining NAC Filters](#) on page 942
 - [Disabling NAC Filters](#) on page 943

ABOUT AUTHENTICATED DHCP

This feature provides the ability to control access to your IPv4 networks. (This feature does not support IPv6 networks.) You can divide a network into segments for unauthenticated, authenticated and guest users, and the DHCP server assigns clients to the appropriate segment based on their MAC addresses and authentication credentials.

For example, you can divide a network into one or more production segments for valid employees and systems, a guest segment with access only to the Internet and/or limited public servers, and a quarantine segment with access to a captive portal only. A captive portal is a web page that can provide an option to register as an authenticated user or as a guest.

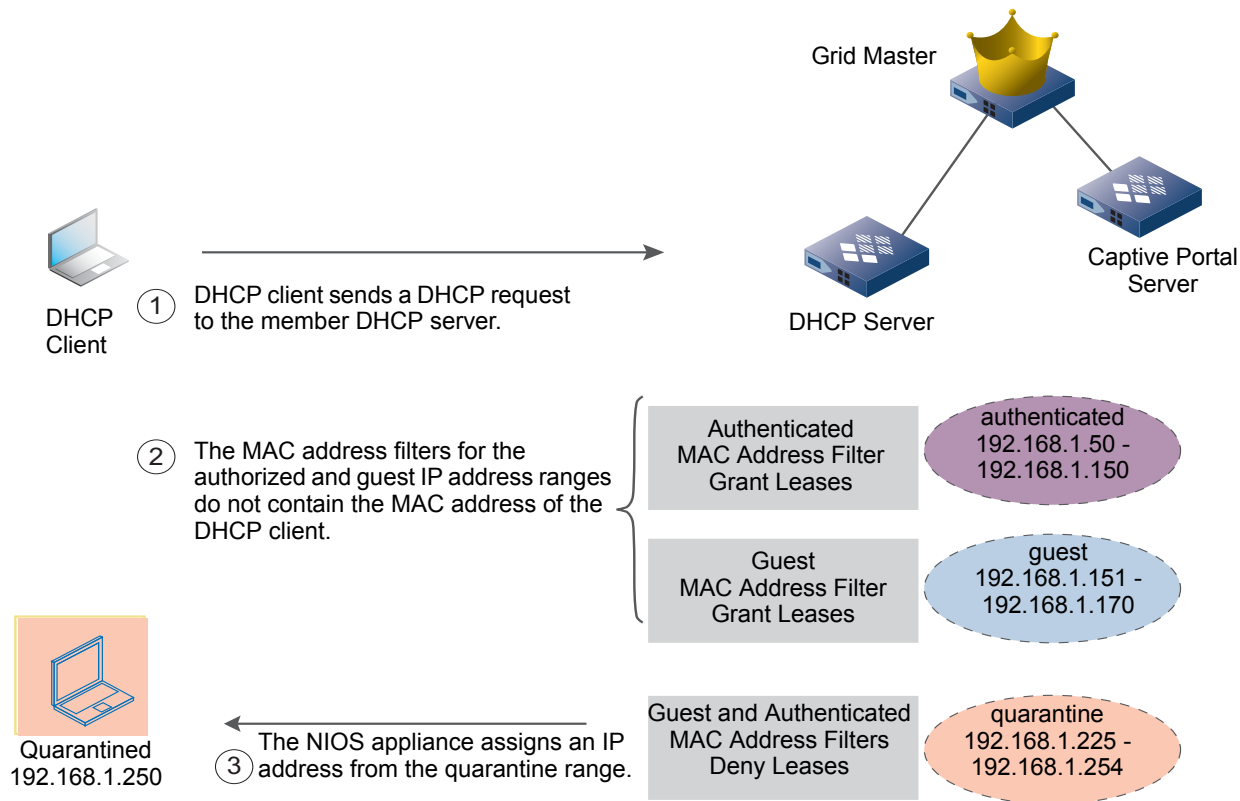
On a member DHCP server, configure DHCP ranges for each access level—quarantine, authenticated, and guest—and create MAC address filters for the DHCP ranges. You can use DHCP options and Access Control Lists (ACLs) on your routers and firewall policies to define the appropriate services for each access level. On another Grid member, configure the captive portal and specify the authentication server group that authenticates the users. You can configure an authentication server group for external servers running RADIUS, LDAP, or Active Directory (AD).

When a DHCP client first sends a request for an IP address, the DHCP server offers an IP address from the quarantine range and directs the client to the captive portal, where the user can register either as an authenticated user or as a guest. When users sign in as guests or are successfully authenticated, the member automatically adds their MAC addresses to the appropriate MAC address filters and assigns addresses out of the appropriate address range.

DHCP Authentication Process

This section illustrates the DHCP authentication process. As illustrated in [Figure 30.1](#), the DHCP authentication process begins when a DHCP client attempts to connect to the network. The member DHCP server checks if the MAC address of the DHCP client matches a MAC address in the guest or authenticated MAC address filters. If the member does not find a match, it assigns an IP address from the quarantine range to the DHCP client. When the client tries to access a web site, it is redirected to the captive portal page.

Figure 30.1 Stage 1: Quarantining an Unauthenticated DHCP Client

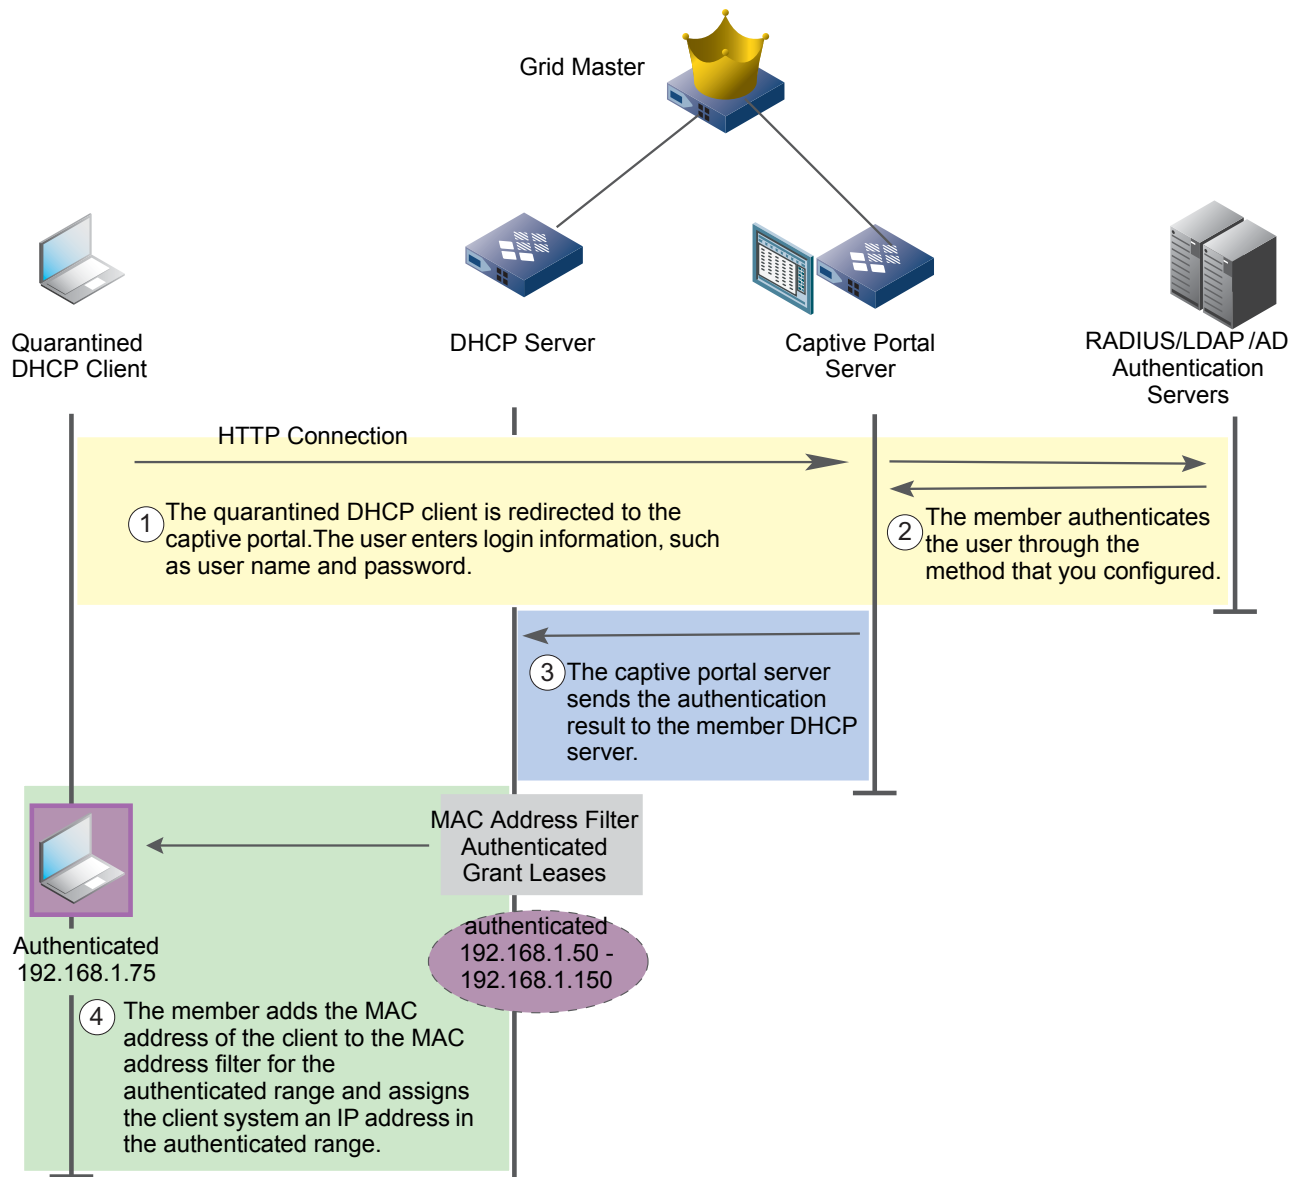


Note that the quarantine range in [Figure 30.1](#) contains MAC address filters to deny leases in the quarantine range to DHCP clients with MAC addresses that match those in the Guest and Authenticated MAC address filters.

When the client connects to the captive portal IP address through its web browser, the user can register and continue the authentication process to obtain an IP address from the authenticated DHCP range, or register as a guest and obtain an IP address from the guest DHCP range.

If the user chooses to continue the authentication process, as shown in [Figure 30.2](#), the member authenticates the user with the authentication service that you configured, which can be RADIUS, LDAP, or AD.

Figure 30.2 Stage 2a: Authenticating the User

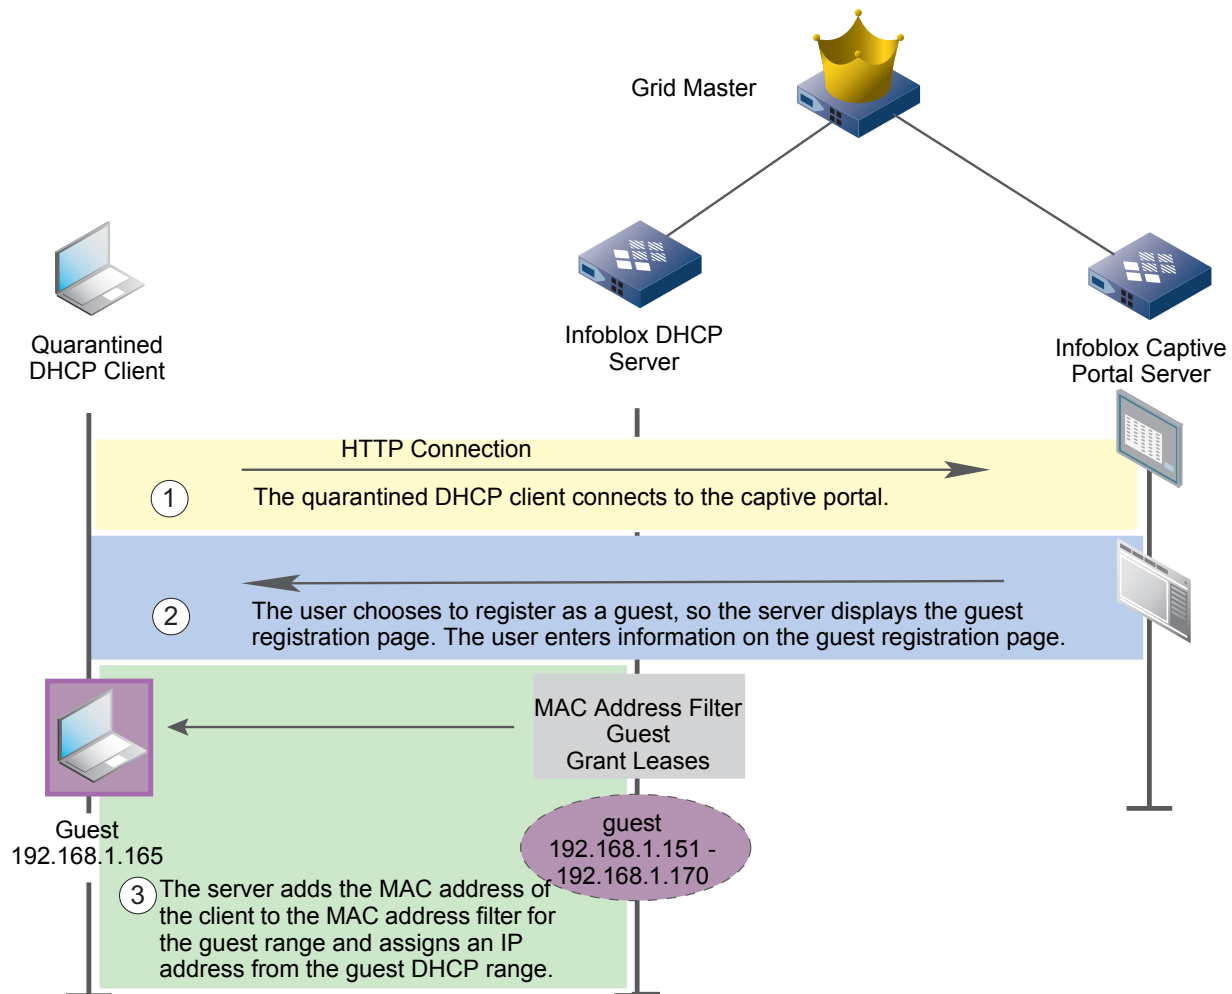


After the client successfully passes the authentication stage, the appliance stores the MAC address of the client in the MAC address filter for the authenticated range. When the client tries to renew its IP address, it receives a new IP address from the authenticated DHCP range.

Note that if the MAC address filter has an expiration period, the member automatically deletes expired MAC addresses from the filter. Therefore, if a DHCP client tries to renew its IP address after the expiration period, the client is redirected to the captive portal because its MAC address is no longer in the MAC address filter. For more information, see [Defining MAC Address Filters](#) on page 838.

If the user chooses to sign in as a guest, as shown in [Figure 30.3](#), the user can fill in the guest registration page provided by the captive portal.

Figure 30.3 Stage 2b: Registering as a Guest



After the user signs in as a guest, the appliance stores the MAC address of the client in the MAC address filter for the guest range. When the DHCP client tries to renew its IP address, it receives a new IP address from the guest DHCP range, unless the MAC address of the client expired and was removed from the filter. In this case, the DHCP client is redirected to the captive portal.

CONFIGURING DHCP AUTHENTICATION

Following are the tasks to configure the DHCP Authentication feature:

1. Configure the authentication server group which the captive portal uses to authenticate DHCP clients. For more information, see [About Authentication Server Groups](#) on page 921.
If the captive portal is used to register guest users and does not authenticate users, then you do not have to configure an authentication server group.
2. Configure the captive portal properties and associate the captive portal with the authentication server group. For more information, see [Configuring Captive Portal Properties](#) on page 924.
3. Optionally, customize the captive portal interface and guest registration page, as described in [Customizing the Captive Portal Interface](#) on page 925. Additionally, if you enabled SSL encryption, upload the required certificates, as described in [Managing Captive Portal Certificates](#) on page 926.
4. Enable the captive portal, as described in [Starting the Captive Portal Service](#) on page 927.
5. Configure the network and a DHCP range for quarantine DHCP clients. Configure DHCP ranges for authenticated and guest DHCP clients, depending on whether you are allowing either one or both types of users to access your network. For information about configuring these DHCP ranges, see [Defining the IPv4 Network and DHCP Ranges](#) on page 928.
6. Run the *Captive Portal* wizard to create MAC address filters for the quarantine range and for the authenticated, and guest DHCP ranges, if configured; and to associate the captive portal server with the member that serves the DHCP ranges. To accomplish these tasks and set other properties, see [Using the Captive Portal Wizard](#) on page 929. Alternatively, you can perform these tasks separately or modify the configured properties, as described in [Adding and Modifying the Filters and Associations](#) on page 930.
7. Enable the DHCP service. For more information, see [Starting DHCP Services on a Member](#) on page 763.

For information about monitoring the captive portal and the DHCP service, see [Monitoring DHCP Authentication](#) on page 931.

ABOUT AUTHENTICATION SERVER GROUPS

Create an authentication server group if you want the captive portal server to authenticate users when they register. You can create an authentication server group with RADIUS servers, LDAP servers, or Active Directory servers, and then associate the group with the member that runs the captive portal and sends the authentication requests. You can associate an authentication server group with multiple captive portals, but you can associate a captive portal with only one authentication server group.

The following sections provide instructions for creating a RADIUS authentication server group, an AD authentication server group and an LDAP server group:

- [Configuring a RADIUS Authentication Server Group](#)
- [Configuring an Active Directory Authentication Server Group](#) on page 923
- [Configuring an LDAP Server Group](#) on page 165

Configuring a RADIUS Authentication Server Group

You can add multiple RADIUS servers to an authentication server group and prioritize them. When the member sends an authentication request, it always selects the first RADIUS server in the list. It only sends authentication requests to the next server on the list if the first server goes down.

To configure the RADIUS authentication server group to which a captive portal server sends authentication requests:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Expand the Toolbar and click **Add -> RADIUS Service**.

3. In the *Add RADIUS Authentication Service* wizard, complete the following:

- **Name:** Enter the name of the server group.
- **RADIUS Servers:** Click the Add icon and enter the following:
 - **Server Name or IP Address:** Enter the RADIUS server FQDN or IP address.
 - **Comment:** You can enter additional information about the server.
 - **Authentication Port:** The destination port on the RADIUS server. The default is 1812.
 - **Authentication Type:** Select the authentication method of the RADIUS server from the drop-down list. You can specify either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol). The default is PAP.
 - **Shared Secret:** Enter the shared secret that the member DHCP server and the RADIUS server use to encrypt and decrypt their messages. This shared secret must match the one you entered on the RADIUS server.
 - **Connect through Management Interface:** Select this to enable the member to use its MGMT port to communicate with just this server.
 - **Disable server:** Select this to disable the RADIUS server if, for example, the connection to the server is down and you want to stop the DHCP server from trying to connect to this server.
 - Click **Test** to validate the configuration and check that the Grid Master can connect to the RADIUS server. Before you can test the configuration though, you must specify the authentication and accounting timeout and retry values.
 If the Grid Master connects to the RADIUS server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the RADIUS server, the appliance displays a message indicating an error in the configuration.
 - Click **Add** to add the RADIUS server to the group.

When you add multiple RADIUS servers to the list, you can use the up and down arrows to change the position of the servers on the list. The member DHCP server connects to the RADIUS servers in the order they are listed.

- **Authentication**
- **Timeout:** The time that the member DHCP server waits for a response from a RADIUS server before considering it unreachable. You can enter the time in milliseconds or seconds. The maximum is 10 seconds.
- **Retries:** The number of times the member DHCP server retries connecting to a RADIUS server before it considers the server unreachable. The default is five.
- **Accounting**
- **Timeout:** The time that the member DHCP server waits for a response from a RADIUS server before considering it unreachable. You can enter the time in milliseconds or seconds. The maximum is 10 seconds.
- **Retries:** The number of times the member DHCP server retries connecting to a RADIUS server before it considers the server unreachable. The default is five.
- **Recovery Interval:** Specifies the duration of time a RADIUS server stays inactive after being down, before becoming eligible to have RADIUS requests sent to it. The recovery interval starts when a RADIUS server is first discovered to be down.
- **Comment:** You can enter additional information about the server group.
- **Disable:** Select this to disable the authentication server group.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring an Active Directory Authentication Server Group

You can add multiple Active Directory servers running Windows Server 2003 or Windows Server 2008 or Windows Server 2012 to an authentication server group and prioritize the servers. When the member sends an authentication request, it always selects the first AD server in the list. It only sends authentication requests to the next server on the list if the first server goes down.

To configure an Active Directory authentication server group for a captive portal server:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Active Directory Services** subtab and click the Add icon.
3. In the *Add Active Directory Authentication Service* wizard, complete the following:
 - **Name:** Enter a name for the service.
 - **Active Directory Domain:** Enter the AD domain name.
 - **Domain Controllers:** Click the Add icon and complete the following to add an AD domain controller:
 - **Server Name or IP Address:** Enter the FQDN or the IP address of the AD server that is used for authentication.
 - **Comment:** Enter additional information about the AD server.
 - **Authentication Port:** Enter the port number on the domain controller to which the member sends authentication requests. The default is 389.
 - **Encryption:** Select **SSL** from the drop-down list to transmit through an SSL (Secure Sockets Layer) tunnel. When you select SSL, the appliance automatically updates the authentication port to 636. Infoblox strongly recommends that you select this option to ensure the security of all communications between the member and the AD server. If you select this option, you must upload a CA certificate from the AD server. Click **CA Certificates** to upload the certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
 - **Connect through Management Interface:** Select this so that the member uses the MGMT port for administrator authentication communications with just this AD server.
 - **Disable server:** Select this to disable an AD server if, for example, the connection to the server is down and you want to stop the Grid member from trying to connect to this server.
 - Click **Test** to test the configuration. If the Grid member connects to the domain controller using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the server, the appliance displays a message indicating an error in the configuration.
 - Click **Add** to add the domain controller to the group.
 - **Timeout(s):** The number of seconds that the Grid member waits for a response from the specified authentication server. The default is 5.
 - **Comment:** Enter additional information about the service.
 - **Disable:** Select this to retain an inactive AD authentication service profile.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

ABOUT THE CAPTIVE PORTAL

The captive portal can be used to register users for authentication, guest users, or both types of users. When a DHCP client attempts to connect to the network and its MAC address is not in any of the configured MAC filters, the member DHCP server assigns it an IP address in the quarantine range. When the quarantined client tries to reach any web site, it is redirected to the captive portal. The captive portal runs a limited DNS server that is used solely to redirect queries to the captive portal web interface.

You can enable the captive portal as a service on any Grid member, except the Grid Master or Grid Master candidate. The Grid member that runs the captive portal cannot run any other service, such as DHCP and DNS. Note that the limited DNS service that the captive portal runs is different from the full-scale DNS service on an Infoblox appliance. The full-scale DNS service must be explicitly disabled on the member that runs the captive portal. For information on disabling DNS service, see [Starting and Stopping the DNS Service](#) on page 478.

You can configure one or more captive portals in the Grid. You can also configure one or more member DHCP servers to use a captive portal to register users. For example, if your organization has two sites, you can configure a captive portal for each site and configure the DHCP servers in each site to use their respective captive portals to authenticate users.

In order for clients to reach the captive portal, you must specify a route to the captive portal. In a network where all IP addresses are on the same subnet, you can configure Option 33 for the quarantine DHCP range. For additional information, see [Quarantine DHCP Range](#) on page 928. On a routed network, you must configure a default route on the router for the subnet.

Following are the tasks to configure a captive portal:

1. Select the Grid member that runs the captive portal and configure its properties, as described in [Configuring Captive Portal Properties](#) on page 924.
2. Optionally, customize the captive portal and registration page. For information about these tasks, see [Customizing the Captive Portal Interface](#) on page 925.
3. If you enabled SSL, generate the CA certificate, as described in [Managing Captive Portal Certificates](#) on page 926.
4. Start the captive portal, as described in [Starting the Captive Portal Service](#) on page 927.

Configuring Captive Portal Properties

When you configure the captive portal properties of a member, you specify if it is used to register users for authentication, guests, or both. If it is used to register guests only, then do not associate it with an authentication server group.

You can specify the VIP address of the Grid member or configure an additional IP address on the loopback interface as the captive portal IP address. Alternatively, if the Grid member supports the LAN2 port and it is enabled, but the NIC failover feature is disabled, you can use the IP address of the LAN2 port as the captive portal IP address. To configure an IP address on the loopback interface, see [Configuring IP Addresses on the Loopback Interface](#) on page 695. For information on the LAN2 port, see [Using the LAN2 Port](#) on page 305.

In addition, you can configure the port on which the appliance listens for authentication requests redirected from the captive portal. When a user logs in to the captive portal, the member sends an authentication request to its associated authentication server group. The member determines future DHCP replies to client requests based on the authentication result.

To configure the properties of the captive portal:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab → **Services** tab.
Grid Manager lists all the members, except for the Grid Master and Grid Master candidate.
3. Select the member that runs the captive portal and click the Edit icon.
4. In the **General Basic** tab of the *Member Captive Portal Properties* editor, complete the following:
 - **Use This Authentication Server Group for Authenticating Captive Portal Users:** Select the authentication server group that authenticates users for this captive portal. For information about authentication server groups, see [About Authentication Server Groups](#) on page 921.
 - **Captive Portal User Types:** Specify whether the captive portal is used to register **Authenticated** users only, **Guest** users only, or **Both**.
 - **Portal IP Address:** Select the IP address of the captive portal server. The appliance lists the VIP address and the IP addresses of the loopback interface and the LAN2 port, if enabled. You can select any of these addresses as the portal IP address.

- **Enable SSL on Portal:** Select this to support encrypted web traffic through SSL/TLS. If you select this option, you must upload a certificate or generate a self-signed certificate. For information about creating and uploading a certificate for the captive portal, see [Managing Captive Portal Certificates](#) on page 926.
 - **Network View:** This field displays if there are multiple network views configured. Select the network view in which the authenticated, quarantine, and guest DHCP ranges belong.
 - **Log Registration Success:** Select to enable the member to log successful registrations in syslog, and then select the logging level from the drop-down list.
 - **Log Registration Failure:** Select to enable the member to log failed registrations in syslog, and then select the logging level from the drop-down list.
5. In the **General Advanced** tab of the editor, you can specify the port on which the member listens for authentication requests redirected from the captive portal. The default port is 4433. Depending on your firewall and network policies, you can configure an unused port greater than 1 and less than 63999.
 6. Save the configuration and click **Restart** if it appears at the top of the screen.

Customizing the Captive Portal Interface

You can customize the captive portal, and if configured, the guest registration page as well. You can upload image files to the appliance and display your own logo, header and footer. In addition, you can upload the acceptable use policies that are displayed on the captive portal and guest registration page.

Following are guidelines for each item you can customize:

- **Logo Image:** The maximum size is 200 pixels wide by 55 pixels high, and the images can be in JPEG, GIF, or PNG format. It displays on top of the header image.
- **Header Image:** The optimal size is 600 pixels wide by 137 pixels high. The image can be in JPEG, GIF, or PNG format. The header displays at the top of the page.
- **Footer Image:** The optimal size is 600 pixels wide by 20 pixels high. The image can be in JPEG, GIF, or PNG format. The footer displays at the bottom of the page.
- **Acceptable Use Policy:** The policy must be saved as a UTF-8 encoded file. It appears below the welcome message in the captive portal. Users can scroll through the policy when they review it. This is used in the captive portal and guest registration page. It must be a .txt file with a maximum of 8000 characters, including white space.

If any of the customizable fields are not configured, then the factory defaults are displayed.

To customize the captive portal:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab-> **Services** tab.
3. Select the member that is running the captive portal and click the Edit icon.
4. Select the **Customization** tab of the *Member Captive Portal Properties* editor.
5. In the **General Captive Portal Customization** section, complete the following:
 - **Company Name:** Enter the name of your company. The company name displays on the title bar of the browser. You can enter a maximum of 256 characters.
 - **Welcome Message:** Type the message that displays on the captive portal. The message can contain a maximum of 300 characters.
 - **Help Desk Message:** Type a message that provides Helpdesk information, such as contact information for technical assistance. The message can contain a maximum of 300 characters.
 - **Logo Image, Header Image, Footer Image, Acceptable Use Policy:** To display the image files and the acceptable use policy on the captive portal, click **Select** beside the item you want to upload. In the *Upload* dialog box, click **Select File** and navigate to the image or text file. Select the file you want to display and click **Upload**. Note that these files have size requirements, as listed earlier in this section.
6. In the **Guest Users Web Page Customization** section, complete the following:

- The appliance displays certain fields on the guest registration page. Select the check boxes of the fields that users are required to complete: **Require First Name**, **Require Middle Name**, **Require Last Name**, **Require Email**, and **Require Phone**.
- **Custom Field 1 — Custom Field 4:** You can display up to four additional fields on the guest registration page. To add a field to the guest registration page, enter a label for that field. The label can have a maximum of 32 characters. Select **Require** to require users to complete the field.

Users can enter a maximum of 128 characters in each of the fields in the captive portal login page and the guest registration page.

7. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing Captive Portal Certificates

When you enable support for encrypted web traffic sent over SSL/TLS, you can do any of the following:

- Generate a self-signed certificate and save it to the certificate store of your browser.
- Request a CA-signed certificate. When you receive the certificate from the CA, upload it on the member running the captive portal.

Generating Self-Signed Certificates

You can generate a self-signed certificate for the captive portal. When you generate a self-signed certificate, you can specify the hostname and change the public/private key size, enter valid dates and specify additional information specific to the captive portal. If you have multiple captive portals, you can generate a certificate for each captive portal with the appropriate hostname.

To generate a self-signed certificate:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab-> **Services** tab.
3. Select the member that is running the captive portal, and then click **HTTPS Cert** -> **Generate Self-signed Certificate** from the Toolbar.
4. In the *Generate Self-signed Certificate* dialog box, complete the following:
 - **Key Size:** Select either **2048** or **1024** for the length of the public key.
 - **Days Valid:** Specify the validity period of the certificate.
 - **Common Name:** Specify the domain name of the captive portal.
 - **Organization:** Enter the name of your company.
 - **Organizational Unit:** Enter the name of your department.
 - **Locality:** Enter a location, such as the city or town of your company.
 - **State or Province:** Enter the state or province.
 - **Country Code:** Enter the two-letter code that identifies the country, such as US.
 - **Admin E-mail Address:** Enter the email address of the captive portal administrator.
 - **Comment:** Enter additional information about the certificate.
5. Click **OK**.

Generating Certificate Signing Requests

You can generate a CSR (certificate signing request) that you can use to obtain a signed certificate from your own trusted CA. Once you receive the signed certificate, you can import it in to the Grid member that runs the captive portal, as described in [Uploading Certificates](#) on page 927.

To generate a CSR:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab-> **Services** tab.

3. Select the member that is running the Captive Portal, and then click **HTTPS Cert -> Create Signing Request** from the Toolbar.
4. In the *Create Signing Request* dialog box, enter the following:
 - **Key Size:** Select either 2048 or 1024 for the length of the public/private key pair.
 - **Common Name:** Specify the domain name of the captive portal.
 - **Organization:** Enter the name of your company.
 - **Organizational Unit:** Enter the name of your department.
 - **Locality:** Enter a location, such as the city or town of your company.
 - **State or Province:** Enter the state or province.
 - **Country Code:** Enter the two-letter code that identifies the country, such as US.
 - **Admin E-mail Address:** Enter the email address of the captive portal administrator.
 - **Comment:** Enter information about the certificate.
5. Click **OK**.

Uploading Certificates

When you upload a certificate, the NIOS appliance finds the matching CSR and takes the private key associated with the CSR and associates it with the newly uploaded certificate. The appliance then automatically deletes the CSR.

If the CA sends an intermediate certificate that must be installed along with the server certificate, you can upload both certificates to the appliance. The appliance supports the use of intermediate certificates to complete the chain of trust from the server certificate to a trusted root CA.

To upload a certificate:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab-> **Services** tab.
3. Select the member that is running the captive portal, and then click **HTTPS Cert -> Upload Certificate** from the Toolbar.
4. In the **Upload** dialog box, click **Select File**, navigate to the certificate location, and click **Open**.
The appliance imports the certificate. When you log in to the appliance again, it uses the certificate you imported.

Downloading Certificates

You can download the current certificate or a self-signed certificate so users can install it in their browsers.

To download a certificate:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab-> **Services** tab.
3. Select the member that is running the captive portal, and then click **HTTPS Cert -> Download Certificate** from the Toolbar.
4. Navigate to where you want to save the certificate and save it.

Starting the Captive Portal Service

Before you start the captive portal service, ensure that the member is not running any other service.

To start the captive portal service:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab -> **Services** tab.
3. Select the member that is configured to run the captive portal service and click the **Start** icon.

DEFINING THE IPv4 NETWORK AND DHCP RANGES

First define the IPv4 network that uses DHCP authentication, and then define the DHCP ranges and services for each access level that you want to provide on the network:

- Quarantine
- Authenticated
- Guest

For information about configuring DHCP IPv4 networks, ranges and services, see [Chapter 26, Managing IPv4 DHCP Data](#), on page 783 and [Chapter 24, Configuring DHCP Properties](#), on page 727.

Quarantine DHCP Range

You must configure a DHCP range for the quarantine level so the member DHCP server can assign IP addresses within that range to unauthenticated DHCP clients. An unauthenticated client is allowed to access the captive portal only and must successfully pass the authentication process before it can receive an IP address from the authenticated range.

Infoblox recommends 30-second leases for addresses in the quarantine DHCP range. This provides enough time for the user authentication process, so when the client attempts to renew the lease at the midpoint of its lease time, the member can then assign the client a new IP address, depending on the result of the authentication process.

When you configure the quarantine DHCP range, you must specify the captive portal IP address as the DNS server for the address range. The captive portal runs a limited DNS server that resolves all queries with the IP address assigned to the web interface on the captive portal.

Note that you can run the *Captive Portal* wizard to automatically set the lease time of the quarantine range to 30 seconds and to add the captive portal IP address as the DNS server. For information about the *Captive Portal* wizard, see [Using the Captive Portal Wizard](#) on page 929. Alternatively, you can set the lease time and the DNS server IP address in the **DHCP** tab of the *DHCP Range* editor. For information about the *DHCP Range* editor, see [Configuring IPv4 Address Ranges](#) on page 796.

To ensure that clients can reach the captive portal, you must specify a route to the captive portal. On a network where all systems can reach each other without going through a router, that is, all IP addresses are on the same subnet, you must configure Option 33 for the quarantine DHCP range. This option specifies a list of static routes that the client should install in its routing cache. The routes consist of a list of IP address pairs. For clients to reach the captive portal, specify the portal IP address first (destination address), and the LAN address of the NIOS appliance second. When the appliance assigns an IP address from the quarantine DHCP range, it also includes the static route that you specified in option 33. For information about configuring DHCP options, see [Configuring IPv4 DHCP Options](#) on page 739. On a routed network, you must configure a default route via the router on the subnet.

Authenticated DHCP Range

Configure a DHCP range for authenticated users if you want the Grid member to assign IP addresses within that range to authenticated DHCP clients. Users that receive an IP address in this range typically are allowed full access to the network.

When a client successfully passes authentication, the member automatically stores its MAC address in the corresponding MAC address filter. When the client attempts to renew the lease at the midpoint of its lease time, the member matches the source MAC address in the request with a MAC address in the filter for the authenticated DHCP address range. The member then assigns the client a new IP address from the authenticated DHCP range.

Guest DHCP Range

Configure a guest DHCP range if you want to provide guest access privileges. You can configure and customize a guest registration page when you configure the captive portal. For information about this feature, see [Customizing the Captive Portal Interface](#) on page 925.

DEFINING MAC ADDRESS FILTERS

After you configure the network and DHCP ranges, you must then configure the MAC address filters and add them to the appropriate DHCP ranges. If you configured DHCP ranges for authenticated and guest users, you must configure MAC address filters for each range with an action of Allow. You must also add those filters to the quarantine range with an action of Deny, to ensure that the member does not allocate an address from the quarantine range to a host whose MAC address matches an entry in the MAC filters for the authenticated and guest DHCP ranges.

When you create the filters, you also specify whether the MAC address entries expire. The member automatically deletes expired MAC address entries from the filter. If a client that registered earlier attempts to renew its IP address or to register after its MAC address has expired, it is redirected to the captive portal because its MAC address is no longer in the filter.

You can run the *Captive Portal* wizard to automatically create the MAC address filters, as described in the next section, [Using the Captive Portal Wizard](#), or you can configure each filter as described in [Defining MAC Address Filters](#) on page 838.

USING THE CAPTIVE PORTAL WIZARD

After you configure the captive portal and the DHCP ranges for each access level, you can use the *Captive Portal* wizard to accomplish the following tasks:

- Associate the captive portal member with the member that serves the DHCP ranges you configured.
- Create MAC address filters and add them to the appropriate DHCP ranges. The wizard allows you to create MAC address filters for the quarantine DHCP range, and for the authenticated and guest DHCP ranges, depending on whether the captive portal is used to register users for authentication, guests, or both. This was specified, when you configured the captive portal properties, described in [Configuring Captive Portal Properties](#) on page 924. For example, if you indicated that the captive portal is used for authenticated users only, then the wizard allows you to create a MAC filter for the authenticated DHCP range only.
 - If the captive portal is used to register users for authentication, the wizard allows you to create a MAC address filter for the authenticated range. The wizard then automatically adds the filter to the authenticated DHCP range with an action of Allow. It also adds the filter to the quarantine range with an action of Deny. This ensures that the member does not allocate an address from the quarantine range to a host whose MAC address matches an entry in the MAC filter.
 - If the captive portal is used to register guest users, the wizard allows you to create a MAC address filter for the guest range. The wizard then automatically adds the filter to the guest DHCP range with an action of Allow. It also adds the filter to the quarantine range with an action of Deny. This ensures that the member does not allocate an address from the quarantine range to a host whose MAC address matches an entry in the MAC filter.
- Add the captive portal IP address as the DNS server for the quarantine address range.
- Set the lease time of the quarantine range to 30 seconds.

To use the *Captive Portal* wizard to complete the tasks for the DHCP authentication feature:

1. From the **Data Management** tab, select the **DHCP** tab, or from the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and click **Configure Captive Portal**.
3. In the *Captive Portal* wizard, complete the following and click **Next**:
 - **Member DHCP:** Select the member DHCP server that uses this captive portal to authenticate users.
 - **Captive Portal:** Select the member that runs the captive portal. Note that the member that runs the captive portal cannot run any other service, such as DHCP or DNS, and cannot be the Grid Master or Grid Master candidate.

4. This panel allows you to create MAC filters for the authenticated and guest DHCP ranges. The MAC filters you can create depend on your entry in the Captive Portal properties of the Grid member. For example, if you indicated that the captive portal is used for authenticated users only, then this panel allows you to create a MAC filter for the authenticated DHCP range only.

You can also specify existing MAC filters, if you want to apply them to the authenticated and guest DHCP ranges.

Complete the following and click **Next**:

- **Authenticated MAC Filter:** Specify a name for the MAC filter that is used for authenticated users.
 - **Expiration Time:** Specify how long a MAC address is stored in the MAC address filter for authenticated users.
 - **Never:** Select this option to store MAC addresses in the MAC address filter until they are manually removed.
 - **Expires in:** Select this option to store MAC addresses in the MAC address filter for the specified period of time.
 - **Guest MAC Filter:** Specify a name for the MAC filter that is used for guest users.
 - **Expiration Time:** Specify how long a MAC address is stored in the MAC address filter for guest users.
 - **Never:** Select this option to store MAC addresses in the MAC address filter until they are manually removed.
 - **Expires in:** Select this option to store MAC addresses in the MAC address filter for the specified period of time.
5. In this panel, you specify the network and address ranges, so the wizard can apply the MAC address filters to the appropriate ranges. Complete the following:
 - **Network:** Select the network that uses DHCP authentication.
 - **Authenticated Range:** Select the IP address range that the appliance uses for authenticated users. The wizard applies the authenticated MAC address filter you specified in the preceding step to this DHCP range with an action of Allow. This effectively allows the member to assign an IP address from the address range to a requesting host whose MAC address matches the MAC address in the filter.
 - **Guest Range:** Select the IP address range that the appliance uses for guest users. The wizard applies the guest MAC address filter you specified in the preceding step to this DHCP range with an action of Allow. This effectively allows the member to assign an IP address from the address range to a requesting host whose MAC address matches the MAC address in the filter.
 - **Quarantine Range:** Select the IP address range that the appliance uses for quarantined addresses. The wizard applies the authenticated and guest MAC address filters to the quarantine DHCP range with an action of Deny. This effectively denies an address request from a host whose MAC address matches an entry in the MAC filters for the authenticated and guest DHCP ranges.
 6. Save the configuration and click **Restart** if it appears at the top of the screen.

ADDING AND MODIFYING THE FILTERS AND ASSOCIATIONS

The *Captive Portal* wizard simplified the configuration process by accomplishing a number of tasks simultaneously. To accomplish each task separately, or to modify the filters or associations after you have run the wizard:

- To define the MAC address filters for each range, see [Defining MAC Address Filters](#) on page 838.
- To bind each filter to the appropriate DHCP range, see [Applying Filters to DHCP Address Ranges](#) on page 849.
- To specify the DNS server IP address for the quarantine range and set the lease time to 30 seconds, see [Configuring General IPv4 DHCP Properties](#) on page 730.
- To associate a member DHCP server with a captive portal and specify the MAC filters for the authenticated and guest DHCP ranges:
 1. From the **Data Management** tab, select the **DHCP** tab → **Members** tab → *member* check box → Edit icon.
 2. In the *Member DHCP Properties* editor, click the **IPv4 Authenticated DHCP** tab and complete the following:

- **Use this Captive Portal for Infoblox Authenticated DHCP:** Select this check box and select the captive portal that you want to associate with the member.
 - **Authenticated User MAC Filter:** Select the MAC filter used for authenticated users. To change your selection, click **Clear** and click **Select** again.
 - **Guest User MAC Filter:** Select the MAC filter for guest users. To change your selection, click **Clear** and click **Select** again.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

MONITORING DHCP AUTHENTICATION

You can monitor the status of the captive portal service, as described in [Monitoring Services](#) on page 961. You can check its status in the *Grid Status* widget and the *Member Status* widget on the Dashboard. For information about these widgets, see [Chapter 2, Dashboards](#), on page 87.

You can also view the MAC addresses that were added to each MAC address filter, as described in [Viewing MAC Address Filter Items](#) on page 856.

Viewing DHCP Ranges and Filters

To view the newly created MAC address filters:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab.
Grid Manager lists all the configured filters.
2. You can select a filter and view or configure its properties, such as extensible attributes.
For more information about the filters and editing their properties, see [Managing DHCP Filters](#) on page 854.

To view the DHCP ranges and the newly added filters:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network*.
2. Select the DHCP range you want to view and click the Edit icon.
3. If the editor is in Basic mode, click **Toggle Expert Mode**.
4. Click the **Filters** tab to view the filters.

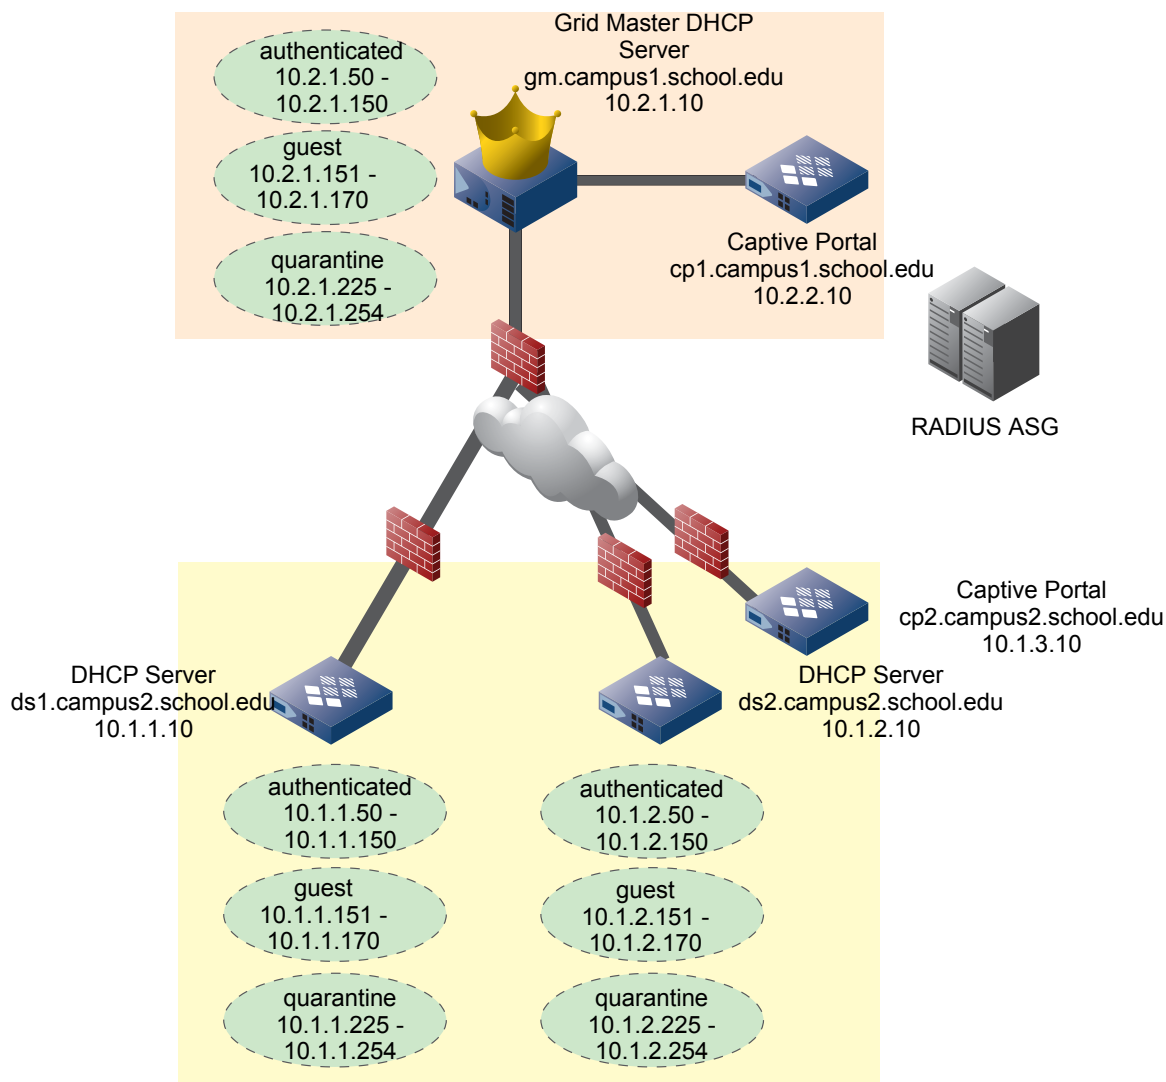
To verify that the captive portal is the DNS server in the quarantine range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network*.
2. Select the quarantine DHCP range and click the Edit icon.
3. In the *DHCP Range* editor, click the **DHCP** tab.
The captive portal IP address is listed in the DNS Servers table.

CONFIGURATION EXAMPLE: CONFIGURING AUTHENTICATED DHCP

In this example, a school (school.edu) has two locations, its main campus, campus1.schoool.edu, and a satellite campus, campus2.school.edu. It has a captive portal server in each location. In the main campus, the Grid Master also functions as a DHCP server and uses a captive portal server to register DHCP clients. In the satellite campus, two members serve DHCP and use the same captive portal server. The captive portal servers use the same RADIUS authentication server group to authenticate users.

Figure 30.4



Create the RADIUS Authentication Server Group

Create the RADIUS authentication server group and add two RADIUS servers to the group.

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Expand the Toolbar and click **Add -> RADIUS Service**.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
 - **Name:** Enter **RADIUS ASG**.
 - **RADIUS Servers:** Click the Add icon and enter the following:
 - **Server Name or IP Address:** Enter the RADIUS server FQDN, which is **rs1.school.edu**.
 - **Authentication Port:** Accept the default port (1812).
 - **Authentication Type:** Select the PAP authentication method.
 - **Shared Secret:** Enter **no1nose**.
 - **Authentication**
 - **Timeout:** Enter 5 seconds.
 - **Retries:** Accept the default, which is five.

- **Accounting**
 - **Timeout:** Enter 5 seconds.
 - **Retries:** Accept the default, which is five.
 - Click **Test** to validate the configuration and check that the Grid Master can connect to the RADIUS server.
Grid Manager displays a message confirming the configuration is valid.
- Click **Add** to add another RADIUS server to the group, and then enter the following:
 - **Server Name or IP Address:** Enter the RADIUS server FQDN, which is **rs2.school.edu**.
 - **Authentication Port:** Accept the default port (1812).
 - **Authentication Type:** Select the PAP authentication method.
 - **Shared Secret:** Enter **no1nose**.
- **Authentication**
 - **Timeout:** Enter 5 seconds.
 - **Retries:** Accept the default, which is five.
- **Accounting**
 - **Timeout:** Enter 5 seconds.
 - **Retries:** Accept the default, which is five.
 - Click **Test** to validate the configuration and check that the Grid Master can connect to the RADIUS server.
Grid Manager displays a message confirming the configuration is valid.

4. Click **Save & Close**.

Configure the Captive Portal Properties

Configure the captive portal properties of cp1.campus1.school.edu.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab → **Services** tab.
3. Select the member **cp1.campus1.school.edu** and click the Edit icon.
4. In the **General Basic** tab of the *Member Captive Portal Properties* editor, complete the following:
 - **Use This Authentication Server Group for Authenticating Captive Portal Users:** Select **RADIUS ASG**.
 - **Captive Portal User Types:** Select **Both**.
 - **Portal IP Address:** Select **10.2.2.10**.
 - **Enable SSL on Portal:** Select this option.
 - **Log Registration Success:** Select **Informational**.
 - **Log Registration Failure:** Select **Informational**.
5. Click **Save & Close**.

Configure the captive portal properties of cp2.campus2.school.edu.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab → **Services** tab.
3. Select the member cp2.campus2.school.edu and click the Edit icon.
4. In the **General Basic** tab of the *Member Captive Portal Properties* editor, complete the following:
 - **Use This Authentication Server Group for Authenticating Captive Portal Users:** Select **RADIUS ASG**.
 - **Captive Portal User Types:** Select **Both**.
 - **Portal IP Address:** Select **10.1.3.10**.
 - **Enable SSL on Portal:** Select this option.
 - **Log Registration Success:** Select **Informational**.

- **Log Registration Failure:** Select **Informational**.

5. Click **Save & Close**.

Customize the Captive Portals

Customize the captive portal cp1.campus1.school.edu.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab→ **Services** tab.
3. Select **cp1.campus1.school.edu** and click the Edit icon.
4. Select the **Customization** tab of the *Member Captive Portal Properties* editor.
5. In the **General Captive Portal Customization** section, complete the following:
 - **Company Name:** Enter **School**.
 - **Welcome Message:** Type the following: **Welcome to School. Please sign in.**
 - **Help Desk Message:** Type: **To reach the Helpdesk, call (408) 111-2222 or email helpdesk@school.edu.**
 - **Logo Image:** Click **Select** beside the logo file and upload it.
6. In the **Guest Users Web Page Customization** section, complete the following:
 - Select the check boxes beside **Require First Name, Require Last Name, Require Email**.
7. Click **Save & Close**.

Select the other captive portal server, cp2.campus2.school.edu, and enter the same information.

Generate a Self-Signed Certificate and Upload It

To generate a self-signed certificate for cp1.campus1.school.edu:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab→ **Services** tab.
3. Select **cp1.campus1.school.edu**, and then click **HTTPS Cert** → **Generate Self-signed Certificate** from the Toolbar.
4. In the *Generate Self-signed Certificate* dialog box, complete the following:
 - **Key Size:** Select **1024** for the length of the public key.
 - **Days Valid:** Enter **60 days**.
 - **Common Name:** Enter **cp1.campus1.school.edu**.
5. Click **OK**.
6. Click **Save & Close**.

To generate a self-signed certificate for the captive portal cp2.campus2.school.edu:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab→ **Services** tab.
3. Select **cp2.campus2.school.edu**, and then click **HTTPS Cert** → **Generate Self-signed Certificate** from the Toolbar.
4. In the *Generate Self-signed Certificate* dialog box, complete the following:
 - **Key Size:** Select **1024** for the length of the public key.
 - **Days Valid:** Enter **60 days**.
 - **Common Name:** Enter **cp2.campus2.school.edu**.
5. Click **OK**.
6. Click **Save & Close**.

Start the Captive Portal Service

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab -> **Services** tab.
3. Select **cp1.campus1.school.edu** and **cp2.campus2.school.edu**, and then click the Start icon.

Configure the Networks and DHCP Ranges

Configure the network on the Grid Master.

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section.
2. Click the Add drop-down list and select **IPv4 Network**.
3. In the *Add IPv4 Network* wizard, select one of the following and click **Next**:
 - **Add Network**: Click this.
4. Complete the following and click **Next**:
 - **Address**: Enter **10.2.1.0/24**.
5. Complete the following to assign the network to the Grid Master:
 - **Add Infoblox Member**: Select **gm.campus1.school.edu**.
6. Click **Save & Close**.

Configure the ranges on the Grid Master.

To create the authenticated range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section.
2. Click the 10.2.1.0/24 network link, and then click the Add drop-down list and select **DHCP Range**.
3. In the *Add IPv4 Range* wizard, select **Add DHCP Range** and click **Next**:
4. Complete the following:
 - **Network**: Click **Select Network** and select **10.2.1.0/24**.
 - **Start**: Enter **10.2.1.50**.
 - **End**: Enter **10.2.1.150**.
 - **Name**: Enter **authenticated range**.
5. Click **Next** and complete the following:
 - **Grid Member**: Select this option and select **gm.campus1.school.edu**.
6. Click **Save & Close**.

To create the guest range:

1. Click the 10.2.1.0/24 network link, and then click the Add drop-down list and select **DHCP Range**.
2. In the *Add IPv4 Range* wizard, select **Add DHCP Range** and click **Next**:
3. Complete the following:
 - **Network**: Click **Select Network** and select **10.2.1.0/24**.
 - **Start**: Enter **10.2.1.151**.
 - **End**: Enter **10.2.1.170**.
 - **Name**: Enter **guest range**.
4. Click **Next** and complete the following:
 - **Grid Member**: Select this option and select **gm.campus1.school.edu**.
5. Click **Save & Close**.

To create the quarantine range:

1. Click the 10.2.1.0/24 network link, and then click the Add drop-down list and select **DHCP Range**.
2. In the *Add IPv4 Range* wizard, select **Add DHCP Range** and click **Next**:

3. Complete the following:
 - **Network:** Click **Select Network** and select 10.2.1.0/24.
 - **Start:** Enter 10.2.1.225.
 - **End:** Enter 10.2.1.254.
 - **Name:** Enter **quarantine range**.
4. Click **Next** and complete the following:
 - **Grid Member:** Select this option and select **gm.campus1.school.edu**.
5. Click **Save & Close**.

Create the network and DHCP ranges for the DHCP servers ds1.campus1.school.edu and ds2.campus2.school.edu.

Run the Captive Portal Wizard

Run the *Captive Portal* wizard to associate the Grid Master with its captive portal, and to configure the MAC address filters:

1. From the **Data Management** tab, select the **DHCP** tab, or from the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and click **Configure Captive Portal**.
3. In the *Captive Portal* wizard, complete the following and click **Next**:
 - **Member DHCP:** Select the Grid Master, **gm.campus1.school.edu**.
 - **Captive Portal:** Select **cp1.campus1.school.edu**.
4. Complete the following and click **Next**:
 - **Authenticated MAC Filter:** Enter **Auth_MAC_Filter**.
 - **Expiration Time:** Select **Never**.
 - **Guest MAC Filter:** Enter **Guest_MAC_Filter**.
 - **Expiration Time:** Select **Never**.
5. Complete the following:
 - **Network:** Select 10.2.1.0/24.
 - **Authenticated Range:** Select 10.2.1.50 - 10.2.1.150.
 - **Guest Range:** Select 10.2.1.151 - 10.2.1.170.
 - **Quarantine Range:** Select 10.2.1.225 - 10.2.1.254.
6. Click **Save & Close**.

Run the *Captive Portal* wizard to associate ds1.campus2.school.edu with the captive portal server cp2.campus2.school.edu, and then run it again to associate ds2.campus2.school.edu with the same captive portal server.

Start the DHCP Service

To start the DHCP service on the Grid Master:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab.
2. Select the Grid Master gm.campus1.school.edu, and the two members, ds1.campus2.school.edu and ds2.campus2.school.edu.
3. Expand the Toolbar and click **Start**.
4. In the *Start Member DHCP Service* dialog box, click **Yes**.
5. Grid Manager starts DHCP services on the Grid Master and on the selected members.

NAC INTEGRATION

You can configure member DHCP servers to send authentication requests to RADIUS servers and to allocate addresses based on the authentication results. This allows you to place DHCP clients into separate network segments.

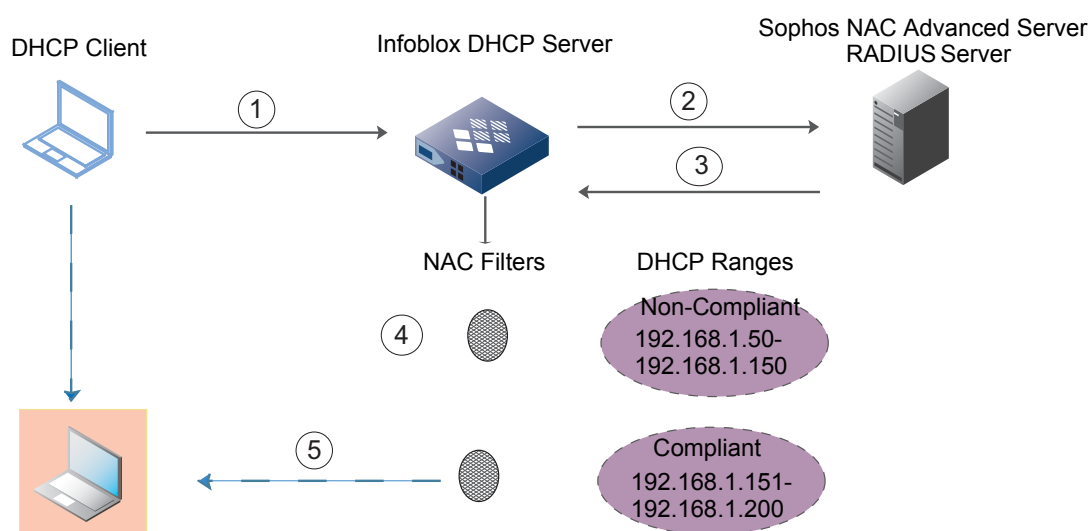
You can divide your network into different segments by configuring address ranges and applying NAC filters to them. NAC filters use authentication results from RADIUS servers as matching criteria for granting or denying address requests.

When a DHCP client requests a lease, the member DHCP server can query a remote backend RADIUS server such as the Sophos NAC Advanced server to determine if the DHCP client is authorized to access the network. A Sophos NAC Advanced server is an access-control and compliance server that supports the RADIUS protocol.

The RADIUS server then checks its database and provides the compliance state and user class, if configured, of the DHCP client. The member DHCP server matches the response with the configured NAC filters, and grants a lease to the appropriate network segment.

[Figure 30.5](#) presents an example illustrating the authentication process and how a member DHCP server matches the response with NAC filters to determine whether to grant or deny a lease. In the example, there are two DHCP ranges configured, each with a NAC filter that specifies RADIUS compliance state of DHCP clients allowed in each range.

Figure 30.5



The following steps relate to [Figure 30.5](#).

1. A DHCP client sends a DHCPDISCOVER, DHCPREQUEST, or DHCPINFORM to the Infoblox DHCP server.
2. The DHCP server sends the RADIUS server a RADIUS Access-Request packet that includes RADIUS VSAs (Vendor Specific Attributes) with the MAC address and DHCP transaction ID of the DHCP client.
3. When the RADIUS server receives the Access-Request packet, it does the following:
 - a. It looks up the MAC address in its database to retrieve the associated compliance state and user class.
 - b. The RADIUS server sends back a RADIUS Access-Accept packet that includes RADIUS VSAs with the compliance status and user class.
4. The DHCP server receives the Access-Accept packet and tries to match the response with a NAC filter.
5. The DHCP server matches the response with the NAC filter for compliant DHCP clients and sends the DHCP client a DHCPOFFER that contains an IP address from the corresponding DHCP range. The server also provides the configuration and options associated with that range.

CONFIGURING NAC WITH RADIUS SERVERS

Complete the following tasks to configure the RADIUS server and the member DHCP server.

On an already functioning RADIUS server:

- Add the member DHCP server as a RADIUS client. Make sure that the shared secret you enter on the RADIUS server matches the shared secret that you specify when you add the server to the authentication server group in Grid Manager.
Note that on Grid Manager, you can enter only one shared secret for each RADIUS server. Therefore, on a RADIUS server, you must define the same shared secret for all Grid members that connect to it.
For information about adding RADIUS clients, refer to the documentation for the RADIUS server.
- Add the Infoblox Grid Master as a RADIUS client, even if it is not going to perform NAC authentication. This enables you to test the connection to the RADIUS server.

On the member DHCP server:

1. Configure the authentication server group for the RADIUS servers. For information, see [Adding a Server Group](#) on page 938.
2. Associate the authentication server group with the Grid member. For information, see [Associating a Server Group with a Member](#) on page 940.
3. Configure the network and the DHCP ranges. For information, see [Configuring DHCP Ranges](#) on page 940.
4. Configure the NAC filters, as described in [About NAC Filters](#) on page 941.
5. Apply the NAC filters to the DHCP ranges, as described in [Applying Filters to DHCP Address Ranges](#) on page 849.
6. Enable the DHCP service. For information, see [Starting DHCP Services on a Member](#) on page 763.

Optionally, you can do the following:

- Manage the authentication cache, as described in [Clearing the Authentication Cache](#) on page 940.

ABOUT AUTHENTICATION SERVERS

You can create a RADIUS authentication server group for Sophos NAC Advanced servers, and then associate the group with the member DHCP server that sends authentication requests. The member DHCP server tries to connect to each Sophos NAC Advanced server in the group using one of the following methods: Ordered List or Round Robin.

In the Ordered List method, the member DHCP server always selects the first Sophos NAC Advanced server in the list when it sends an authentication request. It only sends authentication requests to the next server on the list if the first server goes down.

In the Round Robin method, the member DHCP server selects the first Sophos NAC Advanced server for the first request, the second server for the next request, and so on until it selects the last server in the list. Then it starts with the first server in the list and continues the same selection process.

Each member DHCP server can have only one RADIUS server group assigned, but a RADIUS server group can be assigned to multiple member DHCP servers.

Adding a Server Group

To create a RADIUS authentication server group for Sophos NAC Advanced servers:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Expand the Toolbar and click **Add -> RADIUS Service**.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
 - **Name:** Enter the name of the server group.
 - **RADIUS Servers:** Click the Add icon and enter the following:

- **Server Name or IP Address:** Enter the Sophos NAC Advanced server FQDN or IP address.
- **Comment:** You can enter additional information about the server.
- **Authentication Port:** The destination port on the Sophos NAC Advanced server. The default is 1812.
- **Authentication Type:** Select the authentication method of the RADIUS server from the drop-down list. You can specify either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol). The default is PAP.
- **Shared Secret:** Enter the shared secret that the member DHCP server and the Sophos NAC Advanced server use to encrypt and decrypt their messages. This shared secret must match the one you entered on the Sophos NAC Advanced server.
- **Enable Accounting:** Leave this blank. RADIUS accounting is not supported.
- **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for communications with just this server.
- **Disable server:** Select this to disable the Sophos NAC Advanced server if, for example, the connection to the server is down and you want to stop the DHCP server from trying to connect to this server.
- Click **Test** to validate the configuration and check that the Grid Master can connect to the Sophos NAC Advanced server. Before you can test the configuration though, you must specify the authentication and accounting timeout values.

If the Grid Master connects to the Sophos NAC Advanced server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the Sophos NAC Advanced server, the appliance displays a message indicating an error in the configuration.

- Click **Add** to add the Sophos NAC Advanced server to the server group.

When you add multiple Sophos NAC Advanced servers to the list, you can use the up and down arrows to change the position of the servers on the list. The member DHCP server connects to the Sophos NAC Advanced servers in the order they are listed.

- **Authentication**
- **Timeout:** The time that the member DHCP server waits for a response from a Sophos NAC Advanced server before considering it unreachable. You can enter the time in milliseconds or seconds.
- **Retries:** The number of times the member DHCP server retries connecting to a Sophos NAC Advanced server before it considers the server unreachable. The default is five.
- **Mode:** Specifies how the member DHCP server selects the first Sophos NAC Advanced server to contact.
 - **Ordered List:** The member DHCP server always selects the first Sophos NAC Advanced server in the list when it sends an authentication request. It queries the next server only when the first server is considered down. This is the default.
 - **Round Robin:** The member DHCP server selects the first Sophos NAC Advanced server for the first request, the second server for the next request, and so on. If the last server is reached, then the DHCP server starts with the first server in the list, and so on.
- **Enable Authentication Cache:** The member DHCP server automatically caches authentication results for 120 seconds. When you enable this option, you can override this default in the **Cache Time to Live** field. You must enable this option to clear the cache, as described in [Clearing the Authentication Cache](#) on page 940.
- **Cache Time to Live:** Specifies the duration of time an authentication result is stored. The default is one hour. The maximum is 259200 seconds (3 days).
- **Recovery Interval:** Specifies the duration of time a Sophos NAC Advanced server stays inactive after being down, before becoming eligible to have RADIUS requests sent to it. The recovery interval starts when a Sophos RADIUS server is first discovered to be down.
- **Comment:** You can enter additional information about the server group.
- **Disable:** Select this to disable the authentication server group.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Associating a Server Group with a Member

To associate an authentication server group with a member DHCP server:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box, and click the Edit icon.
2. If the *Member DHCP Properties* editor is in Basic mode, click **Toggle Expert Mode**.
3. Select the **IPv4 Authenticated DHCP** tab.
4. Click the **Use this Authentication Server Group for Sophos/RAIDUIS Authenticated DHCP** check box, and then select a group from the drop-down list.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

Managing Server Groups

To view the list of authentication server groups, from the **Administration** tab, click the **Authentication Server Groups** tab and expand the **RADIUS Service** subtab. For each server group, you can view the server group name, comments, and whether the group is available or disabled. You can then select a server group to modify or delete it.

To modify a server group, select it and click the Edit icon. You can modify any of its properties, and add or delete servers from the group. When you delete a Sophos NAC Advanced server from a group, the appliance permanently deletes it.

To delete a server group, select it and click the Delete icon. When you delete an authentication server group, the appliance permanently deletes it.

Clearing the Authentication Cache

The authentication cache can store authentication results for up to 20,000 DHCP clients. When the cache reaches its limit, the DHCP member logs a message in syslog. To clear the entire cache or the cache entry of a specific MAC address, you must enable the authentication cache in the RADIUS Service wizard or editor.

To clear the entire authentication cache:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box.
2. Expand the Toolbar and select **Clear** -> **Authentication Cache**.
3. When the **Clear Authentication Cache** confirmation dialog appears, click **Purge**.

To delete a specific entry:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* check box.
2. Expand the Toolbar and select **Clear** -> **Authentication Record**.
3. In the **Clear Authentication Record** dialog box, enter the DHCP client MAC address, and then click **Purge**.

CONFIGURING DHCP RANGES

Create the IPv4 network and DHCP ranges as described in [Chapter 26, Managing IPv4 DHCP Data](#), on page 783. You can create multiple DHCP ranges and apply one or more NAC filters to each of them.

Listing DHCP Ranges

By default, DHCP ranges are listed according to their start addresses. You can reorder them according to the order in which you want the member DHCP server to evaluate the ranges.

Consider the following sample DHCP ranges:

- 10.20.30.100-10.20.30.199 (NAC filter that allows leases for compliant DHCP clients)
- 10.20.30.0-10.20.30.99 (No filters)

If the DHCP range with the NAC filter is listed before the range with no filters, then the DHCP server consults the Sophos NAC Advanced server and applies the NAC filter before it grants a lease. It grants leases from the range with no filters only if no NAC filters matched or after all leases from the first range are exhausted. If the first range is the production range and the second range is for the quarantine group, then the server applies the NAC filters for the production range, before it grants leases to the quarantine range.

To change the order of DHCP ranges in a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network*.
2. Expand the Toolbar and click **Order DHCP Ranges**.
3. In the *Order DHCP Ranges* dialog box, click the up and down arrows to move ranges up or down on the list. The Priority value changes accordingly.
4. Click **OK**.

You can view the DHCP objects in a network, including its DHCP ranges by navigating to the **DHCP** tab -> **Networks** tab -> **Networks** panel, and then clicking the network link. You can select the Priority column for display to view the order of the DHCP ranges. For information about editing the columns, see [Customizing Tables](#) on page 53.

ABOUT NAC FILTERS

You can define NAC filters that specify authentication results from a remote, backend RADIUS server such as the Sophos NAC Advanced server. You can then apply each filter to a DHCP range or range template and indicate whether the DHCP server grants or denies a lease when the authentication result matches the filter. You can apply NAC filters to any DHCP range and DHCP range template.

NAC filters are enabled by default. When necessary, you can disable them for the entire Grid so you can perform maintenance on your RADIUS server. When you disable NAC filters, no service interruptions, service down times, configuration changes, or server restarts are required. For information about how to disable NAC filters, see [Disabling NAC Filters](#) on page 943.

In a NAC filter, you can define rules that specify the following:

- The status of the RADIUS authentication server group:
 - Success: At least one of the servers in the RADIUS authentication server group is up.
 - Fail: The MAC address in the DHCP request is not in the authentication cache and all servers in the server group are down.
 - Disabled: The RADIUS authentication server group is disabled, all the servers in the group are disabled, or the member is not assigned a server group.
- The response from the RADIUS server:
 - Accept: The response is an Access-Accept packet.
 - Reject: The response is an Access-Reject packet.
- Whether the Access-Accept packet contains an error. The Infoblox DHCP server expects certain RADIUS VSAs in the Access-Accept packet. An error occurs when any of the RADIUS VSAs are missing. For information about the Access-Accept packet and the RADIUS VSAs, refer to the documentation for the specified RADIUS server.
 - Yes: The Access-Accept packet does not include one or more RADIUS VSAs.
 - No: There are no errors in the Access-Accept packet.
- A compliance state: unknown, non-compliant, compliant or partially compliant
- A RADIUS server user class

When the member DHCP server receives an address request, it checks the DHCP ranges in their priority order. For information about the order of DHCP ranges, see [Listing DHCP Ranges](#) on page 940.

For each DHCP range, it checks if the request matches any MAC filters, relay agent filters, or DHCP option filters that apply to the range. (For information about these filters, see [Chapter 29, Configuring IPv4 DHCP Filters](#), on page 831.) If any of those filters match, then the member either grants or denies a lease to the DHCP client, based on the filter. If none of those filters match and there are NAC filters defined, then the member tries to send an authentication request to a server in the RADIUS authentication server group.

If you want the member DHCP server to grant leases to specific DHCP ranges in case the RADIUS authentication server group is considered disabled (server state = disabled) or if all RADIUS servers are down (server state = failure), create a NAC filter for each situation and apply it to the appropriate range.

Note that when you create a NAC filter, you do not have to include rules that specify prerequisite conditions. For example, when you create a filter that specifies a RADIUS server compliance state or user class, you do not have to include rules that specify the following: server state=success, server response=accept, and server error = no.

Defining NAC Filters

To define a NAC filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **IPv4 Filters** tab, and then expand the Toolbar and click **Add -> IPv4 NAC Filter**.

or

From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 NAC Filter**.

2. In the *Add Filter Wizard*, complete the following and click **Next**:
 - **Name:** Enter a name for the filter. You can enter a maximum of 255 characters. The name must be unique within a specific network. If you want to specify option settings in the filter, the name must be unique among all NAC filters.
 - **Comment:** Optionally, enter additional information about the NAC filter.
3. Create a rule as follows:
 - In the first drop-down list, select one of the following criterion: **Compliance State**, **Server Error**, **Server Response**, **Server State** or **User Class**.
 - In the second drop-down list, select an operator: **equals** or **does not equal**.
 - The selections in the third drop-down list depend on the criterion you selected:

Compliance State: Select one of the following compliance states: **Unknown**, **Non-compliant**, **Compliant** or **Partially Compliant**.

Server Error: The Infoblox DHCP server expects certain RADIUS VSAs in the Access-Accept packet. When any of the VSAs are missing, then the DHCP server considers this an error. For information about the Access-Accept packet and the VSAs, refer to the documentation for the specified RADIUS server. Select one of the following:

 - **Yes:** Create a rule that matches when the RADIUS server sends an Access-Accept packet with a missing VSA.
 - **No:** Create a rule that matches when the RADIUS server sends an Access-Accept packet with no errors.

Server Response: Select one of the following:

 - **Accept:** Create a rule that matches when the server sends back an Access-Accept packet.
 - **Reject:** Create a rule that matches when the server sends back an Access-Reject packet.

Server State: Select one of the following:

 - **Success:** Create a rule that matches when at least one RADIUS server in the group is up.
 - **Fail:** Create a rule that matches when the MAC address of the DHCP client is not in the cache and all RADIUS servers in the server group are down.
 - **Disable:** Create a rule that matches when the RADIUS authentication server group is disabled, all servers in the group are disable, or the member was not assigned a server group.

User Class: Enter the RADIUS user class value, for example, NACDeny. The member DHCP server does not validate the entry. Therefore, you must make sure that the user class you enter matches the user class name on the RADIUS server.

To add another rule:

- Click **+** to add another rule at the same level.
- Click **|<** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level and above the first rule.
- Click **>|** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level.

After you add all the match rules, you can click **Preview** to view the rules or click **Reset** to remove the previously configured rules and start again.

4. Click **Next** and complete the following to define DHCP options:

- **Option Space:** Select an option space from the drop-down list. This field is not displayed if you do not have custom option spaces. The appliance uses the **DHCP** option space as the default.
- **Lease Time:** Enter the value of the lease time in the field and select the time unit from the drop-down list. The lease time applies to hosts that meet the filter criteria.

Options to Merge with Object Options

Click the Add icon. Grid Manager adds a new row to the table with the default **DHCP** option space and option name displayed. Complete the following:

- **Option Space:** Click the down arrow and select an option space from the drop-down list. The selected option space contains the corresponding DHCP options.
- **Option Name:** Click the down arrow and from the drop-down list, select the DHCP option you want to return to the matching client.
- **Value:** Enter the match value that you want the filter to use for the selected DHCP option. For example, enter the value 172.124.3.0 for the SUNW.SrootIP4 option.

To add more options to the filter, click the Add icon and repeat the steps.

5. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 293.

6. Save the configuration and click **Restart** if it appears at the top of the screen.

After you add NAC filters, you must then apply them to DHCP ranges, as described in [Applying Filters to DHCP Address Ranges](#) on page 849. You can also list, modify or delete NAC filters, as described in [Managing DHCP Filters](#) on page 854.

Disabling NAC Filters

NAC filters are enabled by default. When you disable them, the appliance bypasses evaluations of all NAC filters for the entire Grid. There are no configuration changes, service restarts, or service down times when you disable the NAC filters. The appliance keeps the filter configurations so you can enable them at a later time.

To disable NAC filters for the Grid:

1. From the **Data Management** tab -> **DHCP** tab, select **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**, select the **General** tab -> **Advanced** tab, and then complete the following in the Common Properties section:
 - **Disable All NAC Filters:** Select this to disable all NAC filters in the Grid. The appliance keeps the filter configurations so you can enable them when needed.



Chapter 31 Managing Leases

This chapter explains how to manage IPv4 and IPv6 leases. It contains the following sections:

- [*About DHCP Leases*](#) on page 946
- [*Viewing Current Leases*](#) on page 946
- [*Viewing Detailed Lease Information*](#) on page 948
- [*Viewing Lease History*](#) on page 949
- [*Viewing Lease Event Detailed Information*](#) on page 949
- [*Exporting Lease Records*](#) on page 950
- [*Clearing Leases*](#) on page 950

ABOUT DHCP LEASES

Historical DHCP lease records complement the real-time DHCP lease viewer by allowing the appliance to store and correlate DHCP lease information over the lifetime of a lease. You can see critical information such as when the appliance issued or freed an IPv4 or IPv6 address, the MAC address or DUID and host name of the device that received the IP address, the Grid member that supplied the lease, and the start and end dates of the lease.

You can view current leases and lease history in the **Data Management** -> **DHCP** -> **Leases** tab in Grid Manager. To view lease history, you must first enable lease logging at the Grid level. For information, see [Configuring DHCP Logging](#) on page 815 and [Configuring the Lease Logging Member](#) on page 815. You can also export the DHCP lease history log in CSV format for archival and reporting purposes.

In the **Leases** tab, you can do the following:

- View current leases. For information, see [Viewing Current Leases](#) on page 946.
- View detailed information about a specific lease. For information, see [Viewing Detailed Lease Information](#) on page 948.
- View historical lease records. For information, see [Viewing Lease History](#) on page 949.
- View detailed information about a lease event. For information, see [Viewing Lease Event Detailed Information](#) on page 949.
- Export current leases and lease history logs. For information, see [Exporting Lease Records](#) on page 950.
- Clear leases. For information, see [Clearing Leases](#) on page 950.

You can also use the filter and **Go to** functions in the lease panels to retrieve lease information for specific hosts, MAC addresses, and IP addresses. These capabilities are crucial for security auditing and for meeting new compliance regulations such as SOX and HIPAA. You can also sort the lease information by column.

VIEWING CURRENT LEASES

To view current IPv4 and IPv6 leases:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Grid Manager displays the following information:
 - **IP Address:** The IPv4 address or IPv6 prefix or address that the appliance assigned to a DHCP client for this lease.
 - **Protocol:** Indicates whether the lease is for an IPv4 or IPv6 address.
 - **Members/Servers:** The Grid member or Microsoft server (for IPv4 leases only) that granted the lease.
 - **MAC address:** The MAC address of the IPv4 DHCP client that received the lease for an IPv4 address.
 - **DUID:** The DHCP Unique Identifier (DUID) of the IPv6 DHCP client that received the lease for an IPv6 address.
 - **Host Name:** The hostname that the DHCP client sent with its DHCP request. For IPv4 leases, this field displays the hostname of the DHCP client. For IPv6 leases, this field typically displays the FQDN.
 - **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the DHCP client that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.
 - **State:** The binding state of the current lease. The lease state can be one of the following:
 - **Free:** The lease is available for clients to use.
 - **Active:** The lease is currently in use by a DHCP client.
 - **Static:** The lease is a fixed address lease.
 - **Expired:** The lease was in use, but the DHCP client never renewed it, so it is no longer valid.
 - **Released:** The DHCP client returned the lease to the appliance.

- **Abandoned:** The appliance cannot lease this IP address because the appliance received a response when pinging the address.
- **End:** The day, date, and time when the state of the lease ends.
- **Start:** The day, date, and time when the state of the lease starts.
- **Username:** Displays the name of the user who receives the lease for the IP address. The username enables you to differentiate between guest users and authenticated users. If you log in as an authenticated user, your username is whatever you choose when you log in. If you log in as a guest, your username is First: first_name Last: last_name.
For example, if your first name is John and last name is Doe and your username is jdoe, when you log in as an authenticated user, your username is jdoe. If you log in as a guest user, your username is First: John, Last: Doe.
- **ClientID:** The DHCP client identifier (option 61) in an IPv4 lease. The client sends the client identifier as option 61 in the DHCP DISCOVER and REQUEST packets, as described in *RFC2132, DHCP Options and BOOTP Vendor Extensions*. The client identifier is either the MAC address of the network interface card requesting the address or any string uniquely identifying the client. This field is not displayed by default.

Note: The dates and timestamps in the **Leases** tab are determined by the time zone setting of the admin account that you use to log in to the appliance.

You can display the following discovered data for IPv4 leases:

- **Last Discovered:** The timestamp when the IP address was last discovered. This data is read-only.
- **OS:** The operating system of the detected host or virtual entity. The OS can be one of the following:
 - **Microsoft** for all discovered hosts that have a non-null value in the MAC addresses using the NetBIOS discovery method.
 - A value that a TCP discovery returns.
 - The OS of a virtual entity on a vSphere server.
- **NetBIOS Name:** The name returned in the NetBIOS reply or the name you manually register for the discovered host.
- **Discovered Name:** The name of the network device associated with the discovered IP address.
- **Discoverer:** Specifies whether the IP address was discovered by a PortIQ or NIOS discovery process.

You can do the following in this tab:

- Sort the data in ascending or descending order by column.
- View the lease detailed information of a current lease by selecting the check box of the lease, and then clicking the Open icon.
- Change a current lease state to **Free** by selecting the check box of a current lease, and then clicking the Delete icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print and export the data in this tab.

VIEWING DETAILED LEASE INFORMATION

You can view detailed information about a specific lease.

To view detailed information of a specific lease:

1. From the **Data Management** tab, select **DHCP** tab -> **Leases** -> **Current Leases** -> *lease* check box, and then click the **Lease Details** icon.
or
From the **Data Management** tab, select the **IPAM** tab, drill down to the IP Map, IP List, or IP address panel, and then click **Lease Details** from the Toolbar.
 2. In the *Lease Detailed Information* viewer, Grid Manager displays the following for each type of lease:
For IPv4 leases, it displays the fields **Member**, **MAC address**, **Host**, **Start**, **End**, **Binding State**, **Username**, **Binding State**, as described in [Viewing Current Leases](#) on page 946, plus the following information:
 - **Lease Issue:** The date and time when the lease was issued. Displayed in the lease event details report only.
 - **Event:** The action taken. This can be one of the following: **Issued**, **Renewed**, **Freed**, or **Abandoned**. Displayed in the lease event details report only.
 - **Served by:** The member that provides DHCP services to the lease.
 - **Next Binding State:** The subsequent binding state when the current lease expires. The lease state and the next binding state can be one of the following:
 - **Free:** The lease is available for clients to use.
 - **Active:** The lease is currently in use by a DHCP client.
 - **Static:** The lease is a fixed address lease.
 - **Expired:** The lease was in use, but the DHCP client never renewed it, so it is no longer valid.
 - **Released:** The DHCP client returned the lease to the appliance.
 - **Abandoned:** The appliance cannot lease this IP address because the appliance received a response when pinging the address.
 - **Billing Class:** The billing class of the lease.
 - **Option 82 Agent ID:** The agent ID of the relay agent filter (option 82). A relay agent can append DHCP option 82, relay agent information, to a message that it forwards from a DHCP client to a DHCP server.
 - **Option 82 Circuit ID:** The circuit ID of the relay agent filter (option 82).
 - **Option 82 Remote ID:** The remote ID of the relay agent filter (option 82).
 - **Option:** Agent circuit ID and remote ID data sent by a DHCP relay agent in all DHCP options.
 - **UID:** (User ID) The client identifier that the DHCP client sends the appliance (in DHCP option 61) when it acquires the lease. Not all DHCP clients send a UID.
 - **TSFP:** (Time Sent From Partner) The time—from the point of view of a remote DHCP failover peer—when the current lease state ends.
 - **CLTT:** (Client Last Transaction Time) The time of the last transaction with the DHCP client for this lease.
 - **TSTP:** (Time Sent To Partner) The time—from the point of view of the local DHCP failover peer—that the current lease state ends.
- For IPv6 leases, it displays most of the same fields as IPv4 leases, plus the following information:
- **DUID:** The DUID of the IPv6 DHCP client that received the lease for an IPv6 address.
 - **Prefix Bits:** The prefix length.
 - **Preferred Lifetime:** The length of time that a valid address is preferred. A preferred address can be used with no restrictions. When this time expires, the address becomes deprecated.

VIEWING LEASE HISTORY

To view lease history:

- From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Lease History**.

Grid Manager displays a table of historical leases that have been archived in the system. You can export the information in the lease history table. You can also search by the IP address or MAC address of the lease. Grid Manager displays the following read-only information:

- **Lease Issue:** The date and time when the lease was issued.
- **Protocol:** Indicates whether the lease is for an IPv4 or IPv6 address.
- **IP Address:** The IPv4 address or IPv6 prefix or address of the lease.
- **MAC Address:** The MAC address of the IPv4 lease.
- **DUID:** The DUID of the DHCP client that received the lease for an IPv6 address.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the DHCP client that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.
- **Host Name:** The host name that the DHCP client sent to the appliance.
- **Action:** This can be one of the following: Issued, Renewed, Freed, or Abandoned.
- **User Name:** The name of the user who received the lease for the IP address.
- **Start:** The start date of the lease.
- **Stop:** The end date of the lease.
- **Member/Server:** The DHCP member or Microsoft server that granted the lease.
- **Member IP Address:** The IP address of the DHCP member that granted the lease.

You can do the following in this section:

- View the lease event detailed information of a historical lease by selecting the check box of a lease, and then clicking the Open icon.
- Print or export the information in this section.

VIEWING LEASE EVENT DETAILED INFORMATION

You can view detailed information about a historical lease record by clicking the lease in the **Data Management** tab -> **DHCP** tab -> **Leases** tab -> **Lease History**. Grid Manager displays the event, the date and time when the event occurred, plus detailed information about the historical lease record. For information about the fields, see [Viewing Detailed Lease Information](#) on page 948.

You can also export and print the information in this panel. For information, see [Exporting Lease Records](#) on page 950.

EXPORTING LEASE RECORDS

The DHCP lease history log holds a maximum of 100,000 entries. After that maximum is reached, the appliance begins deleting entries, starting with the oldest. To archive DHCP lease history logs, you can export them and save them as CSV (comma separated variables) files. You do not need to export the entire log. You can selectively export a section of the log, such as the lease events for a single day.

As a conservative approach to archiving DHCP lease data, Infoblox recommends exporting the log on a daily basis, perhaps through API (application programming interface) scripting. By exporting the daily log entries every day over a certain period of time and then opening the exported files with a spreadsheet program, you can see the number of entries for each day. You can then estimate how often you need to export the log to ensure that you save all of the entries before the log fills up (at 100,000 entries). As a result, you might discover that you need to export the log more or less frequently than once a day to archive all the records.

A limited-access admin group can view and export the DHCP lease history if it has read-only permission to the DHCP lease history. For information on setting permissions for the DHCP lease history, see [Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories](#) on page 213. In addition, you can export the displayed DHCP current lease information or you can export them to a CSV file.

To export displayed current lease information:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Click the Export icon and select **Export visible data**. For more information on how to export, see [Exporting Displayed Data](#) on page 91.

To export DHCP current lease information to a CSV file:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Click the Export icon and select **Export data in Infoblox CSV Import format**. For more information on how to export, see [Exporting Data to Files](#) on page 89.

To export a lease history log:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases** or **Lease History**.
2. Click the Export icon and select.
3. In the *Export* dialog box, click **Start**
4. Click **Download** when the export is complete. Ensure that you turn off the pop-up blocker in your browser.
5. In the *File Download* dialog box, select the appropriate action to either open or save the CSV file.

CLEARING LEASES

You can clear active leases for which you have read/write permission. When you clear an active lease, its IP address becomes available and its status changes to “Free”. To clear an active lease:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Click the check boxes beside the IP addresses of the leases you want to clear, and then click the Clear Lease icon. Grid Manager clears the selected leases. You can view information about a cleared lease, by selecting it in the Lease History panel and clicking the Edit icon.



PART 6 MANAGING MICROSOFT WINDOWS SERVERS

This section describes how you can centrally manage Microsoft Windows® DNS and DHCP servers from Grid Manager. You can synchronize your DNS and DHCP data from the Microsoft servers to the Grid, and then use IPAM tools to facilitate DHCP and DNS configuration and data management. This section includes the following chapters:

- [Chapter 32, *Managing Microsoft Windows Servers*](#), on page 953
- [Chapter 33, *Managing Microsoft DNS Services*](#), on page 967
- [Chapter 34, *Managing Microsoft DHCP Services*](#), on page 983



Chapter 32 Managing Microsoft Windows Servers

This chapter explains how to configure Grid members to manage Microsoft Windows DNS and DHCP servers from Grid Manager. It includes the following sections:

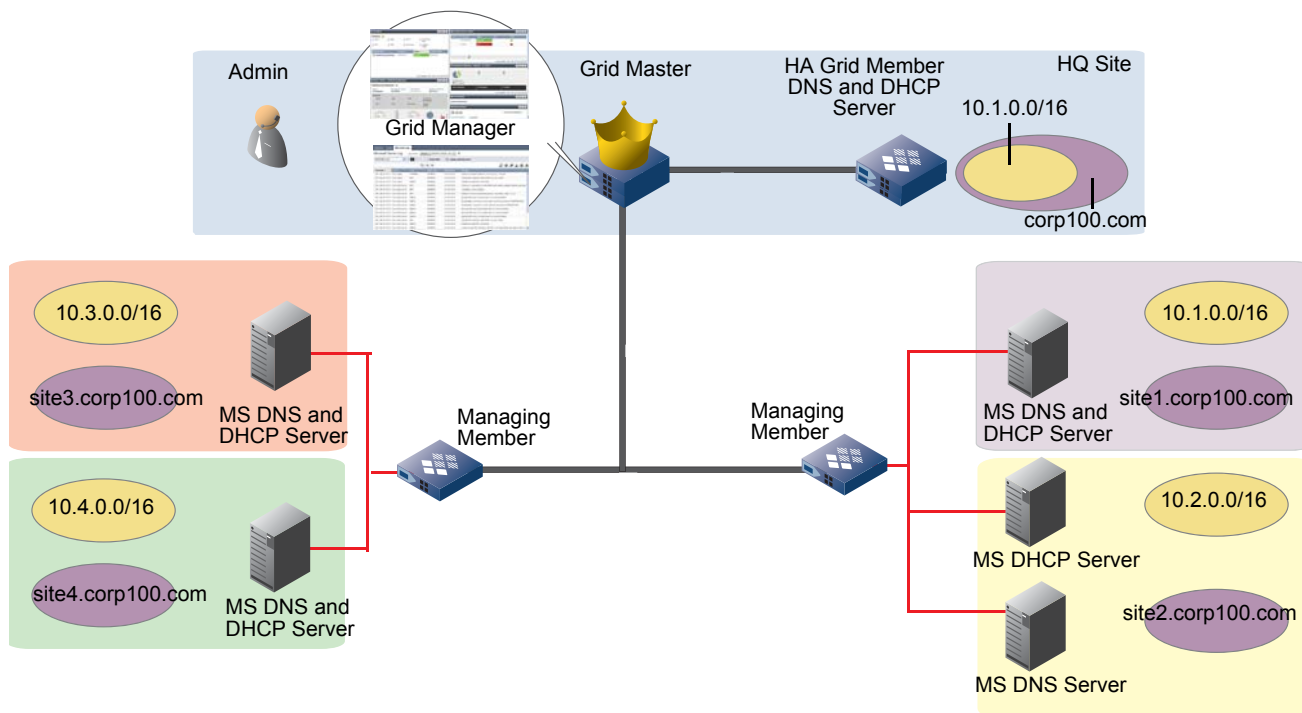
- [*About Managing Microsoft Servers*](#) on page 954
 - [*Requirements*](#) on page 955
 - [*Deployment Guidelines*](#) on page 956
- [*Configuring Members to Manage Microsoft Servers*](#) on page 957
 - [*Setting Microsoft Server Credentials*](#) on page 957
 - [*Configuring a Managing Member*](#) on page 958
- [*Managing Microsoft Servers*](#) on page 961
 - [*Setting Microsoft Server Properties*](#) on page 961
 - [*Changing the Managing Member or Management Mode*](#) on page 961
 - [*Backing Up Synchronized Data*](#) on page 962
 - [*Disabling Synchronization*](#) on page 962
 - [*Removing a Managed Microsoft Server*](#) on page 962
- [*Monitoring Managed Microsoft Servers*](#) on page 962
 - [*Viewing the Status of Servers*](#) on page 963
 - [*Viewing Detailed Status Information*](#) on page 964
 - [*Viewing Synchronization Logs*](#) on page 965

ABOUT MANAGING MICROSOFT SERVERS

You can configure Grid members to manage Microsoft Windows DNS and DHCP servers, and synchronize their DNS and DHCP data to the Grid database, so you can view and optionally, manage the data from Grid Manager. After the data is synchronized, you can use the IPAM tools of Grid Manager to simplify DNS and DHCP configuration and troubleshooting. You can use Smart Folders to organize your data, and monitor your networks and Microsoft servers from the Dashboard. In addition, you can control the DNS and DHCP services of the Microsoft servers from Grid Manager and configure server properties as well.

[Figure 32.1](#) illustrates a Grid that includes a member that provides DNS and DHCP services, and two other members that manage multiple Microsoft DNS and DHCP servers. Assuming the admin has the appropriate permissions, the admin can centrally manage the Microsoft DNS and DHCP servers and Infoblox DNS and DHCP server from a single interface, Grid Manager.

Figure 32.1 Managing Microsoft and Infoblox DNS and DHCP Servers from the Grid Master



You do not have to configure or install any application on the Microsoft servers for the Grid members to communicate with the servers. Infoblox uses MS-RPC (Microsoft Remote Procedure Calls) to manage Microsoft servers.

A Grid member can manage a Microsoft server in either of two modes, read-only or read/write. In read-only mode, the Grid member synchronizes data from the Microsoft server to the Grid so admins can use Grid Manager to view the synchronized data, but not update it. Read/Write mode allows admins to update the synchronized data as well. Updates from Grid Manager are then synchronized to the Microsoft server, and updates from the Microsoft server are synchronized to the Grid.

Configuration changes and data synchronized from the Grid to the Microsoft server apply immediately after the synchronization. You do not have to restart the Microsoft server or for DNS, reload the zones.

Note that due to a field length limit set on the Microsoft DHCP server, after you synchronize DHCP data on the Microsoft server, the “Comment” and “Description” fields for a fixed address and reservation can display only up to 128 characters even though NIOS allows up to 256 characters for these fields.

Requirements

A Grid member must have a Microsoft Management license installed to manage a Microsoft server. The license allows the member to synchronize data with Microsoft servers. It also activates the tabs, dialog boxes and other elements in Grid Manager that you need to manage a Microsoft server.

Note that If you do not see the Microsoft Servers tab after you add a member that has a Microsoft Management license, you might have to restart the Grid Master to view the tab and to manage Microsoft DNS and DHCP servers in the Grid.

Supported Windows Versions

Infoblox Grid members can manage Microsoft servers that support the following Windows versions:

OS	Levels	Platforms
Microsoft Windows 2003 Standard and Datacenter	SP2	32 bits
Microsoft Windows 2003 R2 Standard and Datacenter	Initial Release	32 bits, 64 bits
Microsoft Windows 2008 Standard and Datacenter	SP2	32 bits, 64 bits
Microsoft Windows 2008 R2 Standard and Datacenter	Initial Release	64 bits
Microsoft Windows 2012 Standard and Datacenter	Initial Release	64 bits
Microsoft Windows 2012 R2 Standard and Datacenter	Initial Release	64 bits

Grid members check the Windows version of the Microsoft servers before each synchronization. If a Microsoft server reports an unsupported version before a synchronization, the member logs an error and the synchronization fails.

Note that some Windows versions require certain updates and hotfixes installed, so the Microsoft server can synchronize with the Grid member. Following are the current requirements:

- Windows Server 2003, Enterprise x64 Edition requires the installation of security update 935966.
- Windows Server 2008 R2 requires the hotfix referenced in the Knowledge Base article 981776.
- Windows Server 2008-based DNS servers might not display delegations for reverse lookup zones. For information about this issue, including the available hotfix, refer to Knowledge Base article 958190.
- Windows Server 2012 and Windows Server 2012 R2 require remote workgroup servers for remote management. For more information, see [Adding the Remote Workgroup Server to Server Manager](#) on page 957.

For information about the updates, enter their IDs in the Search field of the Microsoft Support website at <http://support.microsoft.com>.

Administrative Permissions

By default, only superusers can configure Grid members to manage Microsoft servers. Superusers can give limited-access users read-only or read/write permission to Microsoft servers. Read-only permission allows admins to view the properties and data of a Microsoft server from Grid Manager. Write permission is required to configure Grid members to manage Microsoft servers, edit their properties, and start or stop their DNS and DHCP services. For additional information, see [Administrative Permissions for Microsoft Servers](#) on page 197.

Note that to view and manage the DNS and DHCP data synchronized from Microsoft servers, admins must have permissions to the applicable DNS and DHCP resources. For example, to view DNS zones synchronized from Microsoft servers, admins must have read-only permission to zones, and to edit the zones, admins need read/write permission to them. Similarly, to view DHCP ranges synchronized from Microsoft servers, admins must have read-only permission to DHCP ranges, and to edit the DHCP ranges, admins need read/write permission to the DHCP ranges. For information, see [Administrative Permissions for DNS Resources](#) on page 199 and [Administrative Permissions for DHCP Resources](#) on page 205.

The administrative permissions on the Grid are different from those on the Microsoft server. These permissions are independent of each other and are not synchronized.

Deployment Guidelines

Following are some recommendations and considerations when configuring Grid members to manage Microsoft servers:

- Infoblox recommends that you schedule the initial synchronization at a time when your network is less busy, especially if you are synchronizing a large amount of data. In addition, if a Microsoft server reconnects after being disconnected for a long period of time, it could synchronize a significant amount of data and this could impact the Grid Master performance.
- vNIOS Grid members and Grid members running on Infoblox-250, 250-A, Trinzic 100 and Trinzic 810 appliances do not support being configured as managing members.
- The managing member must be close, in terms of network hops, latency and bandwidth, to the Microsoft servers that it manages. This will help reduce the synchronization time and potential retries due to network delays.
- Although a Grid member that manages Microsoft servers can run other protocols and services, to optimize performance, Infoblox recommends that you configure one or more members solely for managing Microsoft servers.

If you are considering running other protocols and services on a managing member, consider using a member that is running on a platform other than the Infoblox-550-A.

- Grid members connect to Microsoft servers using RPC calls over TCP/IP. You must adjust your firewall policies to allow traffic between the managing Grid member and its assigned Microsoft servers. Grid members use the VIP as their source port. In Windows Server 2003, RPC uses the dynamic port range 1025-5000, by default. In Windows Server 2008, RPC uses the dynamic port range 49152-65535, by default. You can reduce the number of available ports as follows:
 - In Windows Server 2003, use the `rpcfcg.exe` tool. For information, refer to <http://support.microsoft.com/kb/908472>.
 - In Windows Server 2008 and later, use the `netsh` tool. For information, refer to <http://support.microsoft.com/kb/929851>.

The minimum number of ports required in the range is 255.

Note that TCP ports 135 and 445 must be open on the Microsoft server, in addition to the dynamic port range. Ports 135 and 445 are used by the port mapper interface, which is a service on the Microsoft server that provides information to clients on which port to use to connect to a specific service, such as the service that allows the management of the DNS service.

- The capacity of the managing member must be greater than or equal to the sum of all its assigned Microsoft servers.
- The capacity of the Grid Master must be greater than or equal to the sum of all managed Microsoft servers
- A Microsoft server can synchronize its data to only one network view, and for DNS data, only one DNS view.

- Multiple Microsoft servers can synchronize their data into the same network view and DNS view, unless there is a conflict in their data. For example, two Microsoft servers in different locations could serve the same private IP address space, such as 10.1.0.0/16, or serve reverse-mapping zones with the same name, such as 10.in-addr.arpa. Synchronizing their data to the same network view and DNS view would cause conflicts which result in the Grid member synchronizing the data of only one Microsoft server and logging an error for the other Microsoft server. In such situations, Infoblox recommends that you synchronize each Microsoft server to a different network view and DNS view to ensure that data from both servers are synchronized.

CONFIGURING MEMBERS TO MANAGE MICROSOFT SERVERS

You can manage Microsoft DNS and DHCP servers on any Grid member. To avoid performance issues, Infoblox strongly recommends that you do not configure Microsoft DNS and DHCP servers on the Grid Master and Grid Master candidate.

When an HA pair manages Microsoft servers, the active node handles synchronization. If an HA failover occurs during a synchronization, the failing node immediately aborts the synchronization. The new active node resumes the next synchronization. Changes that occurred on the Grid since the end of the last synchronization are lost.

For Microsoft DHCP failover, NIOS supports both the hot standby and load sharing modes in read-only mode on DHCP servers running Microsoft Windows 2012 and 2012 R2. For more information about Microsoft DHCP failover, refer to the Microsoft documentation.

Complete the following tasks to configure a Grid member to manage a Microsoft server:

1. On the Microsoft server, create a user account for the Grid member. For information, see [Setting Microsoft Server Credentials](#).
2. On the Grid Master, configure the managing member, as described in [Configuring a Managing Member](#).

Setting Microsoft Server Credentials

To enable a Grid member to synchronize data with a Microsoft server and control DNS and DHCP services, you must do the following on the Microsoft server:

1. Create a user account for the Grid member.
2. Grant the user account the necessary permissions.

You can either add the user account to the Administrators Group or add the user account to specific groups and explicitly set only the permissions necessary to access the DHCP and DNS services of the Microsoft server. The following sections provide general instruction on each method. Note that in order to remotely manage Microsoft Windows Server 2012 and 2012 R2, you must add a remote workgroup server to Server Manager. For more detailed information, refer to the Microsoft documentation or contact Microsoft Technical Support.

Adding the User Account to the Administrators Group

Adding the user account of the Grid member to the Administrators Group provides total control over the AD domain. Do one of the following:

- If the managed Microsoft server is a standalone server or a member server in a domain, open **Computer Management**, click **Groups**, and add the user account to the Administrators Group.
- If the managed Microsoft server is a domain controller, open **Active Directory Users and Computers**, select the domain name, click **Builtin**, and add the user account to the Administrators Group.

Adding the Remote Workgroup Server to Server Manager

To remotely manage Microsoft Windows Server 2012 and 2012 R2, you must add the remote workgroup server name to Server Manager. For information about how to add the remote workgroup, refer to the Microsoft article at <http://technet.microsoft.com/en-us/library/hh831453.aspx>.

Setting Specific Group Memberships and Permissions

If your security policy precludes adding user accounts to the Administrators group, you can add the user account to individual groups and grant only the required permissions. For guidelines and more information, see the following:

<http://support.microsoft.com/kb/325349>

<http://support.microsoft.com/kb/914392>

To add the user account of the Grid member to individual groups and grant specific permissions:

- To enable the member to synchronize DNS data with the Microsoft server, add its user account to the DnsAdmins Group.
- To enable the member to synchronize DHCP data with the Microsoft server, add its user account to the Dhcp Administrators Group.
- To enable the Grid member to monitor, start, and stop the DNS and DHCP services, grant the user account permissions on the Service Control Manager (SCM), as follows:
 1. Grant permissions to the SCM on each managed Microsoft server. For more information, refer to the section *DNS Server Service Permissions* at <http://technet.microsoft.com/en-us/library/gg638675.aspx>.
To find additional information, you can also search for "Least Privilege Setup" on the Microsoft sites.
 2. Grant permissions to the DNS and/or DHCP service on each managed server by doing one of the following:
 - Use the `sc` command line utility to remotely configure each managed DNS or DHCP server.
Note that you need to know the SID of the user account and its current permissions. You can retrieve the SID of the user account by using the `dsquery` and `dsget` commands.
 - Use the Domain Controller Policy editor to define a global policy that applies to all DNS or DHCP services running in a domain or on domain controllers. For additional information, refer to <http://support.microsoft.com/kb/324802>.

Configuring a Managing Member

When you configure a member to manage Microsoft servers, you must specify the following:

- The management mode of the Microsoft server. For information, see *Setting the Management Mode* on page 958.
- A network view, if there is more than one in the Grid, and a DNS view, if there is more than one in the network view. For information, see *Synchronizing to a Network View and DNS View* on page 959.

For the steps on configuring the managing member, see *Assigning Grid Members to Microsoft Servers* on page 959.

Setting the Management Mode

A Grid member can manage a Microsoft server in read-only mode, which is the default, or in read-write mode. In read-only mode, the Grid member copies the DNS and DHCP data from the Microsoft server to the Grid so Grid Manager admins can view the synchronized data. They cannot update the data, control the DNS and DHCP service of the Microsoft server, or configure any properties.

In read/write mode, Grid Manager admins are allowed to update the data of the Microsoft server. Therefore during each synchronization, the Grid member applies changes from the Grid to the Microsoft server and vice versa. Read/Write mode also allows admins to control DNS and DHCP services of the Microsoft server and configure some of their properties.

Note that the management mode of a Microsoft server is separate from the admin permissions that the appliance requires to access the Microsoft servers and DNS and DHCP resources. An admin must still have the applicable permissions to the Microsoft servers and DNS and DHCP resources they want to access. For information on admin permissions, see *Administrative Permissions for Microsoft Servers* on page 197.

Synchronizing to a Network View and DNS View

A Grid has one system-defined default network view, which contains a system-defined default DNS view. A network view is a single routing domain with its own networks and shared networks. A DNS view contains a version of DNS data that it can serve to specified clients. Admins can create additional network views and DNS views, according to their business needs. For information about network views, see [Configuring DHCP for IPv4](#) on page 843. For information about DNS views, see [Chapter 17, DNS Views](#), on page 601.

A Microsoft server can synchronize its data to only one network view and one DNS view. If a Grid contains the default network view and DNS view only, Grid Manager automatically assigns Microsoft servers to these default views. If a Grid has more than one network view, you must select one for the Microsoft server to synchronize its data; and if there are multiple DNS views, you must select one as well. You cannot modify the assigned network view or DNS view of a Microsoft server after its data has been synchronized. Instead, you must remove the Microsoft server and then add it again. For information about removing a server, see [Removing a Managed Microsoft Server](#) on page 962.

Microsoft servers do not support network views and DNS views. Therefore, network view and DNS view properties have no effect on the DNS and DHCP data synchronized from Microsoft servers.

Assigning Grid Members to Microsoft Servers

To configure a Grid member to manage one or more Microsoft servers:

1. From the **Grid** tab -> **Microsoft Servers** tab, click the Add icon.
2. In the *Add Microsoft Server(s)* wizard, complete the following and click **Next**:
 - **Managing Member:** Click **Select Member** and select the Grid member that manages the Microsoft servers. The default is the Grid Master.
 - **Synchronization Interval:** The default synchronization interval is two minutes. This is the time between the completion of one synchronization and the start of a new one. Synchronizing large data sets could take longer than the synchronization interval, causing a delay in the start of the next synchronization. For example, if the synchronization interval is two minutes but a synchronization takes five minutes, the time between the start of the first synchronization and the start of the next one is approximately seven minutes.
 - **Credentials to Connect to the Microsoft Server(s):** Enter the login name and password that the appliance uses to connect to the Microsoft servers. These must be the same as those you specified when you created the user account for the Grid member on the Microsoft servers. Note that you might have to specify the domain name and the user name in the following format: *domain_name\user_name*
 - **Manage Server(s) in:** Select the management mode, which is either **Read-only** or **Read/Write**.
 - **Synchronize Data into Network View:** This field appears only when there is more than one network view in the Grid. Specify to which network view the data from the Microsoft servers is synchronized.
 - **Synchronize DNS Data into DNS View:** This field appears only when there is more than one DNS view in the network view. Specify to which DNS view the data from the Microsoft servers is synchronized.
 - **Comment:** You can enter additional information about the servers.
 - **Disable:** Select this to disable the Microsoft servers. This allows you to preprovision the Microsoft servers and then enable them at a later time.
3. Do the following in the Managed Servers table:
 - **Name or IP Address:** Enter either the FQDN or IP address of the Microsoft server. In order for the member to resolve the FQDN of a Microsoft server, you must define a DNS resolver for the Grid member in the **DNS Resolver** tab of the *Member Properties* editor. Note that if the IP address of the Microsoft server is specified, then the DNS resolver must resolve it when the member and Microsoft server synchronize DHCP data only.
 - **DNS:** Select this option to enable the Grid member to manage the DNS service and synchronize DNS data with this server. Clearing this check box disables DNS service management and data synchronization. This allows you to preprovision specific Microsoft servers and then enable them at a later time.
 - **DHCP:** Select this option to manage the DHCP service of the Microsoft server and synchronize DHCP data with this server. Clearing this check box disables DHCP service management and data synchronization. This allows you to preprovision specific Microsoft servers and then enable them at a later time.

You can assign multiple Microsoft servers to a Grid member and test their connection the Grid member.

- Click the Add icon to add another Microsoft server.
- Select a Microsoft server and click the Test Microsoft Server icon to verify whether the appliance can successfully connect to the Microsoft server. The appliance displays the test results in the *Test Microsoft Server Results* dialog box.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click **Next**: Continue to the next step and define extensible attributes for the Microsoft servers. For information, see [Using Extensible Attributes](#) on page 332.

After you configure a Grid member to manage a Microsoft server, the member automatically connects to the Microsoft server and starts synchronizing data. You can then do the following:

- View the status of the servers in the *Microsoft Servers* panel, as described in [Monitoring Managed Microsoft Servers](#) on page 962. Newly added servers first display a status of **Connecting** as the Grid member contacts the Microsoft servers. The status changes to **OK** after the Grid member successfully connects to the Microsoft server.
- View the data synchronized from the Microsoft servers. To view DNS data, navigate to the DNS view you specified. For information, see [Viewing Zones](#) on page 653. To view DHCP data, navigate to the **Networks** tab of the network view that you specified. For information, see [Managing IPv4 DHCP Data](#) on page 841.
Network conditions and the amount of data can affect the synchronization time. Therefore, you might not be able to view all of the synchronized data immediately.
- Use Smart Folders to organize the Microsoft servers and their data. For example, you can create a folder for DNS zones and another folder for DHCP scopes synchronized from a Microsoft server. For information about Smart Folders, see [Chapter 3, Smart Folders](#), on page 139.
- Update the synchronized data. For information, see [Chapter 33, Managing Microsoft DNS Services](#), on page 967 and [Chapter 34, Managing Microsoft DHCP Services](#), on page 983

You can also use Global Search to search for synchronized data, such as zones and IP addresses. For information, see [Global Search](#) on page 59.

MANAGING MICROSOFT SERVERS

After you configure Grid members to manage Microsoft servers, you can set certain properties and manage the servers as follows:

- Set server properties, as described in [Setting Microsoft Server Properties](#).
- Change the managing member or the management mode, as described in [Changing the Managing Member or Management Mode](#) on page 961.
- Back up the synchronized data, as described in [Backing Up Synchronized Data](#) on page 962.
- Disable synchronization with a Microsoft server, as described in [Disabling Synchronization](#) on page 962.
- Remove a Microsoft server, as described in [Removing a Managed Microsoft Server](#) on page 962.

Setting Microsoft Server Properties

You can modify any of the Microsoft server properties you previously configured, except for the network view and DNS view. You can also set certain properties, including the logging level, extensible attributes, and administrative permissions. Extensible attributes and permissions apply to the data only when they are managed from Grid Manager. Extensible attributes and permissions are not synchronized to the Microsoft server.

To set the properties of a Microsoft server:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> *ms_server* check box, and click the Edit icon.
2. In the *Microsoft Server Properties* editor, you can set properties in the following tabs:
 - **General:** Modify the settings described in [Assigning Grid Members to Microsoft Servers](#) on page 959.
 - **Logging:** Select a logging level for the Microsoft server log.
 - **Low:** Logs only error messages.
 - **Normal:** Logs warning and error messages.
 - **High:** Logs warning, error and information messages.
 - **Debug:** Logs messages about all events associated with synchronization.
 See [Viewing Synchronization Logs](#) on page 965 for a description of each level.
 - **Extensible Attributes:** Define extensible attributes for the server. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** Define administrative permissions that apply to the server. For information see [About Administrative Permissions](#) on page 160.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can edit the General and Logging properties of multiple Microsoft servers at the same time by selecting the Microsoft servers and clicking the Edit icon. When Grid Manager displays the *Microsoft Server Properties* editor, it displays the values that the Microsoft servers have in common. If a property has multiple values, it indicates this. You can then change any of the values and when you click **Save**, Grid Manager applies your changes to all the selected Microsoft servers.

Changing the Managing Member or Management Mode

You can change the managing member and the management mode of a Microsoft server.

If you change the managing member, the previous member aborts any ongoing synchronization, and the newly assigned member resumes the synchronization process.

If you change the management mode of a Microsoft server from read/write to read-only, the Grid member reverts any changes that were made from Grid Manager since the last synchronization. For example, an admin adds a network and DHCP range for a scope. If another admin changes the management mode of the Microsoft server to read-only before the next synchronization, the Grid member deletes the network and DHCP range at the next synchronization.

To change the member or management mode:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> *ms_server* check box, and click the Edit icon.
2. In the *Microsoft Server Properties* editor, select the **General** tab and do any of the following:
 - **Managing Member:** Click **Select Member** and select another Grid member .
 - **Manage Server(s) in:** Select either **Read-only** or **Read/Write**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Backing Up Synchronized Data

When you back up the Grid, it includes all managed Microsoft data. If you restore a backup, the data is restored on the Grid only. It is not synchronized to the Microsoft servers. When the Grid member synchronizes the data after the restore operation, it overrides the data on the Grid with the data from the Microsoft servers. For information about backing up and restoring data, see [Chapter 9, Managing NIOS Software and Configuration Files](#), on page 405.

Disabling Synchronization

When you set the disable option, the Grid member completes any on-going synchronization and does not start a new one. Setting this option only affects data synchronization and does not affect the operations of the Microsoft server. Synchronization resumes when the Microsoft server is re-enabled.

To disable a Microsoft server:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> *ms_server* check box, and click the Edit icon.
2. In the **General** tab, select the **Disable** option.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Removing a Managed Microsoft Server

When you remove a Microsoft server from the Grid, the managing member stops any on-going synchronization and does not start a new one. If the Microsoft server served DNS, the synchronized DNS data remains unchanged in the Grid. If the Microsoft server served DHCP, then Grid Manager deletes all the DHCP ranges, leases, and fixed addresses associated with the server. It also deletes networks that were assigned only to the Microsoft server. It does not delete a network if it was assigned to other Microsoft servers as well.

Removing a managed Microsoft server from the Grid does not affect the operations of the Microsoft server.

To remove a managed server:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> *ms_server* check box, and click the Delete icon.
2. When the *Delete Confirmation* dialog box appears, click **Yes**.

For information about how removing a Microsoft server affects the synchronized DNS and DHCP data, see [Disabling and Removing Microsoft DNS Servers](#) on page 982 and [Disabling and Removing Microsoft DHCP Servers](#) on page 1000.

MONITORING MANAGED MICROSOFT SERVERS

You can monitor the status of managed Microsoft servers from the Dashboard and from various panels in the **Grid** tab. Grid Manager also maintains a log for each managed Microsoft server. You can monitor Microsoft servers and their services as follows:





- You can view the *Microsoft Servers Status* widget on the Dashboard. For information, see [Microsoft Servers Status Widget](#) on page 131.
- You can view the status of Microsoft servers. For information, see [Viewing the Status of Servers](#) on page 963.
- You can view the logs of the Microsoft servers. For information, see [Viewing Synchronization Logs](#) on page 965.

Viewing the Status of Servers





You can view details about the managed Microsoft servers by navigating to the **Grid** tab -> **Microsoft Servers** tab.

For each Microsoft server, the panel displays the following by default:

- **Name:** The FQDN of the Microsoft server
- **Status:** The connection status, which can be one of the following:
 - **Running:** The Grid member is connected to the Microsoft server.
 - **Connecting:** The Grid member is connecting to the Microsoft server.
 - **Error:** The Grid member failed to connect to the Microsoft server. Check the Microsoft log for any messages to determine the reason for the failure.
 - **Unknown:** The Microsoft server is disabled. The Grid member does not try to connect to disabled servers.
- **IP Address:** The IP address of the Microsoft server
- **DNS:** The status of the DNS service on the Microsoft server. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The DNS service is functioning properly.
	Red	The DNS service is stopped.
	Yellow	The DNS service is starting or stopping.
	Gray	Management of the Microsoft DNS server is disabled.

- **DHCP:** The status of the DHCP service on the Microsoft server. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The DHCP service is functioning properly.
	Red	The DHCP service is stopped.
	Yellow	The DHCP service is starting or stopping.
	Gray	Management of the Microsoft DHCP server is disabled.

- **Comment:** Displays any comments that were entered for the Microsoft server.
- **Site:** Displays any values that were entered for this pre-defined attribute.

You can add the following columns for display:

- **Version:** The Windows version of the managed server.
- **Managing Member:** The hostname of the Grid member that manages the server.

You can click **Toggle Synchronization Status View** to display the synchronization status of each managed server. The **Status** column changes to **Synchronization Status** and there is an additional column, **Last Changed**.

- **Synchronization Status:** Displays the synchronization status as follows:
 - **Running:** The Microsoft server is synchronizing data with the Grid member.
 - **Connecting:** The Grid member is trying to connect to the server.

- **Error:** Synchronization failed between the member and server. You can check the messages in the Microsoft server log to determine the reason for the failure.
- **Last Changed:** Displays information about when the status was last updated for Microsoft DNS and DHCP services. It corresponds to the last time information was exchanged with the server.

You can also do the following:

- Add Microsoft servers.
 - Click the Add icon.
- Edit the properties of a Microsoft server.
 - Click the check box beside a server, and then click the Edit icon. For information, see [Setting Microsoft Server Properties](#) on page 961.
- Delete a Microsoft server.
 - Click the check box beside a server, and then click the Delete icon. For information, see [Removing a Managed Microsoft Server](#) on page 962.
- Manage DNS and DHCP services of a Microsoft server.
 - Click the check box beside a server, and then click the Manage Server Services icon to view the service status. You can mouse over the DNS and DHCP service icons and click the Start/Stop service icon to start or stop a service, or click the Edit Service icon to edit the service properties. For information about setting DHCP server properties, see [Setting Microsoft DHCP Server Properties](#) on page 999. For information about setting DNS server properties, see [Specifying Forwarders for Microsoft Servers](#) on page 982.
- View detailed server status information, as described in [Viewing Detailed Status Information](#) on page 964.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Export the list of Microsoft servers to a .csv file.
 - Click the Export icon.
- Print the list of Microsoft servers.
 - Click the Print icon.

Viewing Detailed Status Information

You can view more status information by selecting a server from the Microsoft Servers panel and clicking the Detailed Status icon. The Detailed Status panel displays the following information:

- **Synchronization Status:** The status icon indicates the synchronization status as follows:
 - Green: The Microsoft server is synchronizing data with the Grid member.
 - Red: Synchronization failed between the member and server. You can check the messages in the Microsoft server log to determine the reason for the failure.
- **Last Updated:** Displays information about when the status was last updated for Microsoft DNS and DHCP services. It corresponds to the last time information was exchanged with the server.
- **DNS Service Status:** For information about the status icons, see [Viewing the Status of Servers](#) on page 963.
- **DNS Service Status Last Updated:** The date and time of the last DNS service status update from the Microsoft DNS Server.
- **DHCP Service Status:** For information about the status icons, see [Viewing the Status of Servers](#) on page 963.
- **DHCP Service Status Last Updated:** The date and time of the last DHCP service status update from the Microsoft DHCP Server.

Note the following guidelines about status information:

- Grid Manager does not display any status information if there is no synchronization between DHCP and DNS.
- If the appliance has not received any information when the services are enabled, then the Synchronization Status icon is displayed in red, whereas the DNS and DHCP status icons are displayed in grey.

Viewing Synchronization Logs

Grid Manager maintains a synchronization log file for each Microsoft server managed by a Grid member. It logs events related to the synchronization process, depending on the logging level that you configured in the **Logging** tab of the *Microsoft Server Properties* editor described in [Setting Microsoft Server Properties](#) on page 961.

The log files are rotated and compressed once they reach 40MB.

To view the log file of managed Microsoft server:

1. From the **Administration** tab, select the **Logs** tab -> **Microsoft Logs** tab.
2. If there is more than one managed server in the Grid, you can select the Microsoft server whose logs you want to view.
3. The log file contains information related to the synchronization of the Microsoft DNS and DHCP data, as follows:
 - **Timestamp:** The date and time of the log message. The time zone is the time zone configured in the User Profile.
 - **Source:** Identifies the event that generated the message, such as a server synchronization or zone synchronization.
 - **Level:** Indicates the severity of the message, which can be one of the following:
 - **Debug:** Provides information about all events associated with synchronization.
 - **Information:** The Grid member is synchronizing with the Microsoft server and these messages provide normal status information.
 - **Warning:** The Grid member synchronized the data, but there was an issue, which is detailed in the Message section.
If the Grid member encounters an error during the synchronization, it skips the object with the error, logs the error in the Microsoft log, and continues to synchronize the rest of the data. The Grid member logs the error at each synchronization until you resolve the issue and it can synchronize the object successfully.
 - **Error:** The Grid member failed to synchronize an object, such as a DNS zone or DHCP scope, due to the error described in the Message section.
 - **Object Type:** The type of object that corresponds to the entry, such as FQDN or ADDRESS.
 - **Object Name:** The name of the object that corresponds to the entry
 - **Message:** Detailed information about the event.

You can also do the following in the log viewer:

- Toggle between the single line view and the multi-line view.
- Navigate to the next or last page of the file using the paging buttons.
- Refresh the view.
- Click the Follow icon to have the appliance automatically refresh the log every five seconds.
- Download the log.
- Clear the contents of the log.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Export or print the content of the log.



Chapter 33 Managing Microsoft DNS Services

This chapter provides guidelines for using Grid Manager to manage Microsoft DNS servers and for synchronizing DNS data between Microsoft servers and the Grid. It discusses some features of the Microsoft DNS servers only as they relate to the synchronization of data. Please review the Microsoft documentation for complete information about Microsoft DNS servers and their features.

In addition, if you encounter technical issues with your Microsoft DNS servers, contact Microsoft Technical Support or consult the Microsoft Support site at <http://support.microsoft.com/>. Some Windows versions require certain updates and hotfixes installed, so the Microsoft server can synchronize with the Grid member. For information about these requirements, see *Requirements* on page 955.

The topics in this chapter include:

- *Managing Microsoft DNS Servers* on page 968
 - *Synchronizing DNS Data* on page 968
 - *Synchronizing with Multiple Servers* on page 970
- *Managing Synchronized DNS Data* on page 970
 - *Adding Zones to Microsoft Servers* on page 971
 - *Setting Zone Properties* on page 971
 - *Deleting and Restoring Synchronized Zones* on page 973
 - *Managing Resource Records in Synchronized Zones* on page 973
- *Synchronizing Updates* on page 974
 - *Synchronizing Delegations* on page 977
 - *Synchronizing AD-Integrated Zones* on page 980
 - *Resolving Conflicts* on page 981
- *Viewing Members and Managed Servers* on page 981
- *Specifying Forwarders for Microsoft Servers* on page 982
- *Disabling and Removing Microsoft DNS Servers* on page 982

MANAGING MICROSOFT DNS SERVERS

After you configure a Grid member to manage a Microsoft DNS server, the Grid member connects to the Microsoft server and starts synchronizing DNS data from the Microsoft server to its database. First, it synchronizes the Microsoft server properties and its list of zones. Then it synchronizes each zone individually, including its properties and resource records.

The synchronization time varies, depending on different factors, such as the number of managed Microsoft servers and the amount of data being synchronized. The synchronized data is then replicated to the Grid Master through the Grid replication process.

If the server is managed in read/write mode, admins can update the synchronized DNS data, control the DNS service of the server, and specify forwarders for it as well.

Synchronizing DNS Data

Grid members synchronize the properties and resource records of the following types of DNS zones:

- Authoritative forward-mapping zones
- IPv4 and IPv6 reverse-mapping zones
- Stub zones
- Delegations
- Active Directory-integrated zones

Grid members synchronize most of the resource records supported by Microsoft servers, except for WINS, WINSR, and ATMA records. They synchronize all the resource records supported by Infoblox DNS servers, as well as unsupported records, such as ISDN and X25 records. You can view the unsupported records in Grid Manager and delete them, but you cannot edit them. Note that Grid Manager and Microsoft DNS servers display some resource records, such as SIG records, in a different format.

Grid members do not synchronize the following DNS zones supported by Microsoft servers:

Infoblox Terminology	Microsoft Terminology
Forwarding zones	Conditional forwarders
Cached zones	Stub zones
Root zone	Root zone (Dot zone)
0.in-addr.arpa	0.in-addr.arpa (0.0.0.0)
127.in-addr.arpa	127.in-addr.arpa (127.0.0.1 - loopback)
255.in-addr.arpa	255.in-addr.arpa (255.255.255.255 - broadcast)
TrustAnchors	Trust Anchors

You cannot use Grid Manager to create the unsupported zones and assign them to a Microsoft server. Any zone on the Grid that has the same name as a forwarding, cached or root zone on the Microsoft server is not synchronized. In addition, Grid members do not synchronize the contents of a zone if the Microsoft server is a secondary server.

Subdomains defined within a Microsoft DNS zone are not synchronized unless they contain at least one resource record. For example, in the corp100.com zone, any resource record defined in a subdomain of the corp100.com zone is synchronized. If the subdomain sub.corp100.com zone has no resource record, it is not synchronized.

The following zones and resource records are supported on Microsoft servers running Windows Server 2008 only. Therefore, Grid members can only synchronize these DNS zones and resource records with Microsoft servers running Windows Server 2008.

- IPv6 reverse-mapping zones

- Global Names zones
- DNAME records
- NAPTR records
- DNSSEC records

IDN Support for Synchronized DNS Data

Infoblox Grid supports IDNs for synchronized DNS data between the appliance and Microsoft servers. For more information about IDNs, see [Support for Internationalized Domain Names](#) on page 93.

The appliance stores IDNs in punycode and Microsoft server stores IDNs in \xyz format. Due to this difference at the DNS protocol level, IDNs are not allowed in a zone name when you configure NIOS (primary or secondary) and Microsoft (primary or secondary) servers. For information, see [IDN Support Limitations for Synchronized Data](#) on page 969. If synchronized data between NIOS and Microsoft servers contain IDNs, the IDNs are preserved on the primary server. When a Microsoft server is the secondary server for a zone, MMC (Microsoft Management Console) displays the zone content that contains IDNs in punycode only. Make sure that you use a zone name that complies with the DNS protocol when manually configuring an authoritative zone.

A Microsoft server serves a resource record that contains an IDN in \xyz format when it is configured as the primary server and NIOS as the secondary server. For example, use the \xyz\xyz\xyz.com representation on NIOS for 网络.com, a zone added on the Microsoft server.

You can add resource records that contain IDNs for the following configurations:

- NIOS is the primary server and Microsoft server is the secondary server: You can add records in IDNs or punycode. The appliance preserves IDNs and does not encode punycode to IDNs. Note that you cannot add a resource record that contains an IDN on the Microsoft server when it is configured as the secondary server.
- Microsoft server is the primary server and NIOS is the secondary server: You can add records on both NIOS and Microsoft servers. You can use IDNs or punycode. You can add IDN records on both servers in this scenario.

The following table summarizes how the servers display resource records that contain IDNs after synchronization:

Primary Server	Input	Secondary Server	Input	NIOS displays records in...	Microsoft server displays records in...
NIOS	Punycode	Microsoft	NA	Punycode	Punycode
NIOS	IDN	Microsoft	NA	IDN	Punycode
Microsoft	Punycode	NIOS	Punycode	Punycode	Punycode
Microsoft	IDN	NIOS	IDN	IDN	IDN

IDN Support Limitations for Synchronized Data

You cannot configure an authoritative zone and stub zone that contains IDNs for the following configurations:

- When NIOS is the primary server and Microsoft server is the secondary server.
- When Microsoft server is the primary server and NIOS is the secondary server.

Synchronizing with Multiple Servers

Because a Grid member can manage multiple Microsoft servers, it could potentially manage multiple servers assigned to the same zone. For example, a Grid member could manage a Microsoft server that is the primary server of a zone and one or more Microsoft servers that are secondary servers of the same zone. It could also manage multiple Microsoft servers that are secondary servers for the same zone.

If a Grid member manages the primary server and at least one secondary server of a zone, the Grid member always synchronizes DNS data with the primary server only. It never synchronizes data with the secondary server, even if the primary server fails.

If a Grid member manages several Microsoft servers that are secondary servers of the same zone, it synchronizes DNS data as follows:

- If each Microsoft server is assigned to a different DNS view, the Grid member synchronizes data with each one.
- If the Microsoft servers are synchronized to the same DNS view, the Grid member selects a principal server for synchronization purposes, as follows:
 - The first Microsoft server that is assigned as the DNS secondary server is designated principal server.
 - If the secondary servers are managed in read-only and read/write modes, the Grid member always selects a server that is managed in read/write mode.
 - If a Microsoft server fails three successive synchronizations, it loses its principal server status. The Grid Master checks the date that each server last became a principal server and selects the server that has not been the designated principal server the longest.

Note that a Grid member could fail to synchronize with a Microsoft server due to errors, such as a disabled account or an expired password. In these situations, the failure count is reset and is not increased. This prevents the Microsoft server from losing its master status to another Microsoft server that could experience the same errors.

When a zone is served by multiple Microsoft servers, the **MS Sync Server** column of the **Zones** tab shows which Microsoft server is actually performing the synchronization of that zone with the Grid.

MANAGING SYNCHRONIZED DNS DATA

When Grid members are configured to manage Microsoft servers in read/write mode, you can use Grid Manager to view, edit and delete the DNS data of those servers. You can add new zones and assign them to a Microsoft server. You can modify the properties of zones synchronized from the Microsoft server and edit their resource records as well. All updates are synchronized to the Microsoft servers at regular intervals.

The following sections provide guidelines for managing the zones and resource records served by Microsoft servers:

- [Adding Zones to Microsoft Servers](#) on page 971
- [Setting Zone Properties](#) on page 971
- [Deleting and Restoring Synchronized Zones](#) on page 973
- [Managing Resource Records in Synchronized Zones](#) on page 973

Synchronized zones also support the following features:

- You can import data to zones synchronized with Microsoft servers. Note that the import fails if you try to import unsupported records to a Microsoft zone. For information about importing records, see [Importing Zone Data](#) on page 631.
- You can copy records to and from zones synchronized with Microsoft servers. When copying records to a Microsoft zone, you can copy only those records that are supported by Microsoft servers. For information about copying records, see [Defining a Match Destinations List](#) on page 607.

Adding Zones to Microsoft Servers

From Grid Manager, you can create zones and assign Microsoft servers as their primary or secondary servers. The managing Grid member then synchronizes these zones to the appropriate Microsoft servers.

From Grid Manager, you can add the following types of zones to Microsoft servers:

- Authoritative forward- and reverse-mapping zones—For information, see [Configuring Authoritative Zones](#) on page 616.
- Forward- and reverse-mapping stub zones—For information, see [Configuring Stub Zones](#) on page 644.
- Delegations—For information, see [Configuring a Delegation](#) on page 638.

Note that you cannot add a zone on a Microsoft server and configure it to be served by an Infoblox Grid member. For example, on the Microsoft server, you cannot add a zone and assign a Grid member as its primary server and the Microsoft server as the secondary server. You must add such a zone from Grid Manager.

Following are guidelines for adding zones to a Microsoft server:

- The primary or secondary server of the zone must be a Microsoft server.
- If the primary server is a domain controller, you can enable the option to store the zone in Active Directory, making it an AD-integrated zone. Note that you can enable Active Directory integration only after the Microsoft server has been synchronized at least once because its AD ability is not known before the synchronization.
- You do not have to assign a Grid member as the primary or a secondary server of the zone. For example, a zone can have a Microsoft server as its primary server and an external secondary server.
- The zone must be in the same DNS view to which the DNS data of the Microsoft server was synchronized. You cannot add a zone served by the Microsoft server to a different DNS view.
- The zone does not inherit the properties from the Grid or from the DNS view. It uses the Infoblox-defined defaults. You can change the property values, as described in [Setting Zone Properties](#) on page 971.
- You can set certain zone properties that are not supported and synchronized to the Microsoft server. For example, you can define extensible attributes and administrative permissions. When you set these properties, they apply to the zones only when they are managed from Grid Manager.
- Infoblox does not support all the zone properties of a Microsoft DNS server. When a Grid member synchronizes zones that were created on Grid Manager to the Microsoft server, the zones contain default values for all unsupported properties.
- If you set the Disable option, the zone status is set to “Paused” on the Microsoft server. A zone in a “Paused” status is not served to DNS clients, nor is it available for zone updates.
- Setting the Disable option does not stop synchronization. Grid members synchronize disabled zones.
- The member learns the Windows version of the Microsoft server after its first successful synchronization. Certain zones and resource records are dependent on a specific Windows version. You cannot assign these zones to Microsoft servers whose versions are unknown or insufficient.
- If the member is a secondary server for a zone with a Microsoft primary server, the member obtains the zone data through DNS zone transfers from the Microsoft primary server; not through synchronizations. This ensures that the zone data is always current on the Infoblox secondary server, as it does not have to wait for synchronizations to update its data.

Setting Zone Properties

When the primary server of a zone is a Microsoft server, it does not inherit its properties from the Grid. Zones that are synchronized from a Microsoft server retain their original properties. Zones that Grid Manager admins create assume the Infoblox-defined default values.

To modify the properties of a synchronized zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *DNS_view* -> *zone* check box and click the Edit icon.
2. In the *Authoritative Zone* editor, you can do the following in each tab:

- **General:** You can add or edit comments, and set the Disable and Lock options. Setting the Disable option sets the status of the zone to “Paused” on the Microsoft server. Grid members synchronize disabled zones to Microsoft servers.
- **Name Servers:** You can modify the name servers assigned to the zone. For information, see [Assigning Zone Authority to Name Servers](#) on page 623.
- **Settings:** If the zone was synchronized from a Microsoft server, this tab displays the original settings from the Microsoft server. If the zone was created using Grid Manager, then it inherits the TTL values from the Grid. Note that these values might be different from those on the Microsoft server. To change any of these values, see [Configuring DNS Service Properties](#) on page 557.
- **Zone Transfers:** In this tab, you specify the servers to which zone transfers are allowed. For information about zone transfers, see [Enabling Zone Transfers](#) on page 583. Set the following parameters, depending on whether the primary or secondary servers of the zone are Infoblox or Microsoft DNS servers:
 - If the primary server is an Infoblox, Microsoft or external primary and the secondary servers are both Infoblox and Microsoft DNS servers, this tab displays two separate tables where you can specify zone transfer settings for the Infoblox DNS servers and the Microsoft DNS servers.

Zone Transfer Settings for Infoblox Members: Specify the settings as described in [Configuring Zone Transfers](#) on page 584.

Zone Transfer Settings for Microsoft Servers: Note that you cannot use a named ACL for access control though you can use individual ACEs. For information about named ACLs and access control, see [Configuring Access Control](#) on page 306. You can set access control for zone transfers for Microsoft servers to one of the following:

 - **None:** Does not allow zone transfers to any name server.
 - **Any:** Allows zone transfers to any IP address.
 - **Any Name Server:** Allows zone transfers to any name server in the Name Servers table.
 - **Address:** Allows zone transfers to the IP address that you specify.
 - If both the primary and secondary servers are Microsoft servers, the dialog box displays the **Zone Transfer Settings for Microsoft Servers** table only.
 - If no Microsoft servers are primary or secondary servers, then the dialog box displays the **Zone Transfer Settings for Infoblox Members** table only.
- **Updates:** In this tab, you specify whether the zone can accept dynamic DNS updates. For information about dynamic DNS updates, see [Chapter 20, Configuring DDNS Updates from DHCP](#), on page 689. If the primary server is a Microsoft server, regardless of the secondary servers, the **Updates** tab displays the following:
 - **Dynamic Updates:** Select one of the following:

None: The zone does not accept dynamic updates.

Secure Only: This appears only if the zone is AD-integrated. The zone accepts GSS-TSIG-signed updates only.

Nonsecure and Secure: The zone accepts both nonsecure and GSS-TSIG-signed updates.
- **Active Directory:**
 - **Automatically create underscore zones:** This option allows the appliance to create the following subzones that the DNS server must have to answer AD-related DNS queries:
 - `_msdcs.zone`
 - `_sites.zone`
 - `_tcp.zone`
 - `_udp.zone`
 - `domaindnszones.zone`
 - `forestdnszones.zone`

Note that these zones are automatically generated. You cannot edit these zones or import data into them. They cannot be modified, thus providing protection against forged updates.

- **Extensible Attributes:** Extensible attributes apply to the zones only when they are managed from Grid Manager. For information, see [Using Extensible Attributes](#) on page 332.
- **Permissions:** These permissions apply to Infoblox Grid Manager administrators only. For information, see [About Administrative Permissions](#) on page 160.

Deleting and Restoring Synchronized Zones

When you delete a synchronized zone from the Grid, Grid Manager moves the zone and its resource records to the Recycle Bin. It deletes the zone and its resource records from the Microsoft server at the next synchronization.

Note that if you delete a zone on Grid Manager and plan to add it back to the database with different properties or resource records, ensure that you wait until after the deletion is synchronized to the Microsoft server to add the new zone. Otherwise, if you delete a zone and add a new zone with the same name within a synchronization interval, Grid Manager will synchronize the zone properties and resource records from the Microsoft server to the newly added zone on Grid Manager.

If a zone has subzones, you can choose to remove them and their resource records or “reparent” them to the parent zone of the one you are removing. For information, see [Removing Zones](#) on page 634.

If you restore deleted zones from the Recycle Bin, the Grid member restores it on the Microsoft server as well. For information, see [Restoring Zone Data](#) on page 636.

Managing Resource Records in Synchronized Zones

From Grid Manager, you can add and edit resource records in zones served by Microsoft servers. For information about adding and managing resource records, see [Managing Resource Records](#) on page 660. You can also use IP Map and the IP List to track A, AAAA and PTR records that are synchronized from Microsoft servers. For information, see [Chapter 12, IP Address Management](#), on page 457.

Microsoft DNS servers support all the resource records supported by Infoblox DNS servers, except for hosts, bulk hosts and shared record groups. You cannot add these records to zones served by Microsoft servers or assign zones with these records to Microsoft servers.

Following are guidelines for adding and managing resource records in synchronized zones:

- Infoblox DNS servers support defining a naming policy for the hostnames of A, AAAA, MX, and NS records based on user-defined or default patterns. For information, see [Specifying Hostname Policies](#) on page 592. The hostname policy applies only when records are created from Grid Manager. Resource records that originate from the Microsoft server are synchronized to the Grid member even if they do not comply with the hostname policy of the Grid member. The policy is enforced only if you edit the resource record from Grid manager.
- When you create an A or AAAA resource record on the NIOS appliance with the option to automatically create the corresponding PTR record, Grid Manager uses the deepest reverse zone that can hold the record. For example, a Grid has the following reverse zones: 10.in-addr.arpa, 0.10.in-addr.arpa, and 0.0.10.in-addr.arpa. When you create the A record www A with the IP address 10.0.0.1, Grid Manager creates a PTR record in the zone 0.0.10.in-addr.arpa. If the deepest zone does not allow the creation of the PTR record, Grid Manager creates the A record, but not the PTR record, and displays a warning.
- You can add and edit DNAME records in a DNS zone assigned to a Microsoft server running Windows 2008 or Windows Server 2012. You cannot add or edit DNAME records in zones assigned to Microsoft servers running earlier Windows versions.
- You can disable synchronized resource records from Grid Manager. When you disable a resource record, it is removed from the Microsoft server at the next synchronization.
- If you add a resource record with invalid data from Grid Manager, such as a DNAME record with an alias name that has special characters, the invalid resource record is not synchronized to the Microsoft server and is eventually deleted from the Grid. The error is logged in the Microsoft log.
- If the zone of the resource record was created using Grid Manager, then it and all its resource records inherit their TTL values from the Grid. Note that these values might be different from those on the Microsoft server. You can change these values to match those on the Microsoft server. For information on changing these values, see [Configuring DNS Service Properties](#) on page 557.

- Grid Manager and Microsoft DNS servers display TXT records differently.

On Grid Manager, you enter the text string of TXT records as defined in RFC 1035. You can enter the following:

- A contiguous set of characters without spaces. If you enclose the characters in double quotes, Grid Manager displays the character string without the double quotes. For example, if you enter **"abcdef"**, Grid Manager displays **abcdef**.
- A string that contains any character, including spaces, enclosed in quotes.
 - If the string contains a quote ("), you must precede it with a backslash.
 - If you enter a text string with multiple spaces between each word and the string is not enclosed in double quotes, Grid Manager displays the text string with a single space between each word. For example, if you enter **text string**, the GUI displays **text string**. To preserve multiple spaces, enclose the string in double quotes.

Unlike on Microsoft DNS servers, you cannot enter a text string on multiple lines in Grid Manager. However, each contiguous set of characters or quoted string entered on Grid Manager is equivalent to a separate line entered on a Microsoft DNS server.

On a Microsoft DNS server, you can enter text without quotes and with each line on a separate line. Microsoft DNS servers then display the text in a brief format where the lines are separated by a comma and a space.

For example, if you enter the following in the **Text** field of the TXT Record wizard or editor on Grid Manager:

"this is a line""with another line""and a third one"

It is served by the Microsoft and Infoblox DNS servers as:

"this is a line""with another line""and a third one"

But it is displayed in the Microsoft DNS server as:

this is a line, with another line, and a third one

SYNCHRONIZING UPDATES

A Grid member synchronizes DNS data with each managed Microsoft server at regular intervals. Grid Manager admins with the applicable permissions can then update the synchronized DNS zones and resource records. During each synchronization, updates from Grid Manager are applied to the Microsoft server and updates from the Microsoft server are applied to the Grid as well. Note that the resource records are synchronized only if there is a change to the SOA record on either the Microsoft server or the Grid.

The following examples illustrate how Grid members synchronize DNS data:

- If a Microsoft server admin adds the `finance.corp100.com` zone, it is also added to the Grid after a synchronization.
- If a Grid Manager admin changes the A record of `admin.corp100.com` from `10.2.1.5` to `10.2.1.6`, the IP address of its corresponding A record on the Microsoft server is updated to `10.2.1.6`.
- If a Grid Manager admin deletes a DNS zone that is assigned to a Microsoft server, the corresponding zone on the Microsoft server is deleted as well in the next synchronization.

Because admins can update DNS data from the Microsoft server and from Grid Manager, conflicts can occur during synchronization. In addition, Microsoft servers and Infoblox DNS servers have some differences in the features they support and the way they handle certain zones and resource records.

The following guidelines describe how the Grid member resolves conflicts and handles any differences when DNS data is synchronized between a Microsoft server and the Grid.

- On Microsoft servers, users can enter FQDNs and labels using a mix of upper and lower case characters. The servers preserve the original case when they store the data. When the Grid member synchronizes data with the Microsoft server, it displays the data in lower case in Grid Manager and the Infoblox API. The case of the data is preserved as long as no change is made to the DNS zone or resource record. If a Grid Manager admin modifies a DNS zone or resource record, the next synchronization converts the object name to lower case on the Microsoft server.

- If a Microsoft server admin modifies an object that has a pending scheduled task and synchronization occurs before the scheduled task, the object is modified in both the Microsoft server and the Grid member. When the scheduled task executes at its scheduled time, it fails and an error message is logged in the audit log.
- A situation could arise where two Microsoft servers in different domains are primary servers for zones with the same name. For example, two reverse-mapping zones could be named 1.1.10.in-addr-arpa in two Microsoft servers managed by the same member. If the two Microsoft servers are synchronized to different DNS views, the Grid member synchronizes each one separately. If the Microsoft servers are synchronized to the same DNS view, then the Grid member synchronizes the zone with the first Microsoft server. During the synchronization with the second Microsoft server, the Grid member logs an error and does not synchronize the zone.
- The Grid member does not synchronize the naming policy configured on Microsoft servers. Zones and resource records that fail the policy check on Microsoft servers are reported in the synchronization log file.
- When you remove a Microsoft server that is assigned to a zone, the succeeding synchronization removes the zone from the Microsoft server.
- When a Microsoft server admin and a Grid Manager admin change the same object, the Grid member retains the version that exists on the Microsoft server. Following are some examples:

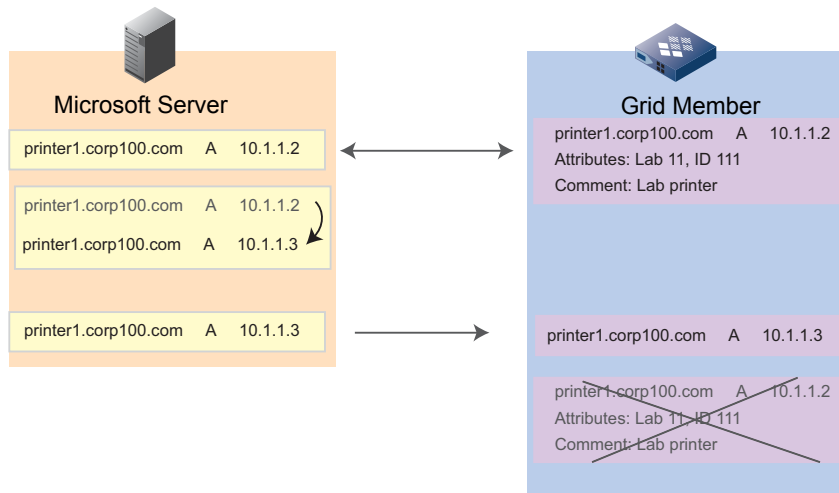
Table 33.1

Grid Manager Admin...	Microsoft Server Admin...	After Synchronization
Deletes the corp100.com zone	Updates the corp100.com zone	The corp100.com zone is created on the Grid with the updates and is assigned to the Microsoft server .
Changes the zone transfer settings of the sales.corp100.com zone.	Deletes the sales.corp100.com zone.	The sales.corp100.com is deleted from the Grid as well.

- Changing the name or IP address of a resource record on the Microsoft server effectively deletes the original resource record and creates a new record with the current information. During the synchronization, the Grid member also deletes the original record, including its associated properties, such as its extensible attributes and administrative permissions, and creates a new record.

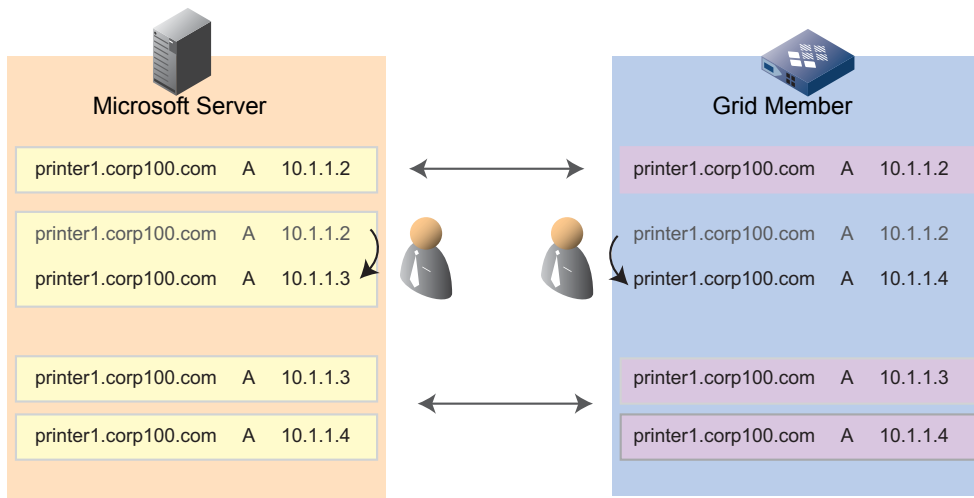
For example, as shown in [Figure 33.1](#), the A record for printer1.corp100.com is on both the Microsoft and Infoblox Grid member. On the Grid, the A record has extensible attributes and a comment. A Microsoft server admin changes the IP address of the A1 resource record from 10.1.1.2 to 10.1.1.3. On the Microsoft server, this is equivalent to deleting the A1 resource record with the IP address 10.1.1.2 and then adding a new A1 resource record with the IP address 10.1.1.3. When the data is synchronized, the Grid member deletes the original record with its extensible attributes and comments and creates a new A record with IP address 10.1.1.3.

Figure 33.1



- If a Microsoft server admin changes the IP address of a resource record and a Grid Manager admin changes the IP address of the same resource record, they are effectively deleting the record and each creating a new one. For example, as shown in [Figure 33.2](#), a Microsoft server admin changes the IP address of the A resource record for printer1.corp100.com from 10.1.1.2 to 10.1.1.3, and a Grid Manager admin changes the IP address of the same resource record to 10.1.1.4. When the data is synchronized, the Grid member deletes the A1 resource record with IP address 10.1.1.2 and creates an A resource record with IP address 10.1.1.3 and another A1 resource record with IP address 10.1.1.4.

Figure 33.2



- Grid members can synchronize classless IPv4 reverse-mapping zones from the Microsoft server to the Grid only if the zone prefix is in one of the following formats: `<subnet>/<subnet mask bit count>` or `<subnet>-<subnet mask bit count>`. For example, 128/26.2.0.192.in-addr.arpa. If the zone prefix is not in the specified format, the Grid member skips the zone and logs an error message. For information, see <http://technet.microsoft.com/en-us/library/cc961414.aspx>. Likewise, Grid Manager admins can add a classless IPv4 reverse-mapping authoritative or stub zone to a Microsoft server only if its prefix is in the specified format. For information about configuring classless IPv4 reverse-mapping zones in Grid Manager, see [Specifying an RFC 2317 Prefix](#) on page 618.

- Grid members synchronize DNS records that contain values that Infoblox does not support. Grid Manager admins can view these records, but they cannot edit or restore such records. For example, if a member synchronizes a NAPTR record that contains an unsupported value in the Service field, admins can view this record but they cannot edit or restore it, as long as it contains an unsupported value.
- When a Grid member synchronizes a zone from a Microsoft server to the appliance and that zone contains UTF-8 characters in the "Responsible Person" field, Grid Manager displays the "Responsible Person" value in the RNAME field of the SOA record of the zone. Note though that you cannot edit the SOA record if the RNAME field contains unsupported UTF-8 characters.
- Synchronizing a new zone from the Microsoft server to the Grid is a two-step process. First the zone name is synchronized, and then its properties and records are synchronized. The zone synchronization from the Microsoft server is not considered complete until both steps are done. NIOS drops any records that are created on the appliance for the synchronized zone before it is completely synchronized.

For example, the corp100.com zone is created on a Microsoft server and then synchronized to the NIOS appliance. If a NIOS admin creates a record, such as an A record, in the corp100.com zone before it is fully synchronized, the record is removed from the corp100.com zone. Ensure that both the zone and its contents are completely synchronized before you add a record to a zone on the NIOS appliance.

- When a Microsoft admin deletes a zone whose primary and secondary servers are Microsoft servers, the zone is deleted from the Grid after a synchronization. If the secondary Microsoft servers are managed by the member in Read-Write mode, the zone is removed from the secondary servers as well. But if some of the secondary Microsoft servers are managed by the member in Read-Only mode, then at the next synchronization the zone is recreated on the Grid with the Microsoft servers as the secondary servers and the masters defined for the zone as external primary servers.
- If you add Grid members or other Microsoft servers as secondary servers to a zone on the Microsoft server, NIOS does not automatically add them as Grid Secondary or Microsoft Secondary servers in the Name Servers tab of the zone after the synchronization. Instead, NIOS creates NS records for them in the zone.

Synchronizing Delegations

When a parent zone delegates a subdomain to one or more name servers, Infoblox DNS servers require the delegation name servers to also be authoritative for the subzone. Microsoft servers do not; they allow the delegation servers of a subzone to be different from its authoritative servers. Infoblox DNS servers support this configuration only if the primary server of the parent zone is a Microsoft server. This configuration is retained when delegations are synchronized from Microsoft servers to the Grid.

For example, as shown in [Figure 33.3](#), on a Microsoft server, corp100.com delegates sales.corp100.com to the name server ns1.corp100.com; but the authoritative server of sales.corp100.com is 2k3r264-2.infoblox.com.

Figure 33.3 Delegation Server and Authoritative Server for corp100.com

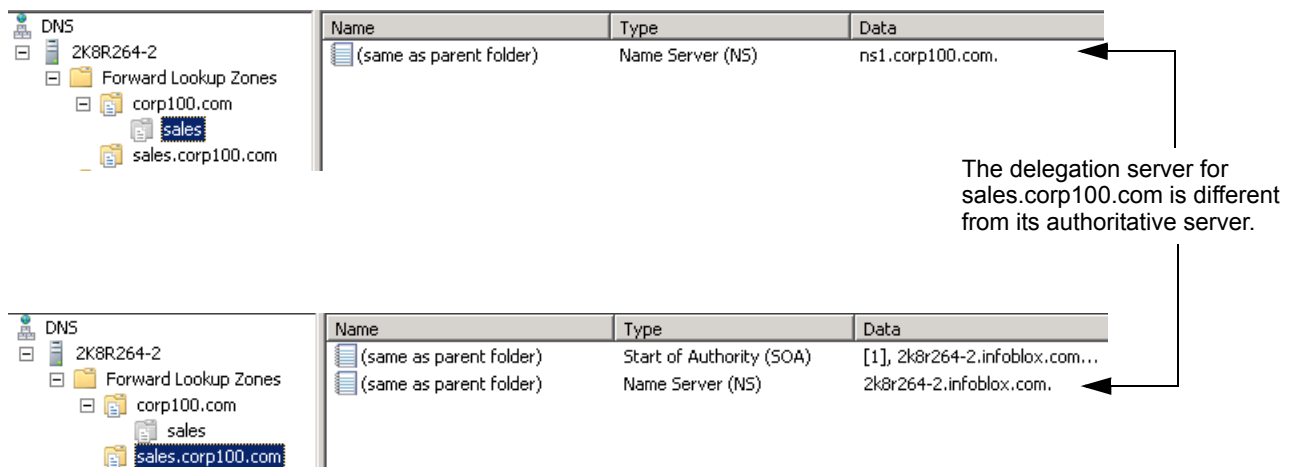







Figure 33.4 shows that after corp100.com and its subzone are synchronized to the Grid, corp100.com contains an NS record for sales.corp100.com and an A record for the delegation name server ns1.corp100.com. The *MS Delegation Addresses* column displays the IP address of the delegation server of the subzone sales.corp100.com.

Figure 33.4 corp100.com Synchronized to the Grid

corp100.com Zone   

Records

Quick Filter: | Filter On | [Show Filter](#) | [Toggle flat view](#)

Go to:  




Name	Type	Data	MS Delegation Addresses	Comment
	SOA Record	Serial: 10 MNAME: 2k8r264-2.intoblox.com RNAME: hostmaster@corp100.com Refresh: 900 Retry: 600 Expire: 86400 Negative Caching TTL: 3600		Auto-created by Add Zone
	NS Record	2k8r264-2.intoblox.com		Auto-created by Add Zone
ns1	A Record	10.2.3.4		Auto-created by Add Zone
sales	NS Record	ns1.corp100.com	10.2.3.4	Auto-created by Add NS

After the synchronization, you can add name servers for the delegation as follows:

1. Select the zone by navigating to the **Data Management** tab -> **DNS** -> **Zones** -> *parent_zone*.
2. Click the Add icon and select **Record** -> **NS Record**.
3. Complete the following and click **Next**:
 - **Name Server**: Enter the hostname you want to configure as the name server for the zone.
 - **Name**: Specify the name of the subzone. Note that you cannot change this name when you edit the record.
4. Enter the IP address of the name server.
5. Save the configuration.



NIOS adds an NS record for the new delegation server and synchronizes this update to the Microsoft server. In Figure 33.5, a new delegation server, ns2.corp100.com, was added.

Figure 33.5 NS Record for ns2.corp100.com

corp100.com Zone   

Records

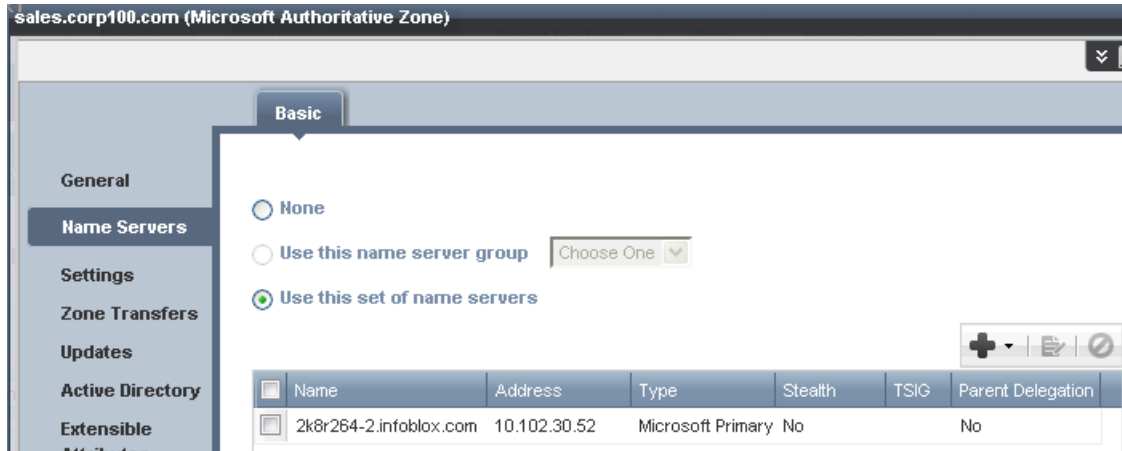
Quick Filter: | Filter On | [Show Filter](#) | [Toggle flat view](#)

Go to:  

Name	Type	Data	MS Delegation Addresses	Comment
	SOA Record	Serial: 11 MNAME: 2k8r264-2.intoblox.com RNAME: hostmaster@corp100.com Refresh: 900 Retry: 600 Expire: 86400 Negative Caching TTL: 3600		Auto-created by Add Zone
	NS Record	2k8r264-2.intoblox.com		Auto-created by Add Zone
ns1	A Record	10.2.3.4		Auto-created by Add Zone
ns2	A Record	10.1.2.3		Auto-created by Add Zone
sales	NS Record	ns1.corp100.com	10.2.3.4	Auto-created by Add NS
sales	NS Record	ns2.corp100.com	10.1.2.3	Auto-created by Add NS

When you navigate to the **Name Servers** tab of sales.corp100.com to view the authoritative name servers for the subzone, note that as shown in [Figure 33.6](#), the table displays 2k8r264-2.infoblox.com as the authoritative server for the subzone. The **Parent Delegation** column indicates if the FQDN and IP address of the authoritative name server for the zone matches the FQDN and IP address in the delegation zone's NS record. In the example, the authoritative name server 2k8r264-2.infoblox.com is different from the delegation name servers (ns1.corp100.com and ns2.corp100.com), so the column displays **No**.

Figure 33.6 Authoritative Name Server of sales.corp100.com



Note though that because Infoblox DNS servers require the delegation servers to also be authoritative for the subzone, if you add another authoritative name server to the subzone from Grid Manager, NIOS also adds it as a delegation server in the parent zone. For example, as shown in [Figure 33.7](#), when an admin adds the name server ns-100.corp100.com as an external secondary server for sales.corp100.com, NIOS automatically adds it as a delegation server by adding an NS record for it in the parent zone.

Figure 33.7 Adding Another Authoritative Server from Grid Manager

The screenshot displays the Grid Manager interface for a DNS zone named **corp100.com**. The **Basic** tab is selected, showing the **Name Servers** configuration. Two name servers are listed:

Name	Address	Type	Stealth	TSIG	Parent Delegation
2k8r264-2.infoblox.com	10.102.30.52	Microsoft Primary	No		No
ns-100.corp100.com	10.3.2.1	Ext Secondary	No		Yes

Below the name servers, the **Records** section is visible, showing a list of DNS records for the zone. The records include:

Name	Type	Data	MS Delegation Addresses	Comment
	SOA Record	Serial: 29 MNAME: 2k8r264-2.infobl RNAME: hostmaster@inf Refresh: 900 Retry: 600 Expire: 86400 Negative Caching TTL: 3600		Auto-created by Add Zone
	NS Record	2k8r264-2.infoblox.com		Auto-created by Add Zone
ns-100	A Record	10.3.2.1		Auto-created by Add Zone
ns1	A Record	10.2.3.4		Auto-created by Add Zone
sales	NS Record	ns-100.corp100.com	10.3.2.1	Auto-created by Add NS
sales	NS Record	ns1.corp100.com	10.2.3.4	Auto-created by Add NS

Synchronizing AD-Integrated Zones

An AD-integrated zone can be served by multiple domain controllers, and a Grid member can manage more than one of the domain controllers. If the domain controllers are configured to synchronize their DNS data to different DNS views, the Grid member synchronizes DNS data with each domain controller. If the domain controllers are configured to synchronize their DNS data to the same DNS view, the member selects a principal server for synchronization purposes and synchronizes data with that principal server only. The selection of the principal server is logged, as well as when it changes. The Grid member selects a principal server as follows:

- The first domain controller that is assigned as the primary server is designated principal server.
- If a domain controller fails three successive synchronizations, it loses its principal status. The Grid Master then checks the date that each domain controller last became a principal server and selects the one that has not been the designated principal the longest.
- If the domain controllers are managed in read-only and read/write modes, the Grid member always selects the domain controller that is managed in read/write mode.

When a zone is served by multiple Microsoft servers, the **MS Sync Server** column of the **Zones** tab shows which Microsoft server is actually performing the synchronization of that zone with the Grid.

The Grid member periodically checks if each zone has one principal server. If it does not find a principal server for a zone, the Grid member selects one among the name servers assigned to the zone. It gives priority to the server that was not the designated principal server the longest.

Following are additional guidelines when synchronizing AD-integrated zones:

- You can create an AD-integrated zone on Grid Manager and assign one domain controller as its primary server. If a domain controller admin adds more primary servers to the zone, they are added to the zone on Grid Manager when the zone is synchronized. If you want to delete the primary servers, you must delete all the primary servers at once. You cannot delete only a subset of the servers.
- A situation could arise where two domain controllers in different AD domains are primary servers for zones with the same name. For example, two reverse-mapping zones could be named 1.1.10.in-addr.arpa in two domain controllers managed by the same member. If the two domain controllers are synchronized to different DNS views, the Grid member synchronizes each one separately. If the domain controllers are synchronized to the same DNS view, then the Grid member synchronizes the zone with the first domain controller. During the synchronization with the second domain controller, the Grid member logs an error and does not synchronize the zone.
- If a Grid Manager admin deletes a CNAME record that has a blank canonical name from an AD-integrated zone, this CNAME record is not deleted from the Microsoft server after the synchronization if the AD-integrated zone is hosted on a Microsoft server running Windows 2008 R2 or Windows Server 2012.
- When a Microsoft server is the primary server of a zone that contains an _msdcs zone, it appends the parent zone name to the server name in the NS record of the _msdcs zone. But when an Infoblox Grid member is the primary server of a zone that contains an _msdcs zone, it specifies the server name only in the NS record. For example, the _msdcs zone is in the corp100.com zone and the name server is nameserver100.com. When a Microsoft server is the primary server of corp100.com, the server name on the NS record of the _msdcs zone is nameserver100.com.corp100.com. When a Grid member is the primary server, the server name on the NS record of the _msdcs zone is nameserver100.com.

Resolving Conflicts

Some conflicts require intervention from an admin. For example, a Grid member cannot synchronize a zone when its primary server on the Microsoft server is different from its primary server on the Grid. When a Grid member is unable to synchronize data due to such conflicts, it logs an error, skips the object with the error and continues synchronizing the rest of the data. You can then view the Microsoft logs to check which objects were not synchronized. If you resolve the problem, the Grid member synchronizes the object on its next attempt. For information about the logs, see [Viewing Synchronization Logs](#) on page 965.

VIEWING MEMBERS AND MANAGED SERVERS

You can view Infoblox and Microsoft DNS servers by navigating to the **Data Management** tab -> **DNS** tab, and then selecting the **Members/Servers** tab. The panel displays the following information about each DNS server:

- **Name:** The hostname of the Grid member or Microsoft server.
- **Status:** The status of the DNS service on the Grid member or Microsoft server.
- **Comment:** Comments that were entered for the Grid member or Microsoft server.
- **Site:** Values that were entered for this pre-defined attribute.
- **Address:** The IP address of the Grid member or Microsoft server.

You can do the following:

- List the DNS views or zones served by the member or Microsoft server.
 - Click a Grid member or Microsoft server name.
- Edit the properties of a Grid member or Microsoft server.
 - Click the check box beside a Grid member or Microsoft server, and then click the Edit icon. To edit the DNS properties of a Grid member, see [Configuring DNS Service Properties](#) on page 557. To edit the DNS properties of a Microsoft server, see [Specifying Forwarders for Microsoft Servers](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.

- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Export the list of Grid members and Microsoft servers to a .csv file.
 - Click the Export icon.
- Print the list of Grid members and Microsoft servers.
 - Click the Print icon.

SPECIFYING FORWARDERS FOR MICROSOFT SERVERS

A forwarder is a name server to which all other name servers first send queries that they cannot resolve locally. The forwarder then sends these queries to DNS servers that are external to the network, avoiding the need for the other name servers in your network to send queries off-site. You can define a list of forwarders for each managed Microsoft server as follows:

1. From the **Data Management** tab, select the **DNS** tab -> **Members/Servers** tab -> *ms_server* check box -> Edit icon.
2. Click the Add icon and enter the IP address of the forwarder in the text field.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

DISABLING AND REMOVING MICROSOFT DNS SERVERS

When you disable synchronization with a Microsoft server, the managing Grid member terminates any ongoing synchronization and restarts it when the Microsoft server is re-enabled. The synchronized DNS data stays in the same state until synchronization resumes. For information, see [Disabling Synchronization](#) on page 962.

When you remove a managed Microsoft server from the Grid, the managing Grid member terminates any ongoing synchronization and does not start a new one. Zones and their content on the Microsoft server remain in the state that existed the moment the Microsoft server was removed. The Grid retains the zones that were assigned to the Microsoft server that was removed, but deletes the Microsoft server from its assigned zones as follows:

- If the Microsoft server is the only primary server and there are no other assigned servers or if the secondary server is an external secondary server, Grid Manager deletes all the server assignments.
- If the Microsoft server is the only primary server and there are Grid secondary servers, an external primary is created with the FQDN and IP address of the removed Microsoft server.
- If the Microsoft server is a secondary server and there is a Grid primary, an external secondary is created with the FQDN and IP address of the removed Microsoft server.
- If the Microsoft server is a server for a stub zone, the server is removed.

To remove a Microsoft DNS server:

1. From the **Data Management** tab, select the **DNS** tab -> **Members/Servers** tab -> *ms_server* check box.
2. Expand the Toolbar and click **Delete**.
3. Click **Yes** when the confirmation dialog box appears.



Chapter 34 Managing Microsoft DHCP Services

This chapter provides guidelines for using Grid Manager to manage Microsoft DHCP servers and for synchronizing DHCP data between Microsoft servers and the Grid. It discusses some features of the Microsoft DHCP servers only as they relate to the synchronization of data. Please review the Microsoft documentation for complete information about Microsoft DHCP servers and their features.

In addition, if you encounter technical issues with your Microsoft DHCP servers, contact Microsoft Technical Support or consult the Microsoft Support site at <http://support.microsoft.com/>. Some Windows versions require certain updates and hotfixes installed, so the Microsoft server can synchronize with the Grid member. For information about these requirements, see *Requirements* on page 955.

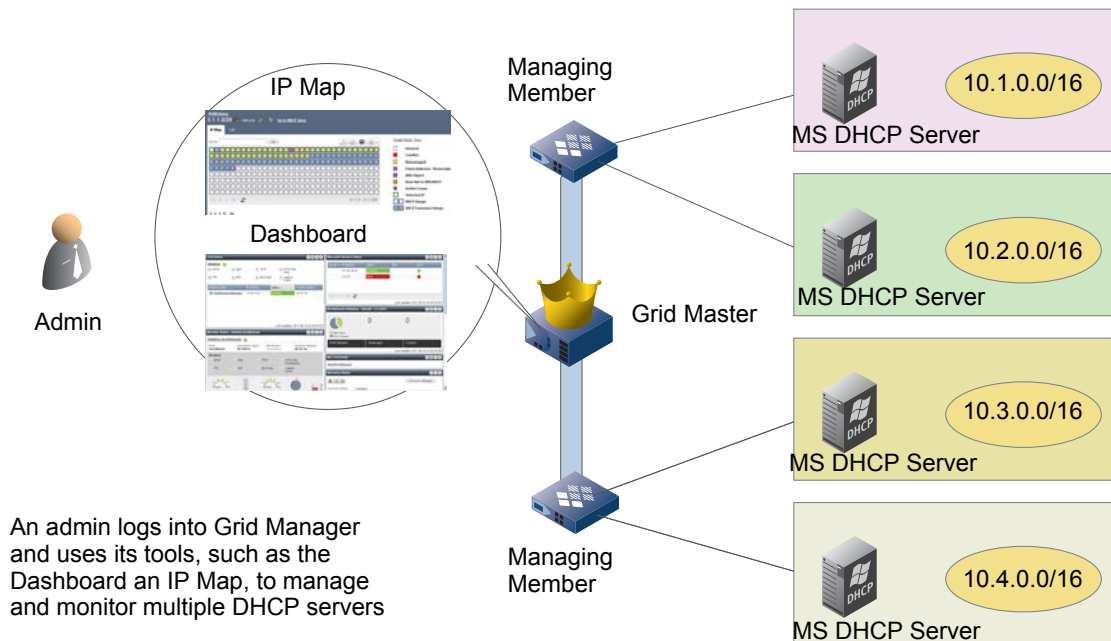
The topics in this chapter include:

- *About Microsoft DHCP Management* on page 984
 - *Synchronizing DHCP Data from Microsoft Servers* on page 984
 - *Viewing Synchronized Leases* on page 986
- *Managing Synchronized DHCP Data* on page 986
 - *Adding and Managing Scopes* on page 987
 - *Adding Fixed Addresses/Microsoft Reservations* on page 993
 - *About Superscopes* on page 995
- *Synchronizing Updates* on page 997
- *Managing Microsoft DHCP Servers* on page 998
 - *Viewing Members and Managed DHCP Servers* on page 998
 - *Setting Microsoft DHCP Server Properties* on page 999
 - *Controlling the DHCP Service of a Microsoft Server* on page 1000
 - *Disabling and Removing Microsoft DHCP Servers* on page 1000
 - *Modifying DHCP Server Assignments* on page 1000

ABOUT MICROSOFT DHCP MANAGEMENT

Grid Manager enables you to centrally manage the DHCP data of multiple Microsoft DHCP servers from a single interface. Once the DHCP data is synchronized, you can use the Dashboard on Grid Manager to monitor DHCP and server operations, or organize DHCP data into Smart Folders. Through IPAM tools, such as network maps and IP maps, you can track and manage IP address usage in your networks and monitor DHCP range utilization. You can also run a network discovery to retrieve IP allocation for both managed and unmanaged devices—including virtualized resources. For information about the IPAM features, see [Chapter 12, IP Address Management](#), on page 457.

Figure 34.1 Managing Microsoft DHCP Servers from Grid Manager



Synchronizing DHCP Data from Microsoft Servers

After you configure a member to manage the DHCP service of a Microsoft server, the Grid member connects to the server and starts synchronizing IPv4 DHCP data from the Microsoft server to the Grid database. It synchronizes the Microsoft server properties, leases, scopes and reservations.

The synchronization time varies, depending on different factors, such as the number of managed Microsoft servers and the amount of data being synchronized.

Note: Synchronizing IPv6 data is not supported.

As shown in [Table 34.1](#), Microsoft servers and Infoblox DHCP servers represent DHCP data differently. Scopes on Microsoft servers are DHCP ranges on Infoblox DHCP servers. Additionally, Microsoft servers support split-scopes, which is a scope assigned to two Microsoft servers. Each scope has an exclusion range on opposite ends to specify the pool of IP addresses that the other Microsoft server allocates. On an Infoblox DHCP server, each scope in the split-scope is represented as a DHCP range with an exclusion range. Note that NIOS also synchronizes scopes assigned to more than two Microsoft servers, but they are not synchronized as split-scopes.

Fixed addresses on Infoblox DHCP servers are the same as reservations on Microsoft servers. Infoblox reservations, which are IP addresses that are excluded from DHCP, are not supported on Microsoft servers. Microsoft superscopes, which are used to group scopes, are represented as superscopes and can be managed from Infoblox DHCP servers.

Table 34.1 DHCP Data in Microsoft and Infoblox DHCP Servers

DHCP Data	Microsoft DHCP Servers	Infoblox DHCP Servers
Address pool from which the server allocates addresses	Scope	DHCP Address Range in a Network
An IP address that is always assigned to the same device	Reservation	Fixed Address
An IP address that is excluded from DHCP because a user intends to configure it manually on a network device	Not supported	Reservation
Administrative group of scopes	Superscope	Microsoft superscope

Note: In this chapter, reservations always refer to Microsoft reservations (Infoblox fixed addresses), unless otherwise specified.

When the member synchronizes a scope to the Grid, it converts the scope to a DHCP range and network. For example, it converts the Microsoft scope 10.1.1.1- 10.1.1.200 with a netmask of /24 to the network 10.1.1.0/24 and DHCP range 10.1.1.1- 10.1.1.200 on Grid Manager. The member associates the DHCP properties of the scope, including its DHCP and Microsoft vendor options, with the DHCP range. It synchronizes the leases within the range and if configured, the exclusion range as well.

NIOS synchronizes two scopes as split-scopes if the following conditions are met:

- Two scopes have the same address range.
- The scopes are assigned to two different Microsoft servers.
- Each scope has an exclusion range and the exclusion ranges are at opposite ends of the scope, so they complement each other. For example, the scope 10.1.1.1-10.1.1.200 on Microsoft server A has an exclusion range of 10.1.1.100-10.1.1.200 and the same scope on Microsoft server B has an exclusion range of 10.1.1.1-10.1.1.99.

When the appliance synchronizes a split-scope, it sets a split-scope flag on each scope to indicate that it is part of a split-scope. For more information, see [Viewing Scopes](#) on page 992. It synchronizes any reservations that are configured in each scope as well.

When the member synchronizes a Microsoft reservation to the Grid, it converts the reservation to a fixed address and static lease on Grid Manager. It associates the DHCP properties and DHCP and Microsoft vendor options of the reservation with the fixed address record.

The Grid member synchronizes superscopes to the Grid as well. The Grid supports Microsoft superscopes, when an MS management license is installed. For information about adding and managing superscopes in Infoblox DHCP servers, see [About Superscopes](#) on page 995.

Following are some guidelines on how a Grid member synchronizes DHCP data from Microsoft servers to the Grid:

- If two superscopes have the same name, but are served by different servers, the member creates two different superscopes on the Grid, each appended with the Microsoft server FQDN.
- The member synchronizes all active and inactive scopes from a managed Microsoft server as long as the scopes do not conflict or include any networks currently served by a Grid member. The member does not synchronize a scope if its network already exists in the Grid and is served by a Grid member. It can synchronize a scope if its network is included in an existing network, only if the network is not served by DHCP.
- Synchronizing scopes that are larger than /12 is not supported.
- NIOS synchronizes all scopes except for those with serving ranges that overlap the serving ranges of existing DHCP ranges.
- If the appliance manages multiple Microsoft servers and synchronizes identical scopes from more than two Microsoft servers, it does not flag the scopes as split-scopes.

- If the appliance synchronizes one or more scopes from Microsoft servers that are identical to an existing split-scope, it removes the split-scope flag from the existing split-scope.
- NIOS does not synchronize partially overlapping scopes inside a single network from different Microsoft servers. It synchronizes only ranges that completely overlap.
- More than two scopes are not synchronized as split-scopes, even if they are identical and have exclusion ranges that complement each other.
- Scopes that have more than one exclusion range are not synchronized as split-scopes, even if the exclusion ranges complement each other. In addition, if a split-scope is synchronized from a Microsoft server and one of the scopes is split again on the Microsoft server, NIOS synchronizes the third scope, but does not set a split-scope flag. In addition, it removes the split-scope flag from the original split-scopes.

You can view the synchronized data as follows:

- To view the networks of the scopes, select the **Data Management** tab -> **DHCP** tab -> **Networks** tab -> **Networks** panel. This panel displays all IPv4 networks. For information about this panel, see [Modifying IPv4 Networks](#) on page 851.
- To view the corresponding DHCP ranges and reservations, select the **Data Management** tab -> **DHCP** tab -> **Networks** tab, and click a network link. For information about this panel, see [Viewing Scopes](#) on page 992.

You can also use the features in the **IPAM** tab, such as the Net Map and IP Map, to view and manage the Microsoft DHCP data. For information, see [Chapter 12, IP Address Management](#), on page 457.

Viewing Synchronized Leases

A Grid member synchronizes all leases from its managed Microsoft server to the Grid. Microsoft servers automatically generate a static lease for each reservation. These static leases are synchronized to the Grid as well. You can view the synchronized leases by navigating to the **Data Management** -> **DHCP** -> **Leases** tab. For information about viewing current leases, see [Viewing Current Leases](#) on page 946. You can do the following:

- View lease details, by selecting a lease and clicking the Lease Details icon. For additional information, see [Viewing Detailed Lease Information](#) on page 948.
- Clear a lease, by selecting it and clicking the Clear Lease icon. Note that Grid Manager clears the lease immediately. It does not wait for the next synchronization.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

Additionally, you can enable a Grid member to log lease related operations, so you can view these events in the Lease History panel. For information, see [Configuring the Lease Logging Member](#) on page 815 and [Viewing Lease History](#) on page 949.

MANAGING SYNCHRONIZED DHCP DATA

When Grid members are configured to manage Microsoft DHCP servers in read/write mode, you can use Grid Manager to view, edit and delete the DHCP data of those servers. You can add and manage networks and DHCP ranges that are synchronized as scopes to the Microsoft server, and add and manage reservations and superscopes. All updates are synchronized to the Microsoft servers at regular intervals.

Grid Manager also allows you to set admin permissions, extensible attributes, and thresholds. These apply only when the DHCP data is managed on Grid Manager. These properties are not synchronized to Microsoft servers.

The following sections provide guidelines for managing Microsoft DHCP data from Grid Manager:

- [Adding and Managing Scopes](#)
- [Adding Fixed Addresses/Microsoft Reservations](#) on page 993
- [About Superscopes](#) on page 995

Adding and Managing Scopes

To add a scope from Grid Manager, you must create an IPv4 network and a DHCP range, and then assign the Microsoft server to the network and range. To add a split-scope from Grid Manager, you must create an IPv4 network and a DHCP range, and then assign two Microsoft server to the network and range.

To edit a scope synchronized from a Microsoft server, you must edit the properties of its corresponding DHCP range. The following sections describe how to add, edit and remove scopes using Grid Manager.

Note: Microsoft servers do not support Infoblox hosts and reservations. You cannot add them to networks and DHCP ranges served by Microsoft servers.

Adding Networks for Scopes

Following are guidelines for adding a network for Microsoft scopes:

- The network must be served by Microsoft servers. It cannot be served by a mix of Microsoft and Infoblox DHCP servers.
- If you are adding a split-scope, you must assign the network to two Microsoft servers that serve the split-scope. A split-scope cannot be served by a mix of Microsoft and Infoblox DHCP servers.
- The network can contain only one DHCP range per Microsoft server. It can contain multiple DHCP ranges as long as they do not overlap and are each served by a different Microsoft server.
- You can set DHCP properties at the DHCP range level only, not the network level.

You can run discoveries on networks served by Microsoft servers. For information about network discoveries, see [Network Discovery](#) on page 493.

Note: Networks served by Microsoft DHCP servers do not support the split, join, and expand functions.

You can create a network from scratch or use a network template. For information about creating network templates, see [Adding IPv4 Network Templates](#) on page 829. To add an IPv4 network for a scope:

1. From the **Data Management** tab, select the **DHCP** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network. It must be the same network view to which the Microsoft server is assigned.
3. Expand the Toolbar and click **Add -> Network**.
4. In the *Add Network* wizard, select one of the following and click **Next**:
 - **Add Network**
 - or
 - **Add Network using Template:** Click **Select Template** and select a network template. For more information, see [About IPv4 Network Templates](#) on page 829. In the *DHCP Network Template Selector* dialog box, select the template you want to use and click the Select icon. Note that when you use a template to create a network, the configurations of the template apply to the new network. The appliance populates the template properties in the wizard when you click **Next**. You can then edit the pre-populated properties, except for **Netmask**.
5. Complete the following and click **Next**:
 - **Address:** Enter the IP address of the network. You can enter the IP address with a CIDR block. For example, enter 10.0.0.0/24, and the netmask slider adjusts the netmask to /24. You can also enter partial IP address with a CIDR block. When you are done, Grid Manager displays the complete IP address with the CIDR block. For example, when you enter 15/24, Grid Manager displays 15.0.0.0/24 and the netmask slider adjusts the netmask to /24. Note that Microsoft DHCP servers do not support /32 subnets.
 - **Netmask:** Use the netmask slider to select the appropriate number of subnet mask bits for the network. Microsoft servers support /1 to /31 netmasks. Note that when you use a template that contains a fixed netmask, you cannot adjust the netmask for this network.
 - **Comment:** Enter additional information about the network, such as the name of the organization it serves.

- **Automatically create reverse-mapping zone in view:** This function is enabled if the netmask of the network equals /8, /16, or /24. Select this to have the appliance automatically create reverse-mapping zones for the network. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network view level.
 - **Disabled:** This option does not apply to networks assigned to Microsoft servers. The member ignores this field when the network is assigned to Microsoft servers. You can disable DHCP ranges assigned to Microsoft servers, but not networks.
6. Click **Next** to add Microsoft servers as DHCP servers for the network. Click the Add icon and select the following:
 - **Add Microsoft Server:** Select the Microsoft server from the *Microsoft Server Selector* dialog box. You can add multiple Microsoft servers, if you are adding multiple DHCP ranges served by different Microsoft servers. For a split-scope, you must assign two Microsoft servers to the network.
 7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 8. Save the configuration and click **Restart** if it appears at the top of the screen.
or
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Setting Network Properties

You can change the Microsoft servers assigned to the network, and define extensible attributes and admin permissions to the network. You can also set thresholds for the network, to enable the appliance to make a syslog entry when address usage goes above or below the thresholds.

To set network properties:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.
2. The *Network* editor contains the following basic tabs from which you can modify data:
 - **General Basic:** You can enter or modify comments.
 - **Member Assignment:** Add or delete Microsoft servers. For information, see [Adding IPv4 Networks](#) on page 845. If the network contains multiple DHCP ranges each managed by a different Microsoft server, then you can add those Microsoft servers here.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Expert Mode** to display the following tabs from which you can modify advanced data.
 - **General Advanced:** You can associate zones with a network. For information, see [Associating Networks with Zones](#) on page 813.
 - **Thresholds:** These watermarks represent thresholds above or below which address usage is unexpected and might warrant your attention. Thresholds are inherited from the Grid.
 - **High-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range exceeds this number, the appliance makes a syslog entry. The default is 95.
 - **Low-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below this number, the appliance makes a syslog entry. The default is 0. Address usage must initially exceed the low-water mark threshold and then dip below it before the appliance considers low address usage an event requiring an alert.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting and Restoring a Network

When you delete a network, Grid Manager moves it and its DHCP ranges and fixed addresses to the Recycle Bin, and permanently deletes its leases. The corresponding scopes and reservations are deleted from the Microsoft server at the next synchronization. If you restore the network on Grid Manager, its DHCP ranges and fixed addresses are restored as well. The Grid member then adds the corresponding scopes and reservations to the Microsoft server on the next synchronization. For information about deleting networks, see [Deleting IPv4 Networks](#) on page 852. For information about restoring data, see [Using the Recycle Bin](#) on page 64.

Adding a DHCP Range/Scope

After you add a network for a scope, you must then define its DHCP range. You can create the DHCP range from scratch or use a DHCP Range template. For information about DHCP templates, see [About IPv4 Range Templates](#) on page 826. You can add multiple ranges to the same network, as long as each range is served by a different Microsoft server and the ranges do not overlap.

When you add a split-scope, you must specify the Microsoft servers that serve the scopes and their exclusion ranges. Each scope inherits its options from its respective Microsoft server. Note that the enabled/disabled setting of the first range automatically applies to the second range. Therefore, if the first range is initially disabled, then the second range is initially disabled as well.

To add a DHCP range for a scope:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Navigate to the network to which you want to add a DHCP range, and then click **Add -> DHCP Range** from the Toolbar. You can also add a DHCP range from any panel in the DHCP tab.
3. In the *Add Range* wizard, select one of the following and click **Next**:
 - **Add DHCP Range**

or

 - **Add DHCP Range using Template**

Click **Select Template** and select the template that you want to use. Note that when you use a template to create a DHCP range, the configurations of the template apply to the new range. The appliance automatically populates the DHCP range properties in the wizard. You can then edit the pre-populated properties.
4. Complete the following:
 - **Network:** Click **Select Network** to select the network to which you want to add the DHCP range. The network must be served by a Microsoft server. If you are adding a DHCP range while viewing the contents of a specific network, the appliance displays the network address here. You can still select a different network by clicking **Select Network**.
 - **Start:** Enter the first IP address in the range.
 - **End:** Enter the last IP address in the range.
 - **Name:** You can enter a name for the DHCP range.
 - **Comment:** You can enter additional information. After the range is synchronized to the Microsoft server as a scope, this text appears in the Description field of the scope on the Microsoft server.
 - **Disabled:** Select this if you do not want the DHCP server to allocate IP addresses from this DHCP range at this time. If you select this, the Grid member synchronizes the range to the Microsoft server as an inactive scope.

5. Click **Next** and select one of the following to provide DHCP services for the DHCP range:
 - None:** Select this if you do not want to synchronize this range to the Microsoft DHCP server.
 - Microsoft Server:** This field displays the Microsoft server that you selected for the network. If several servers were assigned to the network, you can select one from the list.
 - **Microsoft Split-Scope:** Select this to create a split-scope, and then complete the following:
 - **Microsoft Server #1:** Read-only field that displays the Microsoft server that you specified in the preceding step.
 - **Microsoft Server #2:** Select the Microsoft server that will serve the split-scope.
 - **Split Percentage:** Specify the percentage of IP addresses in the scope that is allocated to the exclusion range of each Microsoft server. The default is 50%. You can either move the slider or enter the percentages in the text fields. When you use the slider, you are specifying the percentage of addresses in the exclusion range of the first server. A tooltip window displays the percentage as you adjust the slider. When you set the slider, the **Split Percentage**, **Exclusion Starting Address**, and **Exclusion Ending Address** fields are updated accordingly.
 - **Exclusion Starting Address:** When you set the split percentages, these fields automatically displays the starting address of the exclusion range of each Microsoft server. Alternatively, you can enter the starting address of the exclusion range of the first Microsoft server, and the **Split Percentage** and **Exclusion Ending Address** values adjust accordingly.
 - **Exclusion Ending Range:** When you set the split percentage, these fields automatically display the ending address of the exclusion range of each Microsoft server. Alternatively, you can enter the ending address of the exclusion range of the second Microsoft server, and the **Split Percentage** and **Exclusion Starting Address** values adjust accordingly.
6. Click **Next**, and optionally set operational parameters for the scope. Otherwise, the scope inherits its parameters from the first Microsoft DHCP server.
 - **Lease Time:** Specify the lease time. The default is 8 days. When the range is served by a Microsoft server and you enter a lease time of 1000 days or more, Grid Manager automatically grays out this field and checks the **Unlimited Lease Time option** after you save your entries.
 - **Unlimited Lease Time:** Select this option to set an infinite lease time for the IP addresses leased from this range.
 - **Routers:** In the table, enter the IP address of the router that is connected to the same network as the DHCP clients. Click the Add icon to add more routers.
 - **Domain Name:** Enter the name of the domain for which the Microsoft server serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCP OFFER packet to a DHCPDISCOVER packet from a client.
 - **DNS Servers:** In the table, enter the IP address of the DNS server to which the DHCP clients send name resolution requests. The DHCP server includes this information in the DHCP OFFER and DHCPACK messages.
 - **Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached.
7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
8. Save the configuration and click **Restart** if it appears at the top of the screen.
 - or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Setting DHCP Range/Scope Properties

A Microsoft scope inherits its properties from its Microsoft server. In Grid Manager, you can override the inherited values or set other properties by editing the DHCP range. You can also configure an exclusion range within the scope and set thresholds, to enable the appliance to make a syslog entry when address usage goes above or below the thresholds.

You can modify a scope's properties, including its start and end addresses, servers, and exclusion ranges. If you edit the properties of a split-scope and it results in gaps or overlapping exclusion ranges so that the ranges are no longer identical, Grid Manager displays a warning indicating that continuing with the operation automatically removes the split-scope flag. Grid Manager also removes the flag when the start or end address of a scope changes, so its range is no longer the same.

To set DHCP range properties:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon.
2. The *DHCP Range* editor contains the following basic tabs from which you can modify data:
 - **General:** Modify the fields, including the start and end addresses, as described in [Adding a DHCP Range/Scope](#).
 - **Server Assignment:** Switch to **None** or select a different Microsoft server for the DHCP range.
 - **IPv4 DHCP Options:** Keep the DHCP properties or override them and enter unique settings for the DHCP range. For information about the fields, see [Adding a DHCP Range/Scope](#) on page 989.
 This tab displays DHCP and Microsoft vendor options that were synchronized from the Microsoft server. You can edit any of the options. When you select a different User Class or Vendor Class from the drop-down menus, Grid Manager automatically updates the option definitions in the drop-down list.
 To configure additional DHCP options, click **+** and select a User Class and Vendor Class from the drop-down menus. Select an option from the drop-down list, and enter a value in the field beside it. You can click **-** to remove an option.
 - **Extensible Attributes:** You can add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Expert Mode** to display the following tabs from which you can modify advanced data.
 - **DDNS:** Complete the following to set DDNS parameters for the range:
 - **Enable DDNS Updates:** Click the check box to enable the Microsoft DHCP server to send dynamic DNS updates or clear the check box to disable this function.
 - **Option 81 Support**
DHCP Server Updates DNS If Requested by Client: The DHCP server updates DNS only if it is requested by the client. Otherwise, the client updates DNS.
DHCP Server Always Updates DNS: The DHCP server always updates DNS, regardless of any client request.
 - **Exclusion Ranges:** Configure a range of IP addresses that the server does not use to assign to clients. You can use these exclusion addresses as static IP addresses. For information, see [Configuring IPv4 Fixed Addresses](#) on page 857. In a split-scope, the exclusion range identifies the range of IP addresses that the other Microsoft server serves. If you edit the exclusion range of either of the scopes in a split-scope and the exclusion ranges no longer complement each other, NIOS removes the split-scope flag from both scopes.
 - **Thresholds:** Thresholds are inherited from the Grid. These watermarks represent thresholds above or below which address usage is unexpected and might warrant your attention.
 - **High-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range exceeds this number, the appliance makes a syslog entry. The default is 95.

- **Low-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below this number, the appliance makes a syslog entry. The default is 0. Address usage must initially exceed the low-water mark threshold and then dip below it before the appliance considers low address usage an event requiring an alert.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting and Restoring a DHCP Range/Scope

When you delete a DHCP range, Grid Manager moves it and its exclusion range and fixed addresses to the Recycle Bin, and permanently deletes its leases. At the next synchronization, the member deletes the scope, its exclusion range and reservations from the Microsoft server. If you restore a DHCP range on Grid Manager, then the Grid member adds its corresponding scope, exclusion range and reservations to the Microsoft server at the next synchronization. For information about deleting DHCP ranges, see [Deleting IPv4 Address Ranges](#) on page 857. For information about restoring data, see [Using the Recycle Bin](#) on page 64.

If you delete a scope that is part of a split-scope, Grid Manager automatically removes the split-scope flag from the remaining scope.

Viewing Scopes

To view the scopes in a network, navigate to **DHCP -> Networks -> network**. The panel displays the objects in the network, including the scopes and split-scopes. For split-scopes, the panel displays both scopes with the same start and end address. It displays the following information about each object:

- **IP Address:** The IP address of the DHCP object. For a scope, this field displays the start and end addresses of the scope. Note that the appliance highlights all disabled DHCP objects in gray.
- **Split-Scope:** Displays **Yes** if the scope is a split-scope.
- **MS Server:** Displays the Microsoft server that is serving the scope.
- **Type:** The DHCP object type, such as **DHCP Range** or **Fixed Address**.
- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Comment:** The information you entered for the object.
- **IPv4 DHCP Utilization:** The percentage of the total DHCP usage of a DHCP range. This is the percentage of the total number of fixed addresses, reservations, hosts, and active leases in the DHCP range divided by the total IP addresses in the range, excluding the number of addresses in the exclusion ranges. Note that only enabled objects are included in the calculation.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes.

You can select the following additional columns for display:

- **Static Addresses:** Indicates whether the IP address is a static address.
- **Dynamic Addresses:** Indicates whether the IP address is a dynamically assigned address.
- **Disabled:** Indicates whether the object is disabled.
- **Priority:** Displays the priority of a DHCP range when NAC filters are applied.
- Available extensible attributes.

You can also do the following in this panel:

- Sort the displayed data in ascending or descending order by column.
- Click **Go to IPAM View** to view information about the object in the **IPAM** tab.
- Add new objects, such as DHCP ranges, to the network.
- Delete or schedule the deletion of a selected object or multiple objects.

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print or export the data.

You can also view the scopes in the IP Map.

Adding Fixed Addresses/Microsoft Reservations

To add a reservation from Grid Manager, add a fixed address and Grid Manager synchronizes it to the Microsoft server as a reservation. You can create fixed addresses from scratch or use fixed address templates. For information about fixed address templates, see [Adding IPv4 Fixed Address/Reservation Templates](#) on page 828.

To add a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Expand the Toolbar and click **Add -> Fixed Address**.
3. In the *Add Fixed Address* wizard, select one of the following and click **Next**:
 - **Add Fixed Address**
 - or
 - **Add Fixed Address using Template**
Click **Select Template** and select the template that you want to use.
4. Complete the following:
 - **Network:** Click **Select Network** to select the network to which you want to add the fixed address. If you are adding the fixed address from a specific network, the appliance displays the network address here. You can still select a different network by clicking **Select Network**.
 - **IP Address:** Enter the IPv4 address for the fixed address, or click **Next Available IP** to obtain the next available IP address.

Note: When you save the configuration, the appliance displays an error message if the IP address obtained through **Next Available IP** is being used by another object or operation. You can request another unused IP address or enter a new one.

- **MAC Address:** Enter the MAC address of the host.
 - **Name:** Enter a name for the fixed address. This is required for reservations on Microsoft servers.
 - **Configure On:**
 - None:** Select this if you do not want this synchronized to the Microsoft server.
 - Microsoft Server:** Select the Microsoft server that serves this fixed address.
 - **Comment:** Optionally, enter additional information. The text in this field appears in the Description field of the Microsoft reservation after the fixed address is synchronized. Note that due to a length limit set by the Microsoft DHCP server, after you synchronize DHCP data, the Description field can display only up to 128 characters even though NIOS allows up to 256 characters for this field.
5. Click **Next**, and optionally set operational parameters for the fixed address. Otherwise, the fixed address inherits its parameters from its scope.
 - **Routers:** In the table, enter the IP address of the router that is connected to the same network as the DHCP client. Click the Add icon to add more routers.
 - **Domain Name:** Enter the name of the domain for which the Microsoft DHCP serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCP OFFER packet to a DHCP DISCOVER packet from a client.
 - **DNS Servers:** In the table, enter the IP address of the DNS server to which the DHCP client sends name resolution requests. The DHCP server includes this information in the DHCP OFFER and DHCP ACK messages.
 - **Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached.

6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
7. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Setting Fixed Address/Reservation Properties

Microsoft reservations inherit their properties from their scopes. In Grid Manager, you can override the inherited values or set other properties of a Microsoft reservation, by editing its fixed address.

To modify a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab → **Networks** tab → **Networks** → *network* → *fixed_address* check box, and then click the Edit icon.
2. The *Fixed Address* editor contains the following basic tabs from which you can enter data:
 - **General:** You can modify the fields described in [Adding Fixed Addresses/Microsoft Reservations](#) on page 993.
 - **IPv4 DHCP Options:** Keep the inherited properties, or override them and enter unique settings.
This section displays DHCP and Microsoft vendor options that were synchronized from the Microsoft server. You can edit any of the options. When you select a different User Class or Vendor Class from the drop-down menus, Grid Manager automatically updates the option definitions in the drop-down list.
To configure additional DHCP options, click **+** and select a User Class and Vendor Class from the drop-down menus. Select an option from the drop-down list, and enter a value in the field beside it. You can click **-** to remove an option.
 - **Discovered Data:** If you ran a discovery on the network, Grid Manager displays the discovered data of the fixed address. For information, see [Viewing Discovered Data](#) on page 510. Note that conflicts can occur when discovered data does not match the existing IP address data. For information about resolving these conflicts, see [Resolving Conflicting Addresses](#) on page 513.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 167.
3. Optionally, you can click **Toggle Expert Mode** to display the **DDNS** tab. To set DDNS parameters for the fixed address, complete the following:
 - **Enable DDNS Updates:** Click the check box to enable the Microsoft DHCP server to send dynamic DNS updates or clear the check box to disable this function.
 - **Option 81 Support**
 - **DHCP Server Updates DNS If Requested by Client:** The DHCP server updates DNS only if it is requested by the client. Otherwise, the client updates DNS.
 - **DHCP Server Always Updates DNS:** The DHCP server always updates DNS, regardless of any client request.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
or
 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Deleting and Restoring a Fixed Address/Reservation

When you delete a fixed address, Grid Manager moves it to the Recycle Bin. At the next synchronization, the Grid member deletes its corresponding reservation from the Microsoft server. If you restore fixed address, then the Grid member adds its corresponding reservation to the Microsoft server at the next synchronization. For information about deleting fixed addresses, see [Deleting Fixed Addresses](#) on page 860. For information about restoring data, see [Using the Recycle Bin](#) on page 64.

About Superscopes

In Grid Manager, you can group DHCP ranges served by Microsoft servers into a superscope. You can add multiple DHCP ranges to a superscope, as long as the ranges are all served by the same Microsoft DHCP server. The Grid member then synchronizes the superscope and its associated DHCP ranges as superscopes and scopes to the Microsoft DHCP server.

You can also associate extensible attributes with superscopes in Grid Manager. Extensible attributes are not synchronized to the Microsoft DHCP server.

Only admins with read/write permission to superscopes can add and manage superscopes.

Adding Superscopes

Before you add a superscope, you must first create at least one DHCP range to include in the superscope.

To add a superscope:

1. From the **Data Management** tab, select the **DHCP** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the superscope. The network view must be the same one that is assigned to the Microsoft server.
3. Expand the Toolbar and click **Add -> Superscope**.
4. In the *Add Superscope* wizard, complete the following and click **Next**:
 - **Name**: Enter a name for the superscope.
 - **Comment**: Optionally, enter additional information about the superscope.
 - **Disabled**: Select this to disable the DHCP ranges in the superscope. They are then synchronized as inactive scopes on the Microsoft server.
5. Click the Add icon and select a range from the *Select Range* dialog box. This dialog box lists only the address ranges that are served by a Microsoft server.
6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [About Extensible Attributes](#) on page 322.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

or

 - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Viewing Superscopes

To view superscopes, navigate to the **Data Management** tab -> **DHCP** tab -> **Networks** tab -> **Microsoft Superscopes**. Grid Manager displays the following information about each superscope that is displayed:

- **Name**: The name of the superscope. Grid Manager appends the FQDN of its associated Microsoft server so you can identify which superscope belongs to which server.
- **Comment**: The comment that was entered for the superscope.
- **DHCP Utilization**: The percentage of the total DHCP usage of the ranges in the superscope. Fixed addresses and reservations that are outside of a range are excluded from the calculation.
- **Site**: The site of the superscope. This is one of the predefined extensible attributes.

You can add the following columns for viewing:

- **Static Addresses:** The number of static addresses.
- **Dynamic Addresses:** The number of dynamic addresses.
- **Disabled:** Indicates whether the superscope is enabled.

You can do the following in this section:

- Click the link of a superscope to list its address ranges.
- Add a superscope.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print or export the information in this section.
- Delete a superscope.

Modifying Superscopes

To modify a superscope:

1. From the **Data Management** tab, select the **DHCP** tab -> **Network** tab -> **Microsoft Superscopes** -> *ms_superscope* check box, and then click the Edit icon.
2. The *Superscopes* editor contains the following tabs from which you can modify data:
 - **General:** You can modify the name and comment, and enable or disable the superscope. You can also add and delete address ranges from the superscope. Note that when you delete the last DHCP range in a superscope, Grid Manager automatically deletes the superscope as well.
 - **Extensible Attributes:** Define extensible attributes for the superscope. These apply only when the superscope is managed in Grid Manager. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** Define administrative permissions that apply to the superscope when it is managed in Grid Manager. For information see [About Administrative Permissions](#) on page 160.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Deleting Superscopes

When you delete a superscope in Grid Manager, it is permanently deleted from the database. The superscope is deleted from the Microsoft server at the next synchronization. Note that deleting a superscope does not delete the DHCP ranges in the superscope. These are retained in the database.

To delete a superscope:

1. From the **Data Management** tab, select the **DHCP** tab -> **Network** tab -> **Microsoft Superscopes** -> *ms_superscope* check box, and then click the Delete icon.
2. Click **Yes** when the confirmation dialog appears.

SYNCHRONIZING UPDATES

A Grid member synchronizes DHCP data with each of its managed Microsoft server at regular intervals. During each synchronization, updates from Grid Manager are applied to the Microsoft server and updates from the Microsoft server are applied to the Grid as well.

Because admins can update DHCP data from both the Microsoft server and from Grid Manager, conflicts can occur during synchronization. The following guidelines describe how the Grid member resolves conflicts and handles any differences when DHCP data is synchronized between a Microsoft server and the Grid.

- If a Microsoft server admin modifies an object that has a pending scheduled task in Grid Manager and synchronization occurs before the scheduled task, the object is modified in both the Microsoft server and the Grid member. When the scheduled task executes at its scheduled time, it fails and an error message is logged in the audit log.
- When a Microsoft server admin and a Grid Manager admin change the same object, the Grid member retains the version that exists on the Microsoft server. Following are some examples:

Table 34.2

Grid Manager Admin...	Microsoft Server Admin...	After Synchronization
Deletes the 10.1.1.0/24 network which has two DHCP ranges	Adds a scope that is within the 10.1.1.0/24 network	The 10.1.1.0/24 network is created on the Grid with the updates and is assigned to the Microsoft server.
Changes the DHCP options of a scope	Deletes the scope.	The scope is deleted from the Grid as well.

- If a Grid member manages multiple Microsoft servers, it can synchronize scopes to the same network as long as they are served by different Microsoft servers and they do not overlap. If the Microsoft servers have scopes that overlap, the Grid member synchronizes only one of the scopes, including its reservations. It does not synchronize the other scopes and logs an error message for each scope that is not synchronized. For information about the Microsoft logs, see [Viewing Synchronization Logs](#) on page 965.
Note that a Grid member can synchronize scopes with overlapping reservations because they are served by different Microsoft servers.
- When a Grid member synchronizes a split-scope to its respective Microsoft servers, the scopes use the default value for the DHCP Offer Delay value, since this property is not supported by NIOS.
- If you create a split-scope on a NIOS appliance, synchronization fails if there is an existing scope in the same network on one of the Microsoft servers. Only one scope is allowed in a network, per Microsoft server.
- If a Microsoft admin adds a DHCP range and a NIOS admin is in the process of adding the same range when a synchronization occurs, the NIOS admin will not be able to save the range after the synchronization. Grid Manager will display an error message indicating that the range already exists.
- If both a NIOS admin and a Microsoft admin create a scope or split-scope and conflicts occur, the Microsoft server always takes precedence. All conflicts are logged to the Microsoft log. Following are some examples:
 - If the NIOS admin creates a scope and a Microsoft server admin creates a split-scope for the same DHCP range, the split-scope is synchronized to Grid Manager.
 - If the NIOS admin creates a split-scope on Microsoft servers 1 and 2, and a Microsoft admin creates the same split-scope on Microsoft servers 1 and 3 but with different exclusion ranges, the scope created by the NIOS admin on Microsoft server 1 is dropped upon synchronization.
 - If the NIOS admin creates a split-scope on Microsoft servers 1 and 2, and a Microsoft admin creates the same split-scope on the same Microsoft servers but with different exclusion ranges, the split-scope created by the Microsoft admin is synchronized to NIOS and retained. The split-scope created by the NIOS admin is dropped.

MANAGING MICROSOFT DHCP SERVERS

You can control the DHCP services of managed Microsoft servers and set certain properties as well. This section includes the following topics:

- [Viewing Members and Managed DHCP Servers](#)
- [Setting Microsoft DHCP Server Properties](#) on page 999
- [Controlling the DHCP Service of a Microsoft Server](#) on page 1000
- [Disabling and Removing Microsoft DHCP Servers](#) on page 1000
- [Modifying DHCP Server Assignments](#) on page 1000

Viewing Members and Managed DHCP Servers

You can view Infoblox and Microsoft DHCP servers by navigating to the **Data Management** tab -> **DHCP** tab, and then selecting the **Members/Servers** tab. The panel displays the following information about each DHCP server:

- **Name:** The hostname of the Grid member or Microsoft server.
- **Status:** The status of the DHCP service on the Grid member or Microsoft server.
- **Comment:** Comments that were entered for the Grid member or Microsoft server.
- **DHCP Utilization:** The percentage of the total DHCP utilization of the member or Microsoft server. This is the percentage of the total number of DHCP hosts, fixed addresses, reservations, and leases assigned to the member or Microsoft server versus the total number of IP addresses (excluding IP addresses in the exclusion range) and all DHCP objects assigned to the member or DHCP server. Note that only enabled objects are included in the calculation. The appliance updates the utilization data every 15 minutes. The appliance displays the utilization data in one of the following colors:
 - Red: The DHCP resources are 100% utilized.
 - Yellow: The utilization percentage is over the effective high watermark threshold.
 - Blue: The utilization percentage is below the effective low watermark threshold.
 - Black: The utilization percentage is at any number other than 100%, or within the effective thresholds.
- **Site:** Values that were entered for this pre-defined attribute.

You can select the following additional columns for display:

- **Address:** The IP address of the member or Microsoft server.
- **Static Addresses:** The number of static IP addresses.
- **Dynamic Addresses:** The number of dynamically assigned IP addresses.

You can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Edit the properties of a Grid member or Microsoft server.
 - Click the check box beside a Grid member or Microsoft server, and then click the Edit icon.
- Export the list of Grid members and Microsoft servers to a .csv file.
 - Click the Export icon.
- Print the list of Grid members and Microsoft servers.
 - Click the Print icon.

Setting Microsoft DHCP Server Properties

From Grid Manager, you can set DHCP properties supported by a Microsoft server. These are applied to the server at the next synchronization. You can also set other properties that apply to Grid Manager only, such as thresholds and the logging.

To set properties for a Microsoft DHCP server:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members/Servers** tab -> **Members/Servers** -> *ms_server* check box, and then click the Edit icon.
2. In the *Microsoft Server DHCP Properties* editor, you can configure DHCP properties in each tab as follows:

IPv4 DHCP Options tab: Complete the following to configure basic DHCP options for the server:

- **Routers:** Click the Add icon and enter the IP address of the router that is connected to the same network as the DHCP clients.
- **Domain Name:** Enter the name of the domain for which the server serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCPOFFER packet to a DHCPDISCOVER packet from a client. If DDNS is enabled on the DHCP server, it combines the host name from the client and this domain name to create the FQDN (fully-qualified domain name) that it uses to update DNS.
- **DNS Servers:** Click the Add icon and enter the IP address of the DNS server to which the DHCP client sends name resolution requests. The DHCP server includes this information in the DHCPOFFER and DHCPACK messages.
- **Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached.
- **Custom DHCP Options:** This section displays DHCP and Microsoft vendor options that were synchronized from the Microsoft server. You can edit any of the options. When you select a different User Class or Vendor Class from the drop-down menus, Grid Manager automatically updates the option definitions in the drop-down list.

To configure additional DHCP options, click **+** and select a User Class and Vendor Class from the drop-down menus. Select an option from the drop-down list, and enter a value in the field beside it. You can click **-** to remove an option.

DDNS tab: You can enable or disable dynamic DNS updates and set certain properties.

- **Enable DDNS Updates:** Click the check box to enable the Microsoft DHCP server to send dynamic DNS updates or clear the check box to disable this function.
- **Option 81 Support**
DHCP Server Updates DNS If Requested by Client: The DHCP server updates DNS only if it is requested by the client. Otherwise, the client updates DNS.
DHCP Server Always Updates DNS: The DHCP server always updates DNS, regardless of any client request.

Thresholds tab: Thresholds are inherited from the Grid. These watermarks represent thresholds above or below which address usage is unexpected and might warrant your attention.

- **Enable DHCP Thresholds:** Select this check box to enable the feature.
 - **High-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range exceeds this number, the DHCP server makes a syslog entry. The default is 95.
 - **Low-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below this number, the DHCP server makes a syslog entry. The default is 0. Address usage must initially exceed the low-water mark threshold and then dip below it before the appliance considers low address usage an event requiring an alert.
3. Optionally, you can click **Toggle Expert Mode** to display the **Logging** tab, where you can enable the managing member to log the lease events of the Microsoft server. This setting is inherited from the Grid. You can override that setting by clicking **Override**, and then selecting or clearing the **Log Lease Events from DHCP server** check box.
 4. Save the configuration and click **Restart** if it appears at the top of the screen.

Controlling the DHCP Service of a Microsoft Server

You can start and stop the DHCP service of a managed Microsoft server from Grid Manager as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members/Servers** tab -> **Members/Servers** -> *ms_server* check box.
2. Expand the Toolbar and click **Start** or **Stop**.
3. Click **Yes** when the confirmation dialog appears.

Disabling and Removing Microsoft DHCP Servers

If you remove a Microsoft server as a managed server, Grid Manager deletes all the DHCP ranges, leases, and fixed addresses associated with the server. It also deletes networks that were assigned only to the Microsoft server. It does not delete a network if it was assigned to other Microsoft servers as well.

When you disable a Microsoft server, the managing Grid member terminates any on-going synchronization and restarts synchronization only when the server is re-enabled. The DHCP data associated with that server is preserved in the same state until synchronization resumes.

For information on removing and disabling Microsoft servers, see [Disabling Synchronization](#) on page 962 and [Removing a Managed Microsoft Server](#) on page 962.

Modifying DHCP Server Assignments

If you disable a Microsoft DHCP server or take it offline for maintenance purposes, for example, you can assign its scopes to a member DHCP server.

Following are the tasks to reassign scopes from a Microsoft server to a member DHCP server:

1. Set the server assignments of all fixed addresses in the scope to “None”.
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed_address* check box, and then click the Edit icon. You can change the server assignment in the **General** tab of the *Fixed Address* editor.
2. Set the server assignments of all address ranges served by the Microsoft server to “None”.
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon. You can change the server assignment in the **General** tab of the *DHCP Range* editor.
3. Change the sever assignments of the networks by deleting the Microsoft server and replacing it with a member DHCP server.
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon. You can change the server assignment in the **Member Assignment** tab of the *Network* editor. contains the following basic tabs from which you can modify data:
4. Modify the server assignments of all address ranges and specify the member DHCP server.
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr_range* check box, and then click the Edit icon. You can change the server assignment in the **General** tab of the *DHCP Range* editor.
5. Restart services.

The member DHCP server starts granting lease requests after the restart. Note that you do not need to clear the leases that were active on the Microsoft server, because the member automatically clears them when you change the DHCP server assignment.



PART 7 MONITORING AND REPORTING

This section explains how to use the different monitoring and reporting tools, including DHCP fingerprint detection and SNMP. It includes the following chapters:

- [Chapter 35, *Monitoring the Appliance*](#), on page 1003
- [Chapter 36, *DHCP Fingerprint Detection*](#), on page 1031
- [Chapter 37, *Monitoring with SNMP*](#), on page 1037
- [Chapter 38, *Infoblox Reporting Solution*](#), on page 1113



Chapter 35 Monitoring the Appliance

This chapter describes the status icons that indicate the state of appliances, services, database capacity, Ethernet ports, HA, and Grid replication. It also explains how to use the various logs and the traffic capture tool to monitor a NIOS appliance.

This chapter contains the following sections:

- [Viewing Status](#) on page 1004
 - [Grid Status](#) on page 1004
 - [Member Status](#) on page 1004
 - [Viewing Hardware Status](#) on page 1010
- [Monitoring Services](#) on page 1011
 - [Service Status](#) on page 1011
 - [Monitoring Grid Services](#) on page 1011
 - [Monitoring Member Services](#) on page 1012
- [Using a Syslog Server](#) on page 1012
 - [Specifying Syslog Servers](#) on page 1013
 - [Configuring Syslog for Grid Members](#) on page 1014
 - [Setting DNS Logging Categories](#) on page 1015
 - [Viewing the Syslog](#) on page 1016
 - [Searching in the Syslog](#) on page 1017
 - [Downloading the Syslog File](#) on page 1017
- [Monitoring Tools](#) on page 1018
 - [Using the Audit Log](#) on page 1018
 - [Viewing the Replication Status](#) on page 1020
 - [Using the Traffic Capture Tool](#) on page 1021
 - [Using the Capacity Report](#) on page 1022
 - [Participating in the Customer Experience Improvement Program](#) on page 1023
 - [Monitoring DNS Transactions](#) on page 1024
 - [Viewing DNS Alert Indicator Status](#) on page 1026
 - [Configuring DNS Alert Thresholds](#) on page 1026
 - [DHCP Fingerprint Detection](#) on page 1031

VIEWING STATUS

Grid Manager provides tools for monitoring the status of the Grid, members, and services. You can monitor overall Grid and member status from the Dashboard, which provides a high-level view of your Grid, members and IP address data, and easy access to tasks. For information, see [Dashboards](#) on page 97.

Grid Manager also displays status icons to indicate the state of appliances, services, database capacity, Ethernet ports, HA, and Grid replication. Depending on your appliance, Grid Manager can display status icons for the power supplies as well as icons to indicate the state of the RAID array and disk controller backup battery.

You can monitor detailed status of the Grid, members, and services, and then decide how to manage them. Note that when any member or service encounters issues, the appliance sends SNMP traps. For information, see [Monitoring with SNMP](#) on page 1037.

Grid Status

You can monitor the overall status of the Grid using the *Grid Status* widget on the Dashboard. For information, see [Grid Status](#) on page 121.

You can also view the Grid status from the **Grid Manager** tab. To view Grid status, from the **Grid** tab, select the **Grid Manager** tab. Grid Manager displays the overall Grid status and status of all Grid services. The Grid status represents the status of the most critical members or services in the Grid. When all Grid members are running properly, the overall Grid status is green. When one of the members has operational problems, the overall Grid status is red. Grid Manager lists all Grid members in the **Members** tab so you can identify which member has issues. For information, see [Member Status](#).

In addition, the service bar below the Grid status lists the status of all licensed services—DHCP, DNS, TFTP, HTTP (File Distribution), FTP, NTP, bloxTools, DNS Accelerator—in the Grid. When you click a service link, Grid Manager displays detailed information about the selected service running on all members. For information, see [Monitoring Services](#) on page 1011. Grid Manager also provides icons you can use to edit Grid properties and bookmark the page.

Member Status

You can monitor the overall status, such as the memory usage and system temperature, of a Grid member or an independent appliance using the *Member Status (System Status)* widget on the Dashboard. For information, see [Member Status \(System Status\)](#) on page 123.

To monitor detailed status of a member, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.

In the **Members** tab, Grid Manager displays the Grid Master first and then all other members in alphabetical order. If a member is an HA pair, you can click the arrow next to the member row to view information about the active and passive nodes. Grid Manager can display the following information:

- **Name:** The name of the member.
- **HA:** Indicates whether the member is an HA pair.
- **Status:** The service status of the member. For a vNIOS appliance whose license is revoked and is still operating in the Grid, Grid Manager displays a license violation warning here. You should immediately remove this member from the Grid.
- **IPv4 Address:** The IP address of the appliance, or the VIP of an HA pair. An **IPv6 Address** column is available but is hidden by default.
- **Identify:** This field appears only if your appliance has the unit identification button. This can be On or Off. When you identify the appliance by pressing the UID button on the appliance or through the GUI or CLI command, this field displays On. Otherwise, this is Off.
- **DHCP, DNS, TFTP, HTTP, FTP, NTP, bloxTools, Captive Portal, DNS Accelerator usage, Reporting:** The status icons indicate whether these services are running properly. The DNS accelerator usage feature is only applicable to the IB-4030 appliance. For information, see [Service Status](#) on page 1011.
- **Hardware Type:** The hardware type of the appliance.
- **Hardware Model:** The hardware model of the appliance.

- **Serial Number:** The serial number of the appliance.
- **DB Utilization:** The current percentage of the database in use.
- **Comment:** Information about the member.




To turn the identification button on or off on the member, click the Hardware Identify icon. Grid Manager displays a panel with the appliance name, status, and IP address. Hover your mouse over the row and click **Turn On** to turn the identification button on, or click **Turn Off** to turn it off.

To view detailed status, select a member check box, and then click the Detailed Status icon. Grid Manager displays the *Detailed Status* panel. If the selected member is an HA pair, Grid Manager displays the information in two columns, one for the active node and the other for the passive. The *Detailed Status* panel provides detailed information described in the following sections.

You can modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only.

Appliance Status




The status icon indicates the operational status of a Grid member and a general description of its current operation. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The appliance is operating normally in a “Running” state.
	Yellow	The appliance is connecting or synchronizing with its Grid Master.
	Red	The Grid member is offline, is not licensed (that is, it does not have a DNSone license with the Grid upgrade that permits Grid membership), is upgrading or downgrading, or is shutting down.

The following are descriptions that may appear: Running, Offline, Error, and Warning.



Disk Usage

Grid Manager displays the percentage of the data partition of the hard disk drive that is currently in use on the selected Grid member. It also displays whether the percentage of usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 85% and reset value is 70%. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The disk usage is either below the reset value or has not yet reached the trigger value.
	Yellow	The disk usage is decreasing from the trigger value, but has not yet reached the reset value.
	Red	The disk usage has exceeded the trigger value.




DB Capacity Usage

Grid Manager displays the current percentage of the database in use on the selected Grid member. It also describes whether the usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 80% and reset value is 70%. For information, see [Using the Capacity Report](#) on page 1022. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The database capacity is either below the reset value or has not yet reached the trigger value.
	Yellow	The database capacity is decreasing from the trigger value, but has not yet reached the reset value. When the capacity exceeds the trigger value, the icon changes from green to yellow.



LAN1/LAN2 Ports, HA Port, and MGMT Port

Grid Manager displays the IP address of the port. The status icons for these ports indicate the state of their network connectivity.

Icon	Color	Meaning
	Green	The port is properly connected to a network. Grid Manager displays the IP address of the network.
	Red	The port is not able to make a network connection.
	Gray	The port is disabled.




LCD

The LCD status icon indicates its operational status.

Icon	Color	Meaning
	Green	The LCD is functioning properly.
	Yellow	The LCD process is not running.




Memory Usage

Grid Manager displays the current percentage of system memory in use on the selected Grid member. It also describes whether the usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 90% and reset value is 80%. You can see more details about memory usage through the CLI command: `show memory`. The status icon can be one of the following.

Icon	Color	Meaning
	Red	The memory usage has exceeded the trigger value.
	Yellow	The memory usage is decreasing from the trigger value, but has not yet reached the reset value.
	Green	The memory usage is either below the reset value or has not yet reached the trigger value.

Swap Usage



Grid Manager displays the current percentage of swap area in use on the selected Grid member. It also describes whether the usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 20% and reset value is 10%. The status icon can be one of the following:

Icon	Color	Meaning
	Red	The memory usage has exceeded the trigger value.
	Yellow	The memory usage is decreasing from the trigger value, but has not yet reached the reset value.
	Green	The memory usage is either below the reset value or has not yet reached the trigger value.

FAN



The status icon indicates whether the fan is functioning properly. The corresponding description displays the fan speed. The status icon and fan speed are displayed for Fan1, Fan2, and Fan3.

Note: vNIOS appliances on VMware do not monitor or report the fan speed.

Icon	Color	Meaning
	Green	The fan is functioning properly.
	Red	The fan is not running.





Power Supply

The Infoblox-1552-A, -1852-A, -2000-A, and -4010 have redundant power supplies. The power supply icon indicates the operational status of the power supplies.

Icon	Color	Meaning
	Green	The power supplies are functioning properly.
	Red	One power supply is not running. To find out which power supply failed, check the LEDs of the power supplies.

NTP Synchronization

The status icon indicates the operational status of the current NTP synchronization status.

Icon	Color	Meaning
	Green	The NTP service is enabled and running properly.
	Yellow	The NTP service is enabled, and the appliance is synchronizing its time.
	Red	The NTP service is enabled, but it is not running properly or is out of synchronization.
	Gray	The NTP service is disabled.

CPU Temperature

This icon is always green. The description reports the CPU temperature.

Note: vNIOS appliances on VMware do not monitor or report the CPU temperature.



System Temperature

This icon is always green. The description reports the system temperature.

Note: vNIOS appliances on VMware do not monitor or report the system temperature.




CPU Usage

Grid Manager displays the current percentage of the CPU usage on the selected Grid member. The maximum is 100%. It also describes whether the CPU usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 81% and reset value is 70%. You can see more details about CPU usage through the CLI command: `show CPU`. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The CPU usage is either below the reset value or has not yet reached the trigger value.
	Red	The CPU usage has exceeded the trigger value.

RAID

For the Infoblox-2000-A and -4010, Grid Manager displays one of the following icons to indicate the status of each disk in the RAID array. Next to the status icon is a summary that includes the disk number, the operational status of the disk, and the disk type. Grid Manager also displays a RAID summary with an overall array status icon and the percentage at which the array is currently operating.

Icon	Color	Meaning
	Green	The RAID array or the disk is functioning properly.
	Yellow	A new disk has been inserted and the RAID array is rebuilding.
	Red	The RAID array or the disk is degraded. At least one disk in the array is not functioning properly. Grid Manager lists the disks that are online. Replace only the disks that are offline.




In the event of a disk failure, you must replace the failed disk with one that is qualified and shipped from Infoblox and has the same disk type as the rest of the disks in the array. The appliance displays information about mismatched disks. The disk type of the Infoblox-2000-A can be one of the following:

- IB-Type 1: Infoblox supported disk type
- IB-Type 2: Infoblox supported disk type
- Unk: Unknown disk type that Infoblox does not support

Infoblox-4010 uses only the IB-Type 3 disk type. All disk drives in the array must have the same disk type for the array to function properly. You can have either IB-Type 1, IB-Type 2, or IB-Type 3, but you cannot mix both in the array. When you have a mismatched disk in the array, you must promptly replace the disk with a replacement disk from Infoblox to avoid operational issues.

RAID Battery

The icon indicates the status of the disk controller backup battery on the Infoblox-2000-A or -4010.

Icon	Color	Meaning
	Green	The battery is charged. The description indicates the estimated number of hours of charge remaining on the battery.
	Yellow	The battery is charging.
	Red	The battery is not charged.

Viewing the Grid Node Tree

Navigate to the **Grid** -> **Grid Manager** -> **Visualization** tab to view a graphical representation of the Grid, with its members represented as nodes in the tree. Each member is labeled with its hostname. You can click **Display Node Labels** on the left panel to display or hide the labels.

By default, the Grid Master is the root node at the center of the tree. It is represented by a color-coded icon connected to its members. You can then click a member to re-center the tree on that node. The left panel displays information about the member that is at the center of the node tree.

In the node tree, the shape of the icons indicate the role of the member in the Grid:

- Circle: Grid Master
- Ellipse: Grid Members

The colors of the icons indicate the status of the member:

- Green: The member is online and functioning properly.
- Grey: The member has not joined the Grid.
- Red: The member has operational problems.

The connectors indicate the connection status between the Grid Master and the member.

- Blue Line: Connects the Grid Master with online Grid members
- Thick White Line: Connects the Grid Master with Grid Master Candidates
- Dashed Line Connector: Connects the Grid Master with offline Grid members

The node tree includes zooming and panning capabilities to enable quick navigation and selection among multiple nodes. You can also hover your mouse over a node to view node information. It displays the same information as that displayed on the left panel, when a node is at the center of the tree.

For the Grid Master:

- Grid name
- Standalone or HA
- Number of members in the Grid
- Status of each protocol running on the Grid
- Grid status

For a Member:

- Member name
- Standalone or HA
- HA Status if HA pair
- Status of each protocol running on that member

Viewing Hardware Status

You can view the link activity and connection speed of an Ethernet port by looking at its Link/Act and Speed LEDs on the appliance. The status the LEDs convey through their color and illumination (steady glow or blinking) are presented in the following tables.

For Infoblox-2000-A Appliances

MGMT and HA Ports

Label	Color	Port Status
Link/Act	Steady Orange	Link is up but inactive
	Blinking Orange	Link is up and active
	Dark	Link is down


LAN Ports

Label	Color	Port Status
Link/Act	Steady Green	Link is up but inactive
	Blinking Green	Link is up and active
	Dark	Link is down

MGMT, HA, and LAN Ports

Label	Color	Port Status
Speed	Steady Amber	1000 Mbps
	Steady Green	100 Mbps
	Dark	10 Mbps

For Infoblox--1050-A, -1550-A, -1552-A, and -1852-A Appliances



Label	Color	Port Status
Link/Act	Steady Green	Link is up but inactive
	Blinking Green	Link is up and active
	Dark	Link is down
Speed	Steady Amber	1000 Mbps
	Steady Green	100 Mbps
	Dark	10 Mbps

MONITORING SERVICES

The Grid or device status icon and the service icon indicates whether a service running on a member or an independent appliance is functioning properly or not.

Service Status

After you enable any of the services—DHCP, DNS, TFTP, HTTP (for file distribution), FTP, NTP, bloxTools, and Captive Portal—the appliance indicates their status as follows:

Icon	Color	Meaning
	Green	The service is enabled and running properly.
	Yellow	The service is enabled, but there may be some issues that require attention.
	Red	The service is enabled, but it is not running properly. (A red status icon can also appear temporarily when a service is enabled and begins running, but the monitoring mechanism has not yet notified Grid Manager.)
	Gray	The service is not configured or it is disabled.

Monitoring Grid Services

The status icon of a Grid service represents the status of the most critical service in the Grid. For example, if the Grid DHCP status icon is red, the DHCP service on one of the members in the Grid is not running properly. You can click the DHCP service link to view the service status of all Grid members and identify which member has a service problem. You can then decide to start or stop the service, or modify the service configuration on that member.

To monitor a Grid service:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click a service link.
2. Grid Manager displays the following information in the **Services** tab:
 - **Name:** The name of the member.
 - **Service Status:** The current status of the service.

- **IP Address:** The IP address of the appliance or the VIP of an HA pair.
- **Comments:** Information about the member or service.
- **Site:** The site to which the member belongs. This is one of the predefined extensible attributes.

You can select available extensible attributes for display.

3. Optionally, click the Edit icon next to the service name to edit the Grid properties for the service.

or

Select a member check box, and do one of the following:

- Click the Edit icon to edit the member service configuration. Grid Manager displays the editor for the corresponding service. For example, when you edit the DHCP service, Grid Manager displays the *Member DHCP Configuration* editor.
- Click the Start icon to start the service.
- Click the Stop icon to stop the service.

Grid Manager updates the service status based on your action.

Monitoring Member Services

You can view detailed service status on a selected member. Optionally, you can start and stop a service, and edit the service configuration.

To monitor a member service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click the Manage Member Services icon.

In the *Manage Services* panel, Grid Manager displays the following information:

- **Service:** The name of the service.
- **Status:** The current status of the service running on the member.
- **Description:** The description of the status. Grid Manager displays the percentage of usage for the TFTP, HTTP (File Distribution), FTP, and bloxTools services.

2. Optionally, mouse over a service and do one of the following:

- **Start/Stop Service:** Click this icon to start or stop the selected service. For example, when the DNS service is currently stopped, the appliance starts the service when you click this icon.
- **Edit Service:** Click this icon to edit the selected service. Grid Manager displays the corresponding editor. For example, when you click the Edit Service icon for DNS, Grid Manager displays the *Member DNS Configuration* editor.

Click the Refresh icon to update the service status.

USING A SYSLOG SERVER

Syslog is a widely used mechanism for logging system events. NIOS appliances generate syslog messages that you can view through the Syslog viewer and download to a directory on your management station. In addition, you can configure a NIOS appliance to send the messages to one or more external syslog servers for later analysis. Syslog messages provide information about appliance operations and processes. NIOS appliances include syslog messages generated by the bloxTools service. You can also include audit log messages and specific BIND messages among the messages the appliance sends to the syslog server.

In addition to saving system messages to a remote syslog server, a NIOS appliance also stores the system messages locally. When the syslog file reaches its maximum size, which is 300 MB for Infoblox appliances and VMware virtual appliances, and 20 MB for Riverbed virtual appliances, the appliance automatically writes the file into a new file by adding a .0 extension to the first file and incrementing subsequent file extensions by 1.

Files are compressed during the rotation process, adding a .gz extension following the numerical increment (*file.#.gz*). The sequential incrementation goes from zero through nine. When the eleventh file is started, the tenth log file (*file.9.gz*) is deleted, and subsequent files are renumbered accordingly. For example, the current log file moves to *file.0.gz*, the previous *file.0.gz* moves to *file.1.gz*, and so on through *file.9.gz*. A maximum of 10 log files (0-9) are kept.

You can set syslog parameters at the Grid and member levels. At the member level, you can override Grid-level syslog settings and enable syslog proxy.

This section includes the following topics:

- [Specifying Syslog Servers](#) on page 1013
- [Configuring Syslog for Grid Members](#) on page 1014
- [Setting DNS Logging Categories](#) on page 1015
- [Viewing the Syslog](#) on page 1016
- [Searching in the Syslog](#) on page 1017
- [Downloading the Audit Log](#) on page 1020

Specifying Syslog Servers

To configure a NIOS appliance to send messages to a syslog server:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Monitoring** tab, and then complete the following:

Syslog

In addition to storing the syslog on a Grid member, you can configure the Grid to send the log to an external syslog server.

- **Syslog size (MB):** Specify the maximum size for a syslog file. Enter a value between 10 and 300. The default is 300.
When the syslog file reaches the size you enter here, the appliance automatically writes the file into a new file by adding a .0 extension to the first file and incrementing subsequent file extensions by 1.
- **Log to External Syslog Servers:** Select this to enable the appliance to send messages to a specified syslog server.
- Grid Manager displays the current syslog servers in the table. To define a new syslog server, click the Add icon. Grid Manager adds a row to the table. Enter the following by clicking each field in the row:
 - **Address:** Enter the IP address of a syslog server. Entries may be an IPv4 or IPv6 address.
 - **Transport:** From the drop-down list, select whether the appliance uses TCP or UDP to connect to the external syslog server.
 - **Interface:** From the drop-down list, select the interface through which the appliance sends syslog messages to the syslog server.
 - **Source:** From the drop-down list, select which syslog messages the appliance sends to the external syslog server:
 - **Internal:** The appliance sends syslog messages that it generates.
 - **External:** The appliance sends syslog messages that it receives from other devices, such as syslog servers and routers.
 - **Any:** The appliance sends both internal and external syslog messages.
 - **Port:** Enter the destination port number. The default is 514.
 - **Severity:** Choose a severity filter from the drop-down list. When you choose a severity level, the appliance sends log messages with the selected level and the levels above it. The severity levels range from the lowest, **debug**, to the highest, **emerg**. For example, if you choose **debug**, the appliance sends all syslog messages to the server. If you choose **err**, the appliance sends messages with severity levels **err**, **crit**, **alert**, and **emerg**.

- **emerg:** Panic or emergency conditions. The system may be unusable.
 - **alert:** Alerts, such as NTP service failures, that require immediate actions.
 - **crit:** Critical conditions, such as hardware failures.
 - **err:** Error messages, such as client update failures and duplicate leases.
 - **warning:** Warning messages, such as missing keepalive options in a server configuration.
 - **notice:** Informational messages regarding routine system events, such as “starting BIND”.
 - **info:** Informational messages, such as DHCPACK messages and discovery status.
 - **debug:** Messages that contain information for debugging purposes, such as changes in the latency timer settings and AD authentication failures for specific users.
- **Copy Audit Log Messages to Syslog:** Select this for the appliance to include audit log messages it sends to the syslog server. This function can be helpful for monitoring administrative activities on multiple appliances from a central location.
 - **Syslog Facility:** This is enabled when you select **Copy audit log messages to syslog**. Select the facility that determines the processes and daemons from which the log messages are generated.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Configuring Syslog for Grid Members

You can override Grid-level syslog settings and enable syslog proxy for individual members. When you enable syslog proxy, the member receives syslog messages from specified devices, such as syslog servers and routers, and then forwards these messages to an external syslog server. You can also enable appliances to use TCP for sending syslog messages. Using TCP is more reliable than using UDP; this reliability is important for security, accounting, and auditing messages sent through the syslog. Note that you cannot enable syslog proxy for Grid members, if they are configured on a Grid Master.

To configure syslog parameters for a member:

1. From the **Grid** tab, select the **Grid Manager** tab → **Members** tab → *member* check box, and then click the Edit icon.
2. In the *Grid Member Properties* editor, select the **Monitoring** tab → **Basic** tab, click **Override** in the Syslog section, and then complete the fields as described in [Specifying Syslog Servers](#) on page 1013.

In addition to storing the system log on a Grid member, you can configure a member to send the log to a syslog server.

3. Select the **Advanced** tab and complete the following:
 - **Enable syslog proxy:** Select this to enable the appliance to receive syslog messages from other devices, such as syslog servers and routers, and then forward these messages to an external syslog server.
 - **Enable listening on TCP:** Select this if the appliance uses TCP to receive messages from other devices. Enter the number of the port through which the appliance receives syslog messages from other devices.
 - **Enable listening on UDP:** Select this if the appliance uses UDP to receive messages from other devices. Enter the number of the port through which the appliance receives syslog messages from other devices.
 - **Proxy Access Control:** Select one of the following to configure access control when receiving syslog messages from specific syslog servers or routers:
 - **None:** Select this if you do not want to configure syslog proxy. The appliance receives syslog messages from all syslog servers or routers. This is selected by default.
 - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. This does not support TSIG key based ACEs. When you select this, the appliance permits clients that have **Allow** permission in the named ACL to allow syslog messages from specific syslog servers or routers. You can click **Clear** to remove the selected named ACL.
 - **Set of ACLs:** Select this to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Grid Manager adds a row to the table.

- **IPv4 Address or IPv6 Address:** Select this to add an IPv4 or IPv6 address entry. Click the **Value** field and enter the address. The default permission is **Allow**, which means that the appliance allows access to and from this device. You can change this to **Deny** to block access.
- **IPv4 Network or IPv6 Network:** Select this to add an IPv4 or IPv6 network entry. Click the **Value** field and enter the network. The default permission is **Allow**, which means that the appliance allows syslog messages sent by this network. You can change this to **Deny** to block access.
- **Any Address/Network:** Select this to allow or deny access to all IPv4 and IPv6 addresses and networks. The default permission is **Allow**, which means that the appliance allows syslog messages sent by all addresses and networks. You can change this to **Deny** to block access.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to group and put into a named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an IPv4 network and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Setting DNS Logging Categories

You can specify logging categories you want the syslog to capture. Furthermore, you can filter these messages by severity at the Grid and member levels. For information about severity types, see [Specifying Syslog Servers](#) on page 1013.

Note that logging DNS queries and responses in the syslog will significantly affect system performance. Before you enable query and response logging, ensure that your system has sufficient CPU capacity. Alternatively, if you have a Reporting license, you can capture DNS query and response information and forward it to an external server. For more information about this feature, see [Configuring the Capture of DNS Queries and Responses](#) on page 1123.

To specify logging categories:

1. From the **Data Management** tab, select the **DNS** tab, and then click **Grid DNS Properties** from the Toolbar.
or
From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid_member* check box, and then click the Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Expert Mode** if the editor is in the basic mode, select the **Logging** tab, and then complete the following:
 - **Logging Facility:** Select a facility from the drop-down list. This is the location on the syslog server to which you want to sort the DNS logging messages.
 - **Logging Category:** Select one or more of these log categories:
 - **general:** Records the BIND messages that are not specifically classified.
 - **client:** Records client requests.
 - **config:** Records the configuration file parsing messages.
 - **database:** Records BIND's internal database processes.
 - **dnssec:** Records the DNSSEC-signed responses.
 - **lame servers:** Records bad delegation instances.
 - **network:** Records the network operation messages.
 - **notify:** Records the asynchronous zone change notification messages.
 - **queries:** Records the DNS queries. Note that enabling the logging of queries and responses will significantly affect system performance. Ensure that your system has sufficient CPU capacity before you enable DNS query logging.
 - **resolver:** Records the DNS resolution instances, including recursive queries from resolvers.

- **rpz:** Records log messages when responses are modified through RPZs or for which explicit passthru were invoked in the RPZs. This check box is not selected by default.
- **responses:** Records DNS responses. Note that enabling the logging of queries and responses will significantly affect system performance. Ensure that your system has sufficient CPU capacity before you enable DNS response logging.
- **security:** Records the approved and denied requests.
- **transfer-in:** Records zone transfer messages from the remote name servers to the appliance.
- **transfer-out:** Records zone transfer messages from the NIOS appliance to remote name servers.
- **update:** Records the dynamic update instances.
- **update-security:** Records the security updates.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Viewing the Syslog

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right corner, select the Grid member on which you want to view the syslog.
3. Optionally, use the filters to narrow down the system messages you want to view. Click **Show Filters** to enable the filters. Configure the filter criteria, and then click **Apply**.

Based on your filter criteria (if any), Grid Manager displays the following in the *Syslog* viewer:

- **Timestamp:** The date, time, and time zone of the log message. The time zone is the time zone configured on the member.
- **Facility:** The location on the syslog server that determines the processes and daemons from which the log messages are generated.
- **Level:** The severity of the message. This can be ALERT, CRITICAL, DEBUG, EMERGENCY, ERROR, INFO, NOTICE, or WARNING.
- **Server:** The name of the server that logs this message, plus the process ID.
- **Message:** Detailed information about the task performed.

Note: If the selected member is an HA pair, Grid Manager displays the syslog in two tabs—**Active** and **Passive**. Click the corresponding tab to view the syslog for each node.

You can also do the following in the *Syslog* viewer:

- Toggle between the single line view and the multi-line view for display.
- Navigate to the next or last page of the file using the paging buttons.
- Refresh the syslog output with newly logged messages.
- Click the Follow icon to have the appliance automatically refresh the log every five seconds.
- Clear the contents of the syslog.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Print the report or export it in CSV format.
- Bookmark the syslog page.

Searching in the Syslog

Instead of paging through the syslog to locate messages, you can have the appliance search for syslog messages with certain text strings. To search for specific messages:

- Enter a search value in the search field below the filters, and then click the **Search** icon.
The appliance searches through the syslog and highlights the search value in the viewer. You can use the arrow keys next to the Search icon to locate the previous or next message that contains the search value.

Downloading the Syslog File

You can download the syslog file to a specified directory, if you want to analyze it later.

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab, and then click the Download icon.
2. Navigate to a directory where you want to save the file, optionally change the file name (the default names are *node_1_sysLog.tar.gz* and *node_2_sysLog.tar.gz*), and then click **OK**. If you want to download multiple syslog files to the same location, rename each downloaded file before downloading the next.

Note: If your browser has a pop-up blocker enabled, you must turn off the pop-up blocker or configure your browser to allow pop-ups for downloading files.

MONITORING TOOLS

You can use the audit log, the replication status, the traffic capture tool, and the capacity report in a Grid or HA pair to monitor administrative activities and capture traffic for diagnostic purposes. You can also use CLI commands to monitor certain DNS transactions.

This section includes the following topics:

- [Using the Audit Log](#)
- [Viewing the Replication Status](#) on page 1020
- [Using the Traffic Capture Tool](#) on page 1021
- [Using the Capacity Report](#) on page 1022
- [Participating in the Customer Experience Improvement Program](#) on page 1023
- [Monitoring DNS Transactions](#) on page 1024
- [Viewing DNS Alert Indicator Status](#) on page 1026
- [Configuring DNS Alert Thresholds](#) on page 1026

In addition, if Grid members manage Microsoft servers, Grid Manager creates a synchronization log file for each managed Microsoft server. For information, see [Viewing Synchronization Logs](#) on page 965.

Using the Audit Log

The audit log contains a record of all Infoblox administrative activities. It provides the following detailed information:

- Timestamp of the change. If you have different admin accounts with different time zone settings, the appliance uses the time zone of the admin account that you use to log in to the appliance to display the date and timestamp.
- Administrator name
- Changed object name
- New value of the object. If you change multiple properties of an object, the audit log lists all changes in a comma-separated log entry. You can also search the audit log to find the new value of an object.

The appliance logs the following successful operations:

- Logins to Grid Manager and the API.
- Logout events, including when users log out by clicking the **Logout** button, when the Grid Manager GUI times out, and when users are logged out due to an error.
- Write operations such as the addition, modification, and deletion of objects.
- System management operations such as service restarts and appliance reboots.
- Scheduled tasks such as adding an A record or modifying a fixed address.

Enabling Audit Log Rolling

When the audit log reaches its maximum size, which is 100 MB, the appliance automatically writes the file into a new file by adding a .0 extension to the first file and incrementing subsequent file extensions by 1. Files are compressed during the rotation process, adding a .gz extension following the numerical increment (*file.#.gz*). The sequential incrementation goes from zero through nine. When the eleventh file is started, the tenth log file (*file.9.gz*) is deleted, and subsequent files are renumbered accordingly. For example, the current log file moves to *file.0.gz*, the previous *file.0.gz* moves to *file.1.gz*, and so on through *file.9.gz*. A maximum of 10 log files (0-9) are kept. To list the audit log files and their sizes, log in to the Infoblox CLI and execute the `show logfiles` command.

To enable audit log rolling:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Security** tab, and then select **Enable Audit Log Rolling**.

Specifying the Audit Log Type

Select either the **Detailed** (default) or **Brief** audit log type as follows:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **General** tab, and then select one of the following:
 - **Detailed:** This is the default type. When you select this, Grid Manager displays detailed information on all administrative changes such as the timestamp of the change, administrator name, changed object name, and the new values of all properties in the logged message.
 - **Brief:** Provides information on administrative changes such as the changed object name and action in the log message. The logged message does not show timestamp or admin name.

Viewing the Audit Log

To view an audit log:

1. From the **Administration** tab, select the **Logs** tab -> **Audit Log** tab.
2. Optionally, use the filters to narrow down the audit log messages you want to view. Click **Show Filters** to enable the filters. Configure the filter criteria, and then click **Apply**.

Based on your filter criteria (if any), Grid Manager displays the following in the *Audit Log* viewer:

- **Timestamp:** The date, time, and time zone the task was performed. The time zone is the time zone configured on the member.
- **Admin:** The admin user who performed the task.
- **Action:** The action performed. This can be CALLED, CREATED, DELETED, LOGIN_ALLOWED, LOGIN_DENIED, MESSAGE, and MODIFIED.
- **Object Type:** The object type of the object involved in this task. This field is not displayed by default. You can select this for display.
- **Object Name:** The name of the object involved in this task.
- **Execution Status:** The execution status of the task. Possible values are **Executed**, **Normal**, **Pending Approval** and **Scheduled**.
- **Message:** Detailed information about the performed task.

You can also do the following in the log viewer:

- Toggle between the single line view and the multi-line view for display.
- Navigate to the next or last page of the file using the paging buttons.
- Refresh the audit log view.
- Click the Follow icon to have the appliance automatically refresh the log every five seconds.
- Download the log.
- Clear the contents of the audit log.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Export or print the content of the log.

Searching in the Audit Log

Instead of paging through the audit log file to locate messages, you can have the appliance search for messages with certain text strings.

To search for specific messages:

- Enter a search value in the search field below the filters, and then click the **Search** icon.
The appliance searches through the audit log file and highlights the search value in the viewer. You can use the arrow keys next to the Search icon to locate the previous or next message that contains the search value.

Downloading the Audit Log

You can download the audit log file to a specified directory, if you want to analyze it later.

To download an audit log file:

1. From the **Administration** tab, select the **Logs** tab -> **Audit Log** tab, and then click the Download icon.
2. Navigate to a directory where you want to save the file, optionally change the file name (the default name is *auditLog.tar.gz*), and then click **OK**. If you want to download multiple audit log files to the same location, rename each downloaded file before downloading the next.

Note: If your browser has a pop-up blocker enabled, you must turn off the pop-up blocker or configure your browser to allow pop-ups for downloading files.

Viewing the Replication Status

The *Replication Status* panel reports the status of the database replication between Grid members and Grid Master, and between the two nodes in an independent HA pair. You can use this information to check the health of the Grid and HA pair activity.

To view the current replication status, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Toggle Replication Status View**.

Grid Manager can display the following replication information for each member:

- **Name:** The FQDN (fully qualified domain name) of the appliance.
- **Send Queue:** The size of the queue from the Grid Master to the Grid member.
- **Last Send:** The timestamp of the last replication information sent by the Grid Master.
- **Receive Queue:** The size of the queue from the Grid member to the Grid Master.
- **Last Receive:** The timestamp of the last replication information sent received by the Grid Master.
- **Member Replication Status:** The replication status between the member and the Grid Master. Grid Manager displays the status in green when the status is fine or red when the member is offline.
- **HA Replication Status:** The HA replication status between the active and passive nodes. The status is at the member level, not at the node level. Grid Manager displays the status in red when one of the nodes is offline.
- **Status:** The current operational status of the appliance. The status can be one of the following:
 - **Green:** The appliance is operating normally in a “Running” state.
 - **Yellow:** The appliance is connecting or synchronizing with its Grid Master.
 - **Red:** The Grid member is offline, is not licensed (that is, it does not have a DNSone license with the Grid upgrade that permits Grid membership), is upgrading or downgrading, or is shutting down.
- **IP Address:** The IP address of the appliance.

- **DHCP, DNS, TFTP, HTTP,FTP, NTP, bloxTools, Captive Portal:** The current status of the service. The status can be one of the following:
 - **Green:** The service is enabled and running properly.
 - **Yellow:** The service is enabled, but there may be some issues that require attention.
 - **Red:** The service is enabled, but it is not running properly. A red status icon can also appear temporarily when a service is enabled and begins running, but the monitoring mechanism has not yet notified the Infoblox GUI.
 - **Gray:** The service is not configured or it is disabled.
- **Hardware Type:** The hardware type of the appliance, such as IB-1550-A.
- **Serial Number:** The serial number of the appliance.
- **DB Utilization:** The percentage of the database that is currently in use.
- **Comment:** Information about the appliance.
- **Site:** The location to which the member belongs. This is one of the predefined extensible attributes.
- **HA:** Indicates whether the member is an HA pair. If the member is an HA pair, Grid Manager displays the status of the HA pair.
- **Hardware Model:** The hardware model of the appliance.

You can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- Edit the properties of a member.
 - Click the check box beside a member, and then click the Edit icon.
- Delete a member.
 - Click the check box beside a member, and then click the Delete icon.
- Export or print the list.

Using the Traffic Capture Tool

You can capture the traffic on one or all of the ports on a NIOS appliance, and then view it using a third-party network protocol analyzer application, such as the Wireshark – Network Protocol Analyzer™.

The NIOS appliance saves all the traffic it captures in a .cap file and compresses it into a .tar.gz file. Your management system must have a utility that can extract the .tar file from the .gzip file, and an application that can read the .cap (capture) file format.

Note: This feature captures traffic of all the direct responses received from the cache accelerator on the IB-4030.

This section explains the process of capturing traffic, and how to download the traffic capture file to your management system. After that, you can extract the traffic capture file and view it with a third-party traffic analyzer application.

Note: The NIOS appliance always saves a traffic capture file as *tcpdumpLog.tar.gz*. If you want to download multiple traffic capture files to the same location, rename each downloaded file before downloading the next.

To capture traffic on a member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Traffic Capture** from the Toolbar.
2. In the *Traffic Capture* dialog box, complete the following:
 - **Member:** Grid Manager displays the selected member on which you want to capture traffic. If no member is displayed or if you want to specify a different member, click **Select**. When there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select one. You cannot capture traffic on an offline member.
 - **Interface:** Select the port on which you want to capture traffic. Note that if you enabled the LAN2 failover feature, the LAN and LAN2 ports generate the same output. (For information about the LAN2 failover feature, see [About Port Redundancy](#) on page 356.)
 - **LAN:** Select this to capture all the traffic the LAN port receives and transmits.
 - **MGMT:** Select this to capture all the traffic the MGMT port receives and transmits.
 - **LAN2:** Select to capture all the traffic the LAN2 port (if enabled) receives and transmits.
 - **All:** Select this to capture the traffic addressed to all ports. Note that the NIOS appliance only captures traffic that is addressed to it.
 - **LANx nnnn:** If you have configured VLANs on the LAN1 or LAN2 port, the appliance displays the VLANs in the format LANx nnnn, where x represents the port number and nnnn represents the associated VLAN ID.

Note: Riverbed virtual appliances support capturing traffic only on the LAN port.

- **Seconds to run:** Specify the number of seconds you want the traffic capture tool to run.
3. **Capture Control:** Click the *Start* icon to start the capture. A warning message appears indicating that this report will overwrite the existing file. Click **Yes**. You can click the Stop icon to stop the capture after you start it.
 4. **Uncompressed Capture File Size:** Click **Download** to download the captured traffic after the capture stops and then save the file. You can rename the file if you want. You cannot download the traffic report when the tool is running. Grid Manager updates the size of the report when the capture tool is running.
 5. Use terminal window commands (Linux) or a software application (such as StuffIt™ or WinZip™) to extract the contents of the .tar.gz file.
 6. When you see the traffic.cap file in the directory where you extract the .tar.gz file, open it with a third-party network protocol analyzer application.

Using the Capacity Report

You can view the capacity usage and object type information of an appliance in a capacity report. The capacity report displays capacity and object type information of an independent appliance, a Grid Master, or a Grid member. For an HA pair, the report displays information on the active node.

The top half of the panel displays a capacity summary, and the bottom half displays the object types the appliance supports and the total counts for each object type.

To view a capacity report:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* check box, and then click **Capacity Report** from the Toolbar.

The capacity summary contains the following information:

- **Name:** The name of the appliance.
- **Role:** The role of the appliance. The value can be **Grid Master**, **Grid Master Candidate**, **Grid Member**, or **Standalone**.
- **Hardware Type:** The type of hardware. For an HA pair, the report displays the hardware type for both the active and passive nodes.
- **Object Capacity:** The maximum number of objects the appliance can support.
- **Total Objects:** The total number of objects currently in the database.

- **% Capacity Used:** The percentage of the capacity in use.

The report categorizes object types you can manage through the appliance. It displays the following information for each object type:

- **Object Type:** The type of objects. For example, DHCP Lease, Admin Group, or PTR Record. For objects that are only used for internal system operations, the report groups and shows them under **Other**.
- **Total:** The total number of objects for a specific object type.

You can print the object type information or export it to a CSV file.

Participating in the Customer Experience Improvement Program

Administrators with superuser accounts can configure a Grid Master or an independent appliance to email reports monthly and after each upgrade to Infoblox Technical Support and other specified recipients. The reports are also included in support bundles that you download.

The reports provide status and event information about the Grid or independent appliance and its services. The report is an XML document that includes the following information:

- The phone home feature version.
- The report type, such as periodic and test.
- The time of the report.
- The Infoblox Support ID that was assigned to the account.
- Information about the Grid, such as its NIOS version, name, VIP, Grid Master hostname, LAN IP, and the number of Grid members and appliances in the Grid.
- The upgrade history of the Grid.
- Information about each Grid member, such as the hostname, IP address, status, role (such as standalone, master), and if the member is an HA pair. If the member is a peer in a DHCP failover association, the report also includes the DHCP failover status.
- Hardware information, such as the hardware type, serial number, HA status, and uptime.
- Information about the interfaces, such as the interface name and IP addresses.
- Resource usage information, such as CPU and system temperature, and CPU, database, disk, and memory usage.

Note that if the appliance is configured to send email notifications to an SMTP relay server, as described in [Notifying Administrators](#) on page 191, the appliance sends the phone home reports to the relay server as well.

To configure the Grid Master to email status reports:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **Customer Improvement** tab, and then complete the following:
 - **Participate in Infoblox Customer Experience Improvement Program:** Select the check box to send product usage data to Infoblox on a periodic basis. Infoblox uses this data to improve product functionality.
 - **Support ID:** Enter the Infoblox Support ID that was assigned to your account. It must be a number with four to six digits. Infoblox includes this ID in the data report.
 - **Send notifications to:**
 - **Infoblox Support:** Select this to email the reports to Infoblox Technical Support.
 - **Additional email addresses:** Optionally, you can specify up to 16 additional recipients. Click the Add icon and enter the email addresses of the recipients.
 - **Send Test Report:** Click this to send a test report to the specified recipients.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Monitoring DNS Transactions

The NIOS appliance provides tools for monitoring DNS transactions and mitigating cache poisoning from UDP (User Datagram Protocol) traffic on source port 53. Cache poisoning can occur when a DNS server accepts maliciously created unauthentic data. The DNS server ends up locally caching the incorrect entries and serving them to users that make the same DNS requests. In a maliciously created situation, the attacker can redirect Internet traffic from the legitimate host to another host that the attacker controls.

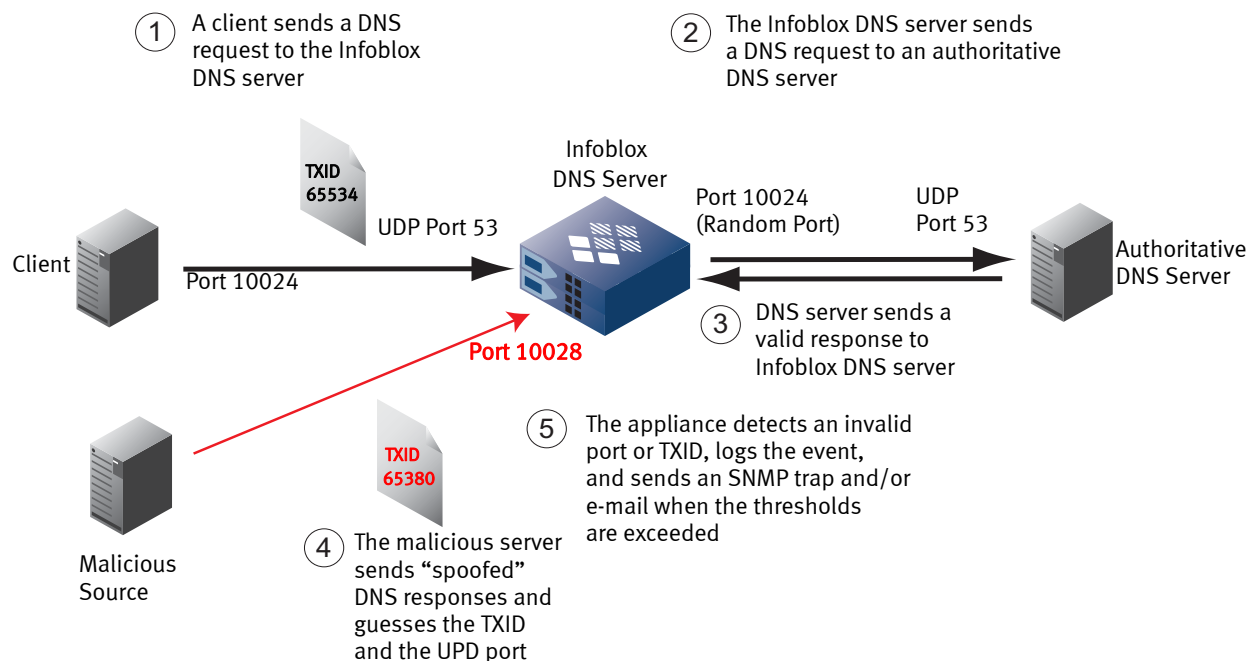
You can configure the appliance to track invalid DNS responses for recursive DNS queries. The appliance tracks DNS responses that arrive on invalid ports or have invalid TXIDs (DNS transaction IDs). Both invalid ports and invalid TXIDs could be indicators of cache poisoning. An invalid port is a DNS response that arrives from UDP (User Datagram Protocol) port 53 with either one of the following conditions:

- There are no outstanding DNS requests from the port on which the response arrives.
- The TXID of the DNS response matches the TXID of an outstanding request. However, the request was sent from a port other than the port on which the response arrives.

An invalid TXID is a DNS response that arrives from UDP port 53, and the TXID does not match the TXID of an outstanding DNS request.

[Figure 35.1](#) illustrates how the appliance detects an invalid port and an invalid TXID.

Figure 35.1 Invalid Port and Invalid TXID



Both invalid ports and invalid TXIDs could be indicators of DNS cache poisoning, although a small number of them is considered normal in situations where valid DNS responses arrive after the DNS queries have timed out. You can configure the appliance to track these indicators, and you can view their status. You can also configure thresholds for them. When the number of invalid ports or invalid TXIDs exceeds the thresholds, the appliance logs an event in the syslog file and sends an SNMP trap and e-mail notification, if you enable them. You can then configure rate limiting rules to limit incoming traffic or completely block connections from primary sources that send the invalid DNS responses.

Rate limiting is a token bucket system that accepts packets from a source based on the rate limit. You can configure the number of packets per minute that the Infoblox DNS server accepts from a specified source. You can also configure the number of packets for burst traffic, which is the maximum number of packets that the token bucket can accept. Once the bucket reaches the limit for burst traffic, it discards the packets and starts receiving new packets according to the rate limit.

The appliance monitors only UDP traffic from remote port 53 for the following reasons:

- The attacks that the appliance monitors do not happen over TCP.
- DNS responses are sent only from port 53. The appliance discards DNS responses that are sent from other ports.

To monitor invalid ports and invalid TXIDs on the Infoblox DNS server, follow these procedures:

1. Enable DNS network monitoring and DNS alert monitoring. For information, see [Enabling and Disabling DNS Alert Monitoring](#) on page 1025.
2. Configure the thresholds for DNS alert indicators. For information, see [Configuring DNS Alert Thresholds](#) on page 1026.
3. Enable SNMP traps and e-mail notifications. For information, see [Configuring SNMP](#) on page 1039.
4. Review the DNS alert status. For information, see [Viewing DNS Alert Indicator Status](#) on page 1026.
5. Identify the source of the attack by reviewing the DNS alert status, syslog file, and SNMP traps. For information on SNMP traps for DNS alerts, see [Threshold Crossing Traps](#) on page 1077.

To mitigate cache poisoning, you can limit incoming traffic or completely block connections from specific sources, as follows:

- Enable rate limiting on the DNS server. For information, see [Enabling and Disabling Rate Limiting from External Sources](#) on page 1027.
- Configure rate limit traffic rules from specific sources. For information, see [Configuring Rate Limiting Rules](#) on page 1028.

You can verify the rate limiting rules after you configure them. For information, see [Viewing Rate Limiting Rules](#) on page 1029.

Enabling and Disabling DNS Alert Monitoring

The appliance monitors only UDP traffic on port 53 for recursive queries, and then reports invalid DNS responses. DNS alert monitoring is disabled by default. For an HA pair, you must enable DNS alert monitoring on both the active and passive nodes.

To enable DNS network monitoring and DNS alert monitoring:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set monitor dns on
```

The appliance displays the following:

```
Turning on DNS Network Monitoring...
```

3. Enter the following command:

```
set monitor dns alert on
```

When you enable DNS alert monitoring and DNS network monitoring is disabled, the appliance automatically enables DNS network monitoring and displays the following:

```
DNS Network Monitoring is disabled. It must be enabled for alerting to function.
```

```
Enable DNS Monitoring now? (y or n):
```

You can also disable DNS network monitoring and DNS alert monitoring using the following commands:

```
set monitor dns off
```

```
set monitor dns alert off
```

Note: When you restart DNS network monitoring, you also reset the SNMP counters for DNS alerts.

You can then view the alert status to identify the primary source of invalid DNS responses. For information, see [Viewing DNS Alert Indicator Status](#) on page 1026.

Viewing DNS Alert Indicator Status

To view DNS alert indicator status:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
show monitor dns alert status
```

The appliance displays historical alert counts and up to five primary sources that generate invalid DNS responses, as shown in the following example:

```
Data last updated: Mon Oct 6 14:47:12 2008
DNS Alert   1m    5m    15m   60m   24h   Ever
=====
port         8     12    12    12    12    12
txid         8     12    12    12    12    12

There were 80 DNS responses seen in the last minute.
10% were to an invalid port.
10% had an invalid TXID.
```

```
Primary sources of invalid responses:
4.4.4.4 (unknown) sent 4
2.2.2.2 (unknown) sent 3
7.7.7.7 (unknown) sent 1
```

The appliance attempts to resolve the hostnames of the sources that sent invalid responses, if the DNS resolver is enabled. If the appliance cannot resolve a hostname, it displays “unknown” as the hostname of the invalid response.

Configuring DNS Alert Thresholds

You can configure thresholds for DNS alerts to control when the appliance tracks DNS attacks on UDP port 53 and issues SNMP traps and e-mail notifications.

Note: Ensure that you enable SNMP traps and e-mail notifications. For information, see [Configuring SNMP](#) on page 1039.

You can configure thresholds for both invalid ports and invalid TXIDs. The default thresholds for both invalid ports and TXIDs are 50%. When the number of invalid ports or invalid TXIDs exceeds the thresholds, the appliance logs the event and sends SNMP traps and notifications. You can configure the thresholds either as absolute packet counts or as percentages of the total traffic during a one minute time interval.

To configure DNS alert thresholds:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:


```
set monitor dns alert modify port | txid over threshold_value packets | percent
```

where

`port` | `txid` = Enter `port` to set the threshold for invalid ports, or enter `txid` to set the threshold for invalid TXIDs.

`threshold_value` = Enter the number of packets or percentage for the threshold.

`packets` | `percent` = Enter `packets` if you want to track the total packet count, or enter `percentage` if you want to track a percentage of the total traffic. For a percentage-based threshold, the appliance does not generate a threshold crossing event if the traffic level is less than 100 packets per minute.

For example, if you want the appliance to send a DNS alert when the percentage of DNS responses arriving on invalid ports from UDP port 53 exceeds 70% per minute, you can enter the following command:

```
set monitor dns alert modify port over 70 percent
```

If you want the appliance to send a DNS alert when the total number of packets with invalid TXIDs from UDP port 53 is over 100 packets per minute, you can enter the following command:

```
set monitor dns alert modify txid over 100 packets
```

When there is a DNS alert, the appliance logs an event in the syslog file and sends an SNMP trap and e-mail notification if enabled.

Viewing DNS Alert Thresholds

You can view the DNS alert thresholds. The appliance displays the current thresholds. If you have not configured new thresholds, the appliance displays the default thresholds, which are 50% for both invalid port and TXID.

To view the DNS alert thresholds:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
show monitor dns alert
```

The appliance displays the threshold information as shown in the following example:

```
DNS Network Monitoring is enabled.
```

```
Alerting is enabled.
```

```
DNS Alert      Threshold (per minute)
```

```
=====
```

```
port           over 70% of packets
```

```
txid           over 100 packets
```

Enabling and Disabling Rate Limiting from External Sources

You can mitigate cache poisoning on your DNS server by limiting the traffic or blocking connections from UDP port 53.

To enable rate limiting from sources:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set ip_rate_limit on
```

The appliance displays the following:

```
Enabling rate limiting will discard packets and may degrade performance.
```

```
Are you sure? (y or n):
```

Note: When you enable rate limiting, the appliance discards packets based on the configured rate limiting rules. This might affect the DNS performance when the appliance discards valid DNS responses.

3. Enter `y` to enable rate limiting.

When you enable rate limiting, the appliance applies the rate limiting rules that you configured. You might want to configure the rate limiting rules before enabling rate limiting. For information on how to configure rate limiting rules, see [Configuring Rate Limiting Rules](#) on page 1028.

You can also disable rate limiting by entering the following command:

```
set ip_rate_limit off
```

When you disable rate limiting, the appliance stops applying the rate limiting rules.

Configuring Rate Limiting Rules

You configure rate limiting rules to limit access or block connections from UDP port 53. The rules take effect when you enable rate limiting.

When adding rules, ensure that you do not include an IP address that matches the IP address of either the Grid Master or Grid member. Doing this could affect VPN connectivity. To configure rate limiting rules:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set ip_rate_limit add source all | ip_address [/mask] limit packets/m [burst
burst_packets]
```

where

`all | ip_address` = Enter `all` or `0.0.0.0` if you want to limit all traffic from all sources, or enter the IP address from which you want to limit the traffic.

`[/mask]` = Optionally, enter the netmask of the host from which you want to limit the traffic.

`packets` = Enter the number of packets per minute that you want to receive from the source.

`[burst burst_packets]` = Optionally, enter `burst` and the number of packets for burst traffic. This is the maximum number of packets accepted.

The following are sample commands and descriptions for rate limiting rules:

- To block all traffic from host 10.10.1.1, enter the following command:

```
set ip_rate_limit add source 10.10.1.1 limit 0
```
- To limit traffic to five packets per minute from host 10.10.1.2, enter the following command:

```
set ip_rate_limit add source 10.10.1.2 limit 5/m
```
- To limit the traffic to five packets per minute from host 10.10.2.1/24 with an allowance for burst traffic of 10 packets, enter the following command:

```
set ip_rate_limit add source 10.10.2.1/24 limit 5/m burst 10
```
- To limit the traffic to 5000 packets per minute from all sources, enter the following command:

```
set ip_rate_limit add source all limit 5000/m
```

Removing Rate Limiting Rules

You can remove the existing rate limiting rules that limit access or block connections from UDP port 53.

To remove all the existing rules:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:
 - To remove the rate limiting rule that limits traffic from all sources, enter:

```
set ip_rate_limit remove source all
```
 - or
 - To remove all of the rate limiting rules from all sources, enter:

```
set ip_rate_limit remove all
```

To remove one of the existing rules for an existing host:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set ip_rate_limit remove source ip-address[/mask]
```

Viewing Rate Limiting Rules

You can view the existing rate limiting rules that limit access or block connections from UDP port 53.

To view rate limiting rules:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
show ip_rate_limit
```

The appliance displays the rules, as shown in the following example:

IP rate limiting is enabled.

Source	Limit	Burst
10.10.1.1	0 packets/minute	0 packets
10.10.1.2	5 packets/minute	5 packets
10.10.2.1/24	5 packets/minute	10 packets
all	5000packets/minute	5000 packets



Chapter 36 DHCP Fingerprint Detection

This chapter explains the Infoblox DHCP fingerprint detection feature and how to configure it on the appliance. It also explains how to configure DHCP fingerprints for IPv4 and IPv6.

It contains the following sections:

- [*Infoblox DHCP Fingerprint Detection*](#) on page 1032
 - [*About DHCP Fingerprints*](#) on page 1033
 - [*Standard and Custom DHCP Fingerprints*](#) on page 1033
 - [*Administrative Permissions*](#) on page 1034
- [*Enabling and Disabling DHCP Fingerprint Detection*](#) on page 1034
- [*Configuring DHCP Fingerprints*](#) on page 1034
 - [*Adding New DHCP Fingerprints*](#) on page 1035
 - [*Modifying Custom DHCP Fingerprints*](#) on page 1035
 - [*Deleting Custom DHCP Fingerprints*](#) on page 1036
 - [*Viewing DHCP Fingerprint Information*](#) on page 1036

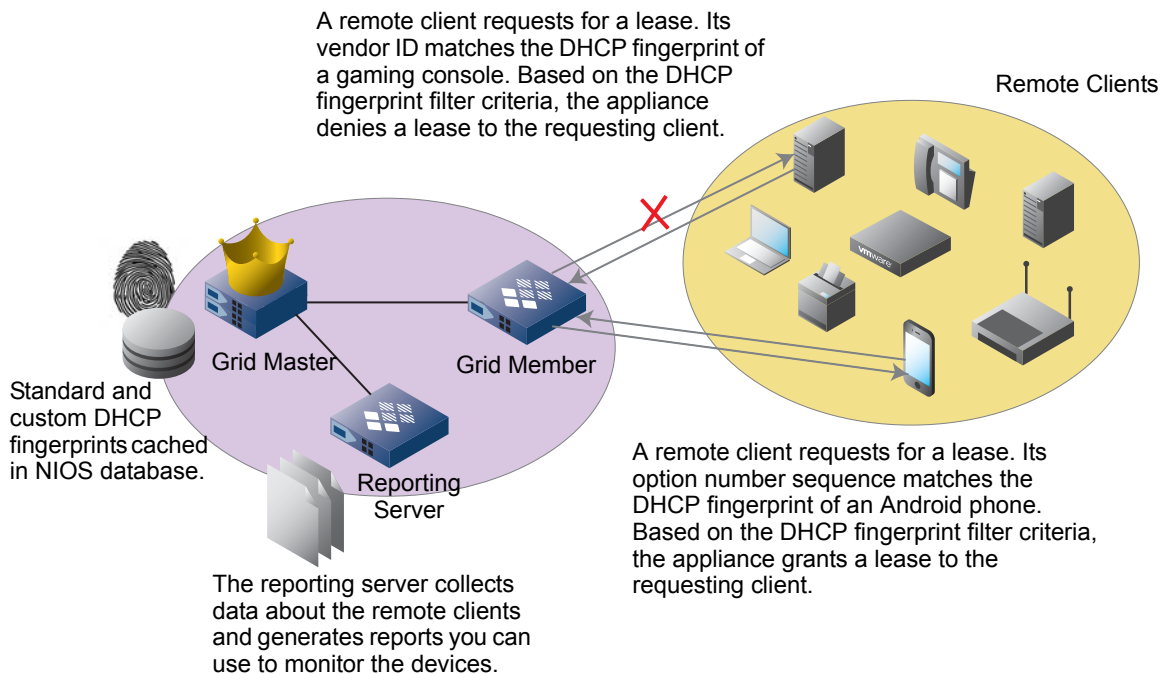
INFOBLOX DHCP FINGERPRINT DETECTION

The NIOS appliance utilizes DHCP fingerprint detection to identify IPv4 and IPv6 mobile devices such as laptop computers, tablets and smart phones, on your network. Due to the broadcast and pervasive nature of DHCP, using DHCP fingerprint detection is an efficient way to perform system identification and inventory. You can use DHCP fingerprint detection to track devices on your network, block those that are not allowed (such as gaming consoles and home routers), and plan for future growth by accessing trending information such as the number of Apple iPhones versus that of Android phones.

When a remote DHCP client sends a DHCP REQUEST message, it includes a set of DHCP options, such as option 55 and 60. Option 55 contains an option number sequence the appliance uses to interpret the list of DHCP options that the client requests. The appliance returns the values of these requested options if the information is available. Option 60 contains a value that indicates the device type of the requesting client. Information in option 55 or 60 is incorporated to form a unique identifier known as the DHCP fingerprint, which the appliance uses to identify the requesting client.

On an Infoblox appliance, DHCP fingerprint detection is enabled by default for all new installations. You can disable this feature at the Grid and member levels. For information, see [Enabling and Disabling DHCP Fingerprint Detection](#) on page 1034. As illustrated in [Figure 36.1](#), the appliance automatically matches option 55 and then option 60 in DHCP REQUEST messages against standard and custom DHCP fingerprints in the database. Once the appliance finds a match, it either grants or denies a lease to the requesting client based on the DHCP fingerprint filters that you apply to the DHCP range. For information about how to configure DHCP fingerprints, see [Configuring DHCP Fingerprints](#) on page 1034. For information about how to define and apply DHCP fingerprint filters, see [Defining DHCP Fingerprint Filters](#) on page 907 and [Applying Filters to DHCP Address Ranges](#) on page 907. To obtain trending information about the top OSs (operating systems) or vendor IDs for remote clients, Infoblox provides a few reports from which you can extract data. For information about reports, see [Infoblox Reporting Solution](#) on page 1116.

Figure 36.1 DHCP Fingerprint Detection



About DHCP Fingerprints

When a DHCP client sends a REQUEST message and includes DHCP option 55 (the parameter request list) and option 60 (the vendor identifier), it provides information about its OS and device type. The combination of the option sequence or vendor ID in option 55 or 60 is used to infer the OS and device type of the remote client. These parameters are then incorporated into a DHCP fingerprint that provides unique information about this client.

For example, the option number sequence for a Microsoft Windows XP system in option 55 can be one of the following:

```
1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43
1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252
1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252, 12
15, 3, 6, 44, 46, 47, 31, 33, 249, 43
15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252
28, 2, 3, 15, 6, 12, 44, 47
```

The option number sequence for an Apple iPhone can be one of the following:

```
1, 3, 6, 15, 119, 78, 79, 95, 252
1, 3, 6, 15, 119, 95, 252, 44, 46, 47
```

In addition, DHCP option 60 tracks vendor ID. This information can be very generic or quite specific. For example, the vendor ID `MSFT 5.0` for a Microsoft Windows XP system and a Windows Vista system can be the same. For certain Cisco VoIP devices, the vendor ID can be `Cisco Systems, Inc. IP Phone`, which is very generic; or it can be `Cisco Systems, Inc. IP Phone 7912`, which is more specific. Depending on how specific the option number sequence and the vendor ID are, this information can form a unique identifier, the DHCP fingerprint, for a remote client.

Note: If you have enabled firewall, and if the corresponding firewall rules or policies are set to modify options 55 and 60 of the remote DHCP client to mask the identity of the client, then NIOS fingerprinting will not be able to fingerprint the clients.

Standard and Custom DHCP Fingerprints

Standard DHCP fingerprints are automatically installed on the appliance when you first set it up or after you have completed an upgrade from previous NIOS releases to NIOS 6.7 and later. Note that new DHCP fingerprints are added to the appliance during a major NIOS upgrade. For more information about upgrades, see [About Upgrades](#) on page 406.

Note: When you upgrade to NIOS 6.7 and later, DHCP fingerprint detection is disabled during the upgrade. You must enable it if you want the appliance to use DHCP fingerprint detection. For information, see [Enabling and Disabling DHCP Fingerprint Detection](#) on page 1034.

You can configure custom DHCP fingerprints for devices whose DHCP fingerprints are not captured in the standard DHCP fingerprints. For information about how to add custom DHCP fingerprints, see [Adding New DHCP Fingerprints](#) on page 1035.

Both standard and custom DHCP fingerprints are cached in memory for matching purposes. Depending on the information provided in a DHCP fingerprint, the appliance first matches the option number sequence sent in the DHCP REQUEST message. If option 55 is not included in the request or if there is no match from the cached DHCP fingerprints, the appliance then tries to match the vendor ID in option 60. When there is an option number sequence match, the appliance displays the name of the DHCP fingerprint in Grid Manager. If there is no option number sequence match but there is a vendor ID match, the appliance displays the vendor ID. For information about how to view fingerprint information, see [Viewing DHCP Fingerprint Information](#) on page 1036.

You can also create IPv4 and IPv6 DHCP fingerprint filters and then apply them as class filters to specific IPv4 and IPv6 DHCP ranges and range templates. For information about how to configure and use DHCP fingerprint filters, see [About DHCP Fingerprint Filters](#) on page 906.

Administrative Permissions

DHCP fingerprint detection is enabled by default for new installations. For upgrades, you must enable this feature after the upgrade is completed. For information, see [Enabling and Disabling DHCP Fingerprint Detection](#). No special licenses are required for this feature.

Superusers can add, modify, and delete DHCP fingerprints and DHCP fingerprint filters. Limited-access users with Read/Write permission to DHCP fingerprints can add, modify, and delete DHCP fingerprints while those who have Read-only permission can only view information in the **Data Management** tab → **DHCP** tab → **Fingerprints** tab. For information about administrative permissions, see [About Administrative Permissions](#) on page 160.

ENABLING AND DISABLING DHCP FINGERPRINT DETECTION

Grid DHCP fingerprint detection is enabled by default for new installations, and no special licenses are required. You can disable this or override the Grid setting at a member level. Note that when you enable DHCP fingerprint detection, there will be a slight impact on DHCP performance.

When you enable DHCP fingerprint on an HA pair, both peers in a failover association maintain the same DHCP fingerprinting state (enabled or disabled) even when one of the peers fails or becomes operational again. Note that both peers must be in the same Grid for the fingerprinting state to stay the same. For information about DHCP failover, see [About DHCP Failover](#) on page 884.

To enable and disable Grid DHCP fingerprinting:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
Member: From the **Data Management** tab, select the **DHCP** tab → **Members** tab → **Members** → *member* check box, and then click the Edit icon.
2. In the *Grid DHCP Properties* or *Member DHCP Properties* editor, select the **Fingerprinting** tab.
3. Complete the following
 - **Enable Fingerprint Detection:** Deselect this check box to disable the feature. You can enable DHCP fingerprint detection again by selecting the check box. Click **Override** to change the configuration for a member, or click **Inherit** to inherit the Grid setting.
4. Save the configuration and click **Restart** at the top of the screen.

CONFIGURING DHCP FINGERPRINTS

The appliance installs standard DHCP fingerprints when you first install or upgrade to NIOS 6.7 and later. You cannot modify standard DHCP fingerprints nor delete them, but you can disable them. When you disable a DHCP fingerprint, the appliance disables the associated option number sequence and vendor ID, and it cannot match a remote device against the disabled DHCP fingerprint.

When you add a new DHCP fingerprint, the appliance marks it as a custom DHCP fingerprint. For information about adding custom DHCP fingerprints, see [Adding New DHCP Fingerprints](#) on page 1035. You can modify information about custom DHCP fingerprints, and you can delete them. For information, see [Modifying Custom DHCP Fingerprints](#) and [Deleting Custom DHCP Fingerprints](#) on page 1036. When you delete a custom DHCP fingerprint, the appliance moves it to the Recycle Bin, if enabled. You can later restore it from the Recycle Bin if needed.

Activities, such as additions, modifications, and deletions of DHCP fingerprints, are recorded in the audit log. For information about how to use the audit log, see [Using the Audit Log](#) on page 1018.

Note: The appliance periodically updates the cached DHCP fingerprints. When you add, modify, or delete a DHCP fingerprint, you do not need to restart services but it may take up to two minutes before the appliance updates the DHCP fingerprint.

Adding New DHCP Fingerprints

To add a custom DHCP fingerprint:

1. From the **Data Management** tab, select the **DHCP** tab -> **Fingerprints** tab, and then click the Add icon.
2. In the *Add DHCP Fingerprint* wizard, complete the following:
 - **Name:** Enter the name of the custom DHCP fingerprint. The name must be unique, and it cannot contain any UTF-8 characters.
 - **Device Class:** From the drop-down list, select the device category to which this new fingerprint belongs. You can also enter a new device class here. When you enter a device class that already exists, the appliance matches the entry and uses the class from the current list. Device class is used for filtering purposes. For information about DHCP fingerprint filters, see [Defining DHCP Fingerprint Filters](#) on page 907.
 - **Protocol:** From the drop-down list, select the protocol used for custom DHCP fingerprint.
 - **Option Number Sequence:** Click the Add icon in the table. The appliance adds a row to the table. Click the row and enter the DHCP option number you want the appliance to validate. Valid values are from 0 to 255. When you enter more than one option, you must use commas (without spaces) to separate the numbers. For example, you can enter 1, 15, 3, 6, 44, 46, 47, 31, 33 for a Windows XP system.
You can also select an option sequence and click the Delete icon to delete it. Note that if you enter an option sequence that already exists in a standard DHCP fingerprint, you must disable that standard fingerprint before you can add the option sequence to the new DHCP fingerprint.
 - **Vendor Identifier:** Click the Add icon in the table. The appliance adds a row to the table. Click the row and enter a vendor ID for this fingerprint. You can add more than one vendor ID. You can also select a vendor ID and click the Delete icon to delete it.
 - **Comment:** Enter additional information about the custom DHCP fingerprint.
 - **Disabled:** Select this if you want to save the configuration for the DHCP fingerprint but do not want to activate it yet. You can clear this check box when you are ready to use this DHCP fingerprint.
3. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

To schedule this task, click the Schedule icon at the top of the wizard. In the Schedule Change panel, click **Later**, and then specify a date, time, and time zone.

Modifying Custom DHCP Fingerprints

You can modify custom DHCP fingerprints, but not standard ones. Note that the appliance periodically updates the cached DHCP fingerprints. When you modify a DHCP fingerprint, you do not need to restart services but it may take up to two minutes before the appliance updates the DHCP fingerprint.

To modify a custom DHCP fingerprint:

1. From the **Data Management** tab, select the **DHCP** tab -> **Fingerprints** tab -> *custom_fingerprint* check box, and then click the Edit icon.
2. The *DHCP Fingerprint* editor provides the following tabs from which you can modify information:
 - **General:** Modify general information, such as the name and device class, as described in [Adding New DHCP Fingerprints](#) on page 1035. Note that when you change the name of a DHCP fingerprint, the old name no longer exists, and you cannot use it for searching or filtering purposes. You may not be able to modify all fields in a standard DHCP fingerprint.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the DHCP fingerprint. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332. Note that you cannot modify extensible attributes for standard fingerprints.
3. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the Schedule Change panel, click **Later**, and then specify a date, time, and time zone.

Deleting Custom DHCP Fingerprints

When you delete a custom DHCP fingerprint, the appliance moves it to the Recycle Bin, if enabled. You can later restore the DHCP fingerprint if needed. Note that you cannot delete standard DHCP fingerprints.

To delete a custom DHCP fingerprint:

1. From the **Data Management** tab, select the **DHCP** tab -> **Fingerprints** tab -> *custom_fingerprint* check box, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes** to delete the DHCP fingerprint.

To schedule this task, click the Delete icon -> **Schedule Delete**. In the *Schedule Deletion* dialog box, click **Delete Later**, and then specify a date, time, and time zone.

Viewing DHCP Fingerprint Information

The following are a few ways you can view DHCP fingerprint information:

- In a Grid with a reporting server, you can view reports that contain information about the top OSs and device types of the leasing clients in your network. For more information, see [DHCP Fingerprint Reports](#) on page 1149.
- The appliance provides a few predefined smart folders from which you can view lease information about specific device groups, such as gaming consoles and Microsoft Windows devices. For more information, see [Predefined Smart Folders](#) on page 142.
- When the appliance finds a DHCP fingerprint match for a client, Grid Manager displays either the fingerprint name or the vendor ID in the following panels of Grid Manager: IP List, Current Lease, Lease History, and DHCP Range panels. You can see this information in the **Fingerprints** column in these panels.
- The appliance records all DHCP fingerprint related activities in the audit log. For more information, see [Using the Audit Log](#) on page 1018.



Chapter 37 Monitoring with SNMP


This chapter describes how you can use SNMP (Simple Network Management Protocol) to monitor NIOS appliances in your network. It contains the following sections:

- [*Understanding SNMP*](#) on page 1038
 - [*About SNMPv1 and SNMPv2*](#) on page 1039
 - [*About User-Based Security Model in SNMPv3*](#) on page 1039
- [*Configuring SNMP*](#) on page 1039
 - [*Configuring SNMPv3 Users*](#) on page 1040
 - [*Modifying SNMPv3 Users*](#) on page 1041
 - [*Deleting SNMPv3 Users*](#) on page 1041
 - [*Accepting Queries*](#) on page 1041
 - [*Adding Trap Receivers*](#) on page 1042
 - [*Defining Thresholds for Traps*](#) on page 1043
 - [*Setting SNMP and Email Notifications*](#) on page 1044
 - [*Setting SNMP System Information*](#) on page 1043
 - [*Testing the SNMP Configuration*](#) on page 1047
- [*SNMP MIB Hierarchy*](#) on page 1048
 - [*MIB Objects*](#) on page 1049
 - [*System Object IDs*](#) on page 1049
- [*Infoblox MIBs*](#) on page 1051
 - [*Loading the Infoblox MIBs*](#) on page 1051
 - [*ibTrap MIB*](#) on page 1053
 - [*ibPlatformOne MIB*](#) on page 1086
 - [*ibDHCPOne MIB*](#) on page 1100
 - [*ibDNSOne MIB*](#) on page 1107
 - [*IB-DNSSERV-MIB*](#) on page 1111
 - [*IB-DNSHITRATIO-MIB*](#) on page 1111
 - [*IB-DNSQUERYRATE-MIB*](#) on page 1111
 - [*IB-DHCPSESV-MIB*](#) on page 1111

UNDERSTANDING SNMP

You can use SNMP (Simple Network Management Protocol) to manage network devices and monitor their processes. An SNMP-managed device, such as a NIOS appliance, has an SNMP agent that collects data and stores them as objects in MIBs (Management Information Bases). The SNMP agent can also send traps (or notifications) to alert you when certain events occur within the appliance or on the network. You can view data in the SNMP MIBs and receive SNMP traps on a management system running an SNMP management application, such as HP OpenView, IBM Tivoli NetView, or any of the freely available or commercial SNMP management applications on the Internet.

Figure 37.1 *SNMP Overview*



The NIOS appliance supports SNMPv1, SNMPv2, and SNMPv3. It also adheres to the following RFCs:

- *RFC 3411, An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*
- *RFC 3412, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
- *RFC 3413, Simple Network Management Protocol (SNMP) Applications*
- *RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMP)*
- *RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- *RFC 3418, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- *RFC 1155, Structure and identification of Management information for TCP/IP-based internets*
- *RFC 1213, Management Information Base for Network Management of TCP/IP-based internets:MIB-II*
- *RFC 2578, Structure of Management Information Version 2 (SMIv2)*

About SNMPv1 and SNMPv2

SNMPv1 is the initial implementation of SNMP. It operates over protocols such as UDP (User Datagram Protocol) and IP (Internet Protocol). SNMPv2 includes improvements in performance and security. It adds new protocol operations such as GetBulk and Inform, which allow the management system to request larger blocks of data from the agent. Both SNMPv1 and SNMPv2 use common strings that are sent in clear text to authenticate clients.

The NIOS appliance supports SNMPv1 and SNMPv2 in which the SNMPv2 agent acts as a proxy agent for the SNMPv1 management systems. When an SNMPv1 management system sends a query to the appliance, the SNMPv2 proxy agent forwards the request to the SNMPv1 agent. The proxy agent maps the SNMPv1 trap messages to the SNMPv2 trap messages, and then forward the messages to the management system.

You can enable the appliance to receive queries from SNMPv1 and SNMPv2 management systems. You can also add SNMPv1 and SNMPv2 management systems to receive traps from the appliance. For information about how to configure SNMPv1 and SNMPv2 on the appliance, see [Configuring SNMP](#) on page 1039.

About User-Based Security Model in SNMPv3

SNMPv3 adds security and remote configuration enhancements to SNMPv1 and SNMPv2. The NIOS appliance supports the USM (User-based Security Model) in SNMPv3 for the authentication, encryption, and decryption of SNMP data. SNMPv3 uses the same MIB objects as those supported in SNMPv1 and SNMPv2.

SNMPv3 provides the following security measures:

- **Data integrity:** Ensure that SNMP data is not maliciously modified by unauthorized entities during its transmission through the network. This protects against unauthorized management operations, such as falsifying the value of a MIB object.
- **Authentication:** Verify the identities of the origin of the SNMP data to protect against masquerade threats that may temper the identity of users who have the appropriate authorization to send and receive SNMP data.
- **Confidentiality:** Ensure that unauthorized users cannot eavesdrop on any data exchanges between SNMP agents and management systems, depending on local policies of the systems.
- **Timeliness:** Ensure that the SNMP data is received in a timely manner to prevent malicious reordering of data by unauthorized entities.

To enable SNMPv3 on the NIOS appliance to provide user-based security, you must first configure SNMPv3 users on the appliance to enable access by SNMP management systems. The appliance supports HMAC-MD5-96 and HMAC-SHA-96 hash functions as the authentication protocols, and DES (Data Encryption Standard) and AES (Advanced Encryptions Standard) as the encryption methods for SNMPv3 users. For information, see [Configuring SNMP](#) on page 1039.

CONFIGURING SNMP

Note: SNMP operation is not supported across the NIOS appliance's LAN2 interface.

You can configure the appliance to receive SNMP queries from specific management systems and send SNMP traps to specific trap receivers. The appliance supports SNMPv1, SNMPv2, and SNMPv3. You can set up either SNMPv1/SNMPv2 or SNMPv3, or all of them for the Grid. You can also override the Grid settings at a member level.

To configure SNMPv1 and SNMPv2 on the appliance, do the following:

- Enable the NIOS appliance to accept queries, as described in [Accepting Queries](#) on page 1041.
- Specify the management systems to which the appliance sends traps, as described in [Adding Trap Receivers](#) on page 1042.
- Specify system information using managed objects in MIB-II, the standard MIB defined in *RFC 1213*. For information, see [Setting SNMP System Information](#) on page 1043.

To configure SNMPv3 on the appliance, do the following:

- Add an SNMPv3 user and set up authentication and privacy protocols. For information, see [Configuring SNMPv3 Users](#) on page 1040. After you set up an SNMPv3 user, you can modify and delete it. For information, see [Modifying SNMPv3 Users](#) on page 1041 and [Deleting SNMPv3 Users](#) on page 1041.
- Enable the NIOS appliance to accept queries, as described in [Accepting Queries](#) on page 1041.
- Specify the management systems to which the appliance sends traps, as described in see [Adding Trap Receivers](#).
- Specify system information using managed objects in MIB-II, the standard MIB defined in *RFC 1213*. For information, see [Setting SNMP System Information](#) on page 1043.

Configuring SNMPv3 Users

To enable SNMPv3, you must first configure SNMPv3 users on the appliance. For information about SNMPv3, see [About User-Based Security Model in SNMPv3](#) on page 1039.

To configure an SNMPv3 user:

1. From the **Administration** tab, select the **SNMPv3 Users** tab, and then click the Add icon.
2. In the *Add SNMPv3 User* wizard, complete the following:
 - **Name:** Enter a user name for the SNMPv3 management system.
 - **Authentication Protocol:** Select one of the following:
 - **MD5:** Select this to use the HMAC-MD5-96 authentication protocol to authenticate the SNMPv3 user. This protocol uses the MD5 (Message-Digest algorithm 5) hash function in HMAC (Hash-based Message Authentication Code) and truncates the output to 96 bits. The output is included as part of the SNMP message sent to the receiver. For detailed information about the protocol, refer to *RFC1321, The MD5 Message-Digest Algorithm*.
 - **SHA:** Select this to use the HMAC-SHA-96 authentication protocol to authenticate the SNMPv3 user. This protocol uses the SHA (Secure Hash Algorithm) hash function and truncates the output to 96 bits. The output is included as part of the SNMP message sent to the receiver.
 - **None:** Select this to decline using any authentication protocol for this SNMPv3 user. When you select this option, you are not required to enter a password.
 - **Password:** Enter a password for the selected authentication protocol.
 - **Confirm Password:** Enter the same password.
 - **Privacy Protocol:** Select one of the following:
 - **DES:** Select this to use DES for data encryption. DES is a block cipher that employs a 56-bit key size and 64-bit block size in the encryption.
 - **AES:** Select this to use AES for data encryption. AES is a symmetric-key encryption standard that comprises three block ciphers, AES-128, AES-192, and AES-256. Each of these ciphers has a 128-bit block size and a key size of 128, 192, and 256 bits, respectively.
 - **None:** Select this to decline using any privacy protocol for this SNMPv3 user. When you select this option, you are not required to enter a password.
 - **Password:** Enter a password for the privacy protocol.
 - **Confirm Password:** Enter the same password.
 - **Comment:** Enter useful information about the SNMP user, such as location or department.
 - **Disable:** Select this check box to retain an inactive profile for this SNMP user in the configuration. You can clear this check box to activate the profile.

Note: If an SNMPv3 user is configured to send SNMP queries, you cannot delete the user.

3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Modifying SNMPv3 Users

1. From the **Administration** tab, select the **SNMPv3 Users** tab -> *snmpv3user*, and then click the Edit icon.
2. The *SNMPv3 User* editor provides the following tabs from which you can edit data:
 - **General:** Modify the data as described in [Configuring SNMPv3 Users](#) on page 1040.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the SNMPv3 user account. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
3. Save the configuration.

Deleting SNMPv3 Users

When you delete an SNMPv3 user that is configured to send queries or receive traps, a warning message states that the SNMPv3 is associated with the corresponding function. You can then decide whether you want to delete the user or not.

To delete an SNMPv3 user:

1. From the **Administration** tab, select the **SNMPv3 Users** tab -> *snmpv3user*, and then click the Delete icon.
2. In the *Delete confirmation* dialog box, click **Yes**.

Note: You cannot schedule the deletion of an SNMPv3 user.

Accepting Queries

You can allow specific management systems to send SNMP queries to a NIOS appliance. For SNMPv1 and SNMPv2, you must specify a community string. The appliance accepts queries only from management systems that provide the correct community string. You can also specify SNMPv3 users to send queries. For information about configuring SNMPv3 users, see [Configuring SNMPv3 Users](#) on page 1040.

To configure an appliance to accept SNMP queries:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, select the **SNMP** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP** section.
 - **Enable SNMPv1/SNMPv2 Queries:** Select this to accept SNMPv1 and SNMPv2 queries from management systems.
 - **Community String:** Enter a text string that the management system must send together with its queries to the appliance. A community string is similar to a password in that the appliance accepts queries only from management systems that send the correct community string. Note that this community string must match exactly what you enter in the management system.
 - **Engine ID:** Displays the engine ID of the appliance that manages the SNMP agent. The management system needs this ID to send traps to the appliance. If the appliance is an HA pair, this field displays the engine IDs for both the active and passive nodes.
 - **Enable SNMPv3 Queries:** Select this to enable queries from SNMPv3 management systems. Click the Add icon to add SNMPv3 users that you have configured on the appliance. In the *SNMPv3 User Selector* dialog box, click the SNMPv3 user you want to add. The appliance displays the selected SNMPv3 users in the table. You can add comments in the table. You can also select an SNMPv3 user and click the Delete icon to remove it from the table. Note that a disabled SNMPv3 user cannot send queries to the appliance.
4. Save the configuration.

Adding Trap Receivers

You can enable the NIOS appliance to send traps to specific management systems using either SNMPv1/SNMPv2 or SNMPv3, or all versions of SNMP. You can then add management systems that are allowed to receive traps from the appliance. Note that you cannot enable both SNMPv1/SNMPv2 and SNMPv3 on the same trap receiver. The appliance sends traps when certain events occur. You can enable SNMP traps and add trap receivers to the Grid. You can also override the Grid settings at the member level.

To enable the appliance to send traps and to add trap receivers, do the following:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, select the **SNMP** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP** tab:
 - **Enable SNMPv1/SNMPv2 Traps:** Select this to enable the appliance to send traps to specified management systems.
 - **Community String:** Enter a text string that the NIOS appliance sends to the management system together with its traps. Note that this community string must match exactly what you enter in the management system.
 - **Enable SNMPv3 Traps:** Select this to enable the appliance to send traps to specified SNMPv3 users.
4. Click the Add icon and select one of the following from the drop-down menu to add an SNMP trap receiver:
 - **SNMPv1/SNMPv2:** Select this to add an SNMPv1 or SNMPv2 management system as a trap receiver. Grid Manager adds a row to the table. In the **Address** field, enter the IP address of the SNMP management system to which you want the SNMP agent on the appliance to send traps. You can enter more than one trap receiver. To remove a trap receiver from the list, select the address, and then click the Delete icon.
 - **SNMPv3:** Select this to add an SNMPv3 management system as a trap receiver. Grid Manager displays the *SNMPv3 User Selector* dialog box. Click the name of the SNMPv3 user in the dialog box. Grid Manager adds the user to the table. In the **Address** field, enter the IP address of the SNMP management system to which you want the SNMP agent on the appliance to send traps. You can add more than one trap receiver. To remove a trap receiver from the list, select the address, and then click the Delete icon.

Trap receiver IP addresses may be in IPv4 or IPv6 format.

In the Trap Receiver table, Grid Manager displays the following information about the trap receivers:

- **Address:** The IPv4 or IPv6 address of the trap receiver. Note that when an SNMPv3 user is disabled, SNMPv1/SNMPv2 traps are disabled. You can modify the IP address of the trap receiver even when the following are disabled: SNMPv3 users, SNMPv1/SNMPv2 traps, and SNMPv3 traps.
 - **SNMPv3 User:** The user name of the SNMPv3 trap receiver. This is for SNMPv3 only.
 - **Comment:** Information you entered about the management system.
5. Save the configuration.

Setting SNMP System Information

You can enter values for certain managed objects in MIB-II, the standard MIB defined in *RFC 1213*. Management systems that are allowed to send queries to the appliance can query these values. You can enter these values for the Grid and specific members. You can also override the Grid values at a member level.

To enter system information:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, select the **SNMP** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP** tab. For an HA member, click **Override Node 2 settings** to enter information for node 2 of the HA pair.
 - **sysContact:** Enter the name of the contact person for the appliance.
 - **sysLocation:** Enter the physical location of the appliance.
 - **sysName:** Enter the fully qualified domain name of the appliance.
 - **sysDescr:** Enter useful information about the appliance, such as the software version it is running.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Defining Thresholds for Traps

Threshold events for appliance performance are configurable. For each event, you can set a value that triggers the appliance to send a trap and another value at which the appliance sends a CLEAR trap. The appliance sends a CLEAR trap the first time the event value reaches the reset value after it reached the trigger value.

To define the threshold values:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, click **Toggle Advanced Mode**, and then select the **SNMP Threshold** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP Threshold** tab. Each of the following event types have default Trigger and Reset values. You can change the values for any of them. You can set SNMP thresholds above or below which the appliance sends SNMP traps and email notifications, if configured to do so. When any allocated usage exceeds the Trigger value, the appliance sends an SNMP trap and email notification to the designated destination, and the status icon for that usage turns red. When usage drops to the Reset value, the status color goes back to normal and turns green.
 - **Network Capacity:** When the Grid is part of a Master Grid, this is the percentage of the Master Grid's network capacity that is used by the Grid's networks. The default Trigger value is 85% and default Reset value is 75%.
 - **Database Objects:** The percentage of database capacity that is currently in use. The default Trigger value is 80%, and the default Reset value is 70%.
 - **Disk:** The percentage of the primary hard disk that is currently in use. The default Trigger value is 85%, and the default Reset value is 70%.
 - **Memory:** The percentage of the system memory that is currently in use. The default Trigger value is 90%, and the default Reset value is 80%.
 - **Swap Usage:** The percentage of the swap area that currently is in use. The default Trigger value is 20%, and the default Reset value is 10%.
 - **Root File System:** The percentage of the root file system ("/") that is currently in use. The default Trigger value is 85%, and the default Reset value is 70%.

- **CPU Usage:** The percentage of the CPU that is currently in use. The default Trigger value is 81%, and the default Reset value is 70%. Note that these default values are set to disable the CPU usage trap. You can enable this trap and configure the trigger and reset values using the CLI command `set thresholdtrap`.
- **Reporting:** The number of reports created on the system that can trigger an SNMP trap. The default Trigger value is 85, and the default Reset value is 70. Note that the maximum number of reports supported per Grid is 300. This field is displayed only when you have configured a reporting server.
- **Reporting Volume:** The percentage of data transmissions to the reporting server. The default Trigger value is 80%, and the default Reset value is 71%. This field is displayed only when you have configured a reporting server.
- **File Distribution Usage:** The percentage of the file distribution storage capacity that is currently in use on the selected member. The default Trigger value is 90%, and the default Reset value is 70%.

If you have installed Threat Protection licenses on the appliance and are using the Infoblox Advanced DNS Protection feature, Grid Manager displays the following for **Trigger events per second** and **Reset events per second**:

- **Alert Rate:** The number of SNMP traps sent per second when the appliance sends alerts while passing packets based on threat protection rule configuration. The default Trigger value is 1 and the default Reset value is 0.
- **Drop Rate:** The number of SMMP traps sent per second when the appliance drops packets based on the threat protection rule configuration. The default Trigger value is 1 and the default Reset value is 0.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Setting SNMP and Email Notifications

You can specify the event types that trigger trap and email notifications.

To set SNMP trap and email notifications:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** → **Edit** from the Toolbar.
Member: From the **Grid** tab, select the **Grid Manager** → **Members** tab → *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, click **Toggle Advanced Mode**, and then select the **Notification** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following:
 - **Enable All SNMP Notifications:** Select this check box if you want the appliance to send SNMP notifications (traps) for all events to the configured trap receivers. This is selected by default. To send SNMP notifications for specific events to the configured trap receiver, select the check box for respective event type.
For information on configuring trap receivers, see [Adding Trap Receivers](#) on page 1042.
 - **Enable All Email Notifications:** Select this check box if you want the appliance to send email notifications (traps) for all events to the configured email recipients. This is deselected by default. To send email notifications for specific events to the configured email recipients, select the check box for each respective event type. For more information, see [Selecting Email Notification Types](#) on page 1044.
For information on enabling email notifications and specifying recipients, see [Notifying Administrators](#) on page 191.
 - Alternatively, you can select specific event types from the table, and specify whether you want the appliance to send SNMP Notifications and Email notifications for each type of event.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Selecting Email Notification Types

You can configure email notifications in the NIOS appliance to receive alert messages on your email when a hardware or software fails in the NIOS appliance. To receive notifications when a specific hardware or software fails, select the corresponding notification types, as follows:

Table 37.1 Hardware Failures

Hardware Failure	Description
Fan	Indicates the status of the system fan. For more information, see Equipment Failure Traps on page 1070.
Disk	Indicates the status of primary disk. For more information, see Equipment Failure Traps on page 1070.
Memory	Indicates the status of the system memory. For more information, see Threshold Crossing Traps on page 1077.
CPU	Indicates the status of CPU usage. For more information, see Threshold Crossing Traps on page 1077.
MGM	Indicates the status of MGM port. For more information, see Object State Change Traps on page 1083.
HSM DNSSEC	Indicates the status of HSM operation. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
Login	Indicates that the login details are incorrect. For more information, see Processing and Software Failure Traps on page 1071.
PowerSupply	Indicates the status of power supply. For more information, see Equipment Failure Traps on page 1070.
FTP	Indicates the status of FTP service. For more information, see Processing and Software Failure Traps on page 1071.
TFTP	Indicates the status of TFTP service. For more information, see Processing and Software Failure Traps on page 1071.
MSServer	Indicates the status of Microsoft Server. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
Database	Indicates the status of database. For more information, see Processing and Software Failure Traps on page 1071.
RootFS	Indicates the status of root file system. For more information, see Threshold Crossing Traps on page 1077.
RAID	Indicates the status of RAID array. For more information, see Equipment Failure Traps on page 1070.
ENAT	Indicates the status of Ethernet port. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.

Table 37.2 Software Failures

Software Failure	Description
bloxTools	Indicates the status of bloxTools. For more information, see Object State Change Traps on page 1083.
Backup	Indicates the status of backup operation. For more information, see Processing and Software Failure Traps on page 1071.
Clear	Indicates that the SNMP trap is cleared. When you select the check box, the CLEAR trap is sent for the following software failures: LDAP servers, OSCP responders, LCD, Serial Console, OSPF, OSPF6, BGP, HSM, Controld, SSH, HTTP, Cluster, Login, and Duplicate IP. For file distribution, the trap is sent when the service is restored. If you deselect the check box, the CLEAR trap is not sent when any of the mentioned software fails. For more information, see ibProbableCause Values (OID 3.1.1.1.2.4.0) on page 1059.
SNMP	Indicates the status of SNMP server. For more information, see Processing and Software Failure Traps on page 1071.
LCD	Indicates the status of LCD process. For more information, see Processing and Software Failure Traps on page 1071.
SSH	Indicates the status of sshd process. For more information, see Processing and Software Failure Traps on page 1071.
SerialConsole	Indicates that the serial console login has failed or the admin failed to login to the serial console. For more information, see Processing and Software Failure Traps on page 1071.
Network	Indicates the status of the LAN port. For more information, see Threshold Crossing Traps on page 1077.
Cluster	Indicates the status of NIOS clusterd process. For more information, see Processing and Software Failure Traps on page 1071.
Controld	Indicates that the NIOS controld process has failed. For more information, see Processing and Software Failure Traps on page 1071.
OSPF	Indicates that the ospfd process has failed. For more information, see Processing and Software Failure Traps on page 1071.
IFMAP	Indicates the status of IF-MAP service. For more information, see ibProbableCause Values (OID 3.1.1.1.2.4.0) on page 1059.
BGP	Indicates that the BGP software has failed. For more information, see ibProbableCause Values (OID 3.1.1.1.2.4.0) on page 1059.
CaptivePortal	Indicates that the Captive Portal service has failed. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
DuplicateIP	Indicates that there are duplicate IP addresses. For more information, see ibProbableCause Values (OID 3.1.1.1.2.4.0) on page 1059.
License	Indicates that the license has been revoked. For more information, see Revoked License Trap on page 1086.
System NIOS	Indicates the status of the NIOS system. For more information, see Process Started and Stopped Traps on page 1085.

Software Failure	Description
Syslog NIOS	Indicates that the syslog process has stopped. For more information, see Processing and Software Failure Traps on page 1071.
DiscoveryConflict DHCP status	Indicates there is a conflict between the DHCP address and the existing IP address. For more information, see Processing and Software Failure Traps on page 1071.
Reporting Volume Database	Indicates the status of reporting database. For more information, see Threshold Crossing Traps on page 1077.
HTTP	Indicates the status of HTTP service. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
NTP	Indicates the status of NTP service. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
DNS	Indicates the status of DNS service. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
DHCP	Indicates the status of DHCP service. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
HA	Indicates the status of HA port link. For more information, see ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068.
File Distribution Usage	Indicates that the HTTP file distribution process has failed. For more information, see Processing and Software Failure Traps on page 1071.
OCSP Responders	Indicates the status of OCSP responders.
Disconnected Grid	Indicates that a Grid has been disconnected from the Master Grid. For more information, see Object State Change Traps on page 1083.
LB Device	Indicates whether the LB device is in sync or not.

Testing the SNMP Configuration

After you configure SNMP on the appliance, you can do the following to test your SNMP configuration:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid_member* check box, and then select **Test SNMP** from the Toolbar.

The appliance sends a “test trap” string to the trap receiver and displays a confirmation message at the top of the screen if your SNMP configuration is properly set up. If your SNMP configuration is not complete or if it is invalid, the appliance displays an error message. You can check your configuration and try again.

The following is a sample test trap that the trap receiver can get:

```
2011-04-04 17:37:14 10.32.2.80 [UDP: [10.32.2.80]:49244->[10.32.2.80]]:
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::snmpTrapOID
SNMPv2-MIB::sysName.0 = STRING: 'Test trap'
```

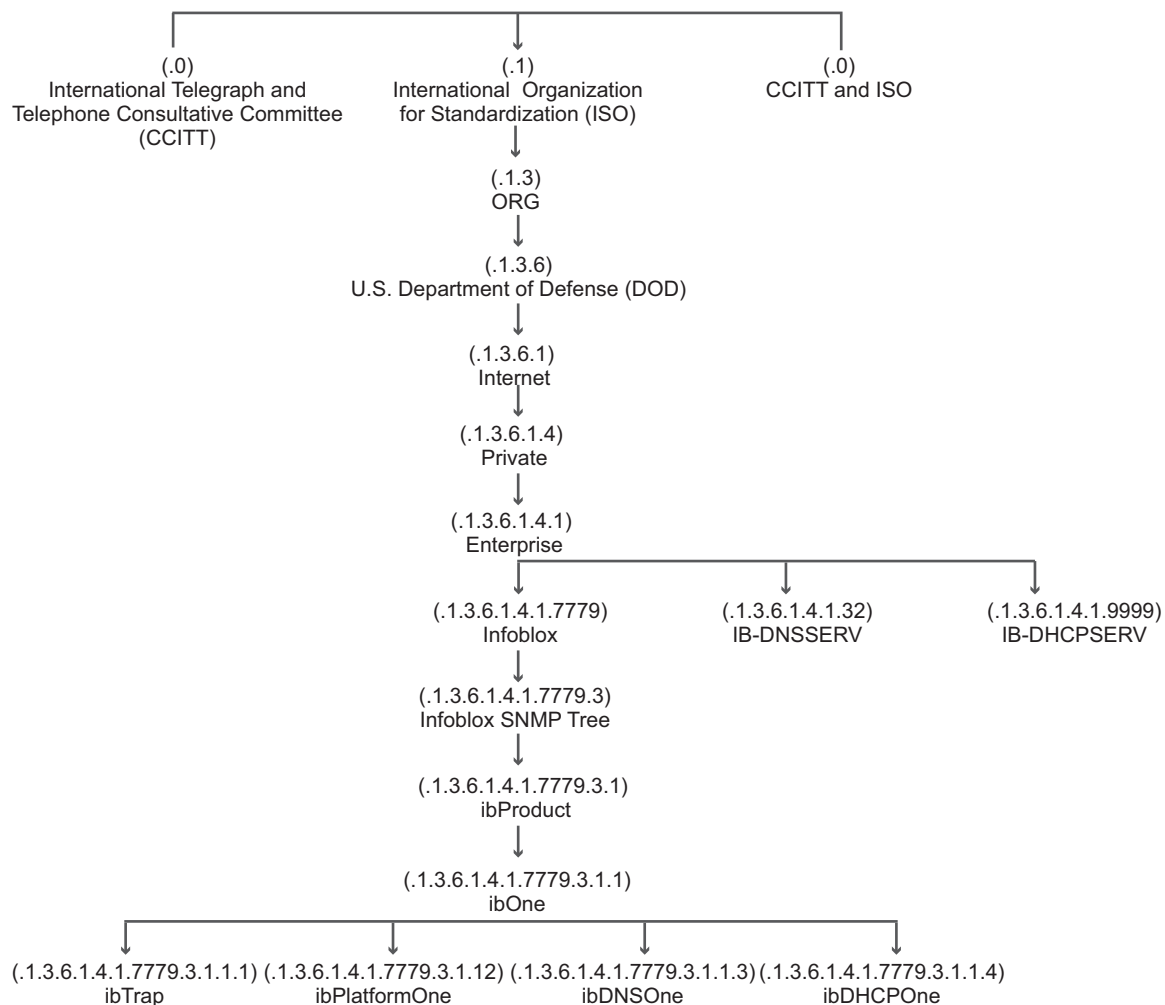
SNMP MIB Hierarchy

In addition to implementing its own enterprise MIBs, Infoblox supports the standard MIBs defined in *RFC-1213, Management Information Base for Network Management of TCP/IP-based internets: MIB-II*.

The Infoblox MIBs are part of a universal hierarchical structure, usually referred to as the MIB tree. The MIB tree has an unlabeled root with three subtrees. [Figure 37.2](#) illustrates the branch of the MIB tree that leads to the Infoblox enterprise MIBs. Each object in the MIB tree has a label that consists of a textual description and an OID (object identifier). An OID is a unique dotted-decimal number that identifies the location of the object in the MIB tree. Note that all OIDs begin with a dot (.) to indicate the root of the MIB tree.

As shown in [Figure 37.2](#), Infoblox is a branch of the Enterprise subtree. IANA (Internet Assigned Numbers Authority) administers the Enterprise subtree, which is designated specifically for vendors who define their own MIBs. The IANA-assigned enterprise number of Infoblox is 7779; therefore, the OIDs of all Infoblox MIB objects begin with the prefix .1.3.6.1.4.1.7779. In addition, IB-DNSSERV and IB-DHCPserv are branches of the Enterprise subtree as well. The Infoblox SNMP subtree branches down through two levels, ibProduct and ibOne, to the Infoblox MIBs: ibTrap, ibPlatformOne, ibDNSOne, and ibDHCPOne. The ibTrap MIB defines the traps that NIOS appliances send, and the ibPlatformOne, ibDNSOne, and ibDHCPOne MIBs provide information about the appliance. For detailed information about these MIBs, see [Infoblox MIBs](#) on page 1051.

Figure 37.2 MIB Hierarchy



MIB Objects

The Infoblox MIB objects were implemented according to the guidelines in RFCs 1155 and 2578. They specify two types of macros for defining MIB objects: OBJECT-TYPE and NOTIFICATION-TYPE. These macros contain clauses that describe the characteristics of an object, such as its syntax and its status. OBJECT-TYPE macros describe MIB objects, and NOTIFICATION-TYPE macros describe objects used in SNMP traps.

Each object in the ibPlatformOne, ibDNSone, and ibDHCPone MIBs contains the following clauses from the OBJECT-TYPE macro:

- OBJECT-TYPE: Provides the administratively-assigned name of the object.
- SYNTAX: Identifies the data structure of the object, such as integers, counters, and octet strings.
- MAX-ACCESS: Identifies the type of access that a management station has to the object. All Infoblox MIB objects provide read-only access.
- STATUS: Identifies the status of the object. Values are current, obsolete, and deprecated.
- DESCRIPTION: Provides a textual description of the object.
- INDEX or AUGMENTS: An object that represents a conceptual row must have either an INDEX or AUGMENTS clause that defines a key for selecting a row in a table.
- OID: The dotted decimal object identifier that defines the location of the object in the universal MIB tree.

The ibTrap MIB defines the SNMP traps that a NIOS appliance can send. Each object in the ibTrap MIB contains the following clauses from the NOTIFICATION-TYPE macro:

- NOTIFICATION-TYPE: Provides the administratively-assigned name of the object.
- OBJECTS: Provides an ordered list of MIB objects that are in the trap.
- STATUS: Identifies the status of the object. Values are current, obsolete, and deprecated.
- DESCRIPTION: Provides the notification information.

System Object IDs

Infoblox uses the SNMP system object identifier **sysObjectID** to identify Infoblox appliances. The following is a definition of **sysObjectID** from the SNMPv2 MIB, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*:

OBJECT-TYPE	sysObjectID
SYNTAX	Object Identifier
MAX-ACCESS	read-only
STATUS	current
DESCRIPTION	"The vendor's authoritative identification of the network management subsystem contained in the entity. This value is allocated within the SMI enterprises subtree (1.3.6.1.4.1) and provides an easy and unambiguous means for determining 'what kind of box' is being managed. For example, if vendor 'Flintstones, Inc.' was assigned the subtree 1.3.6.1.4.1.424242, it could assign the identifier 1.3.6.1.4.1.424242.1.1 to its 'Fred Router'."

[Table 37.3](#) lists the enterprise IDs and their corresponding Infoblox hardware platforms that an SNMP query can return when you request the sysObjectID value. Note that the IDs shown in the table do not include 1.3.6.1.4.1.7779.1. (the infobloxProducts prefix).

Table 37.3 *sysObjectID* for Infoblox Hardware

ID	Description	Definition
1000	ibDefault	Default environments, such as chroot
1001	ibRsp2	vNIOS appliances on Riverbed Services Platforms
1002	ibCisco	Cisco servers
1003	ibvm	vNIOS appliances on VMware ESX or ESXi servers
1004	ibvnios	Virtual NIOS
1031	ib4030	Infoblox-4030 appliances
1032	ib4030	Infoblox-4030 appliances
1101	ib1000	Infoblox -1000 appliances
1102	ib1200	Infoblox-1200 appliances
1103	ib500	Infoblox-500 appliances
1201	ib550	Infoblox-550 appliances
1202	ib1050	Infoblox-1050 appliances
1203	ib1550	Infoblox-1550 appliances
1204	ib1552	Infoblox-1552 appliances
1205	ib2000	Infoblox-2000 appliances
1206	ib250	Infoblox-250 appliances
1207	ib1220	Infoblox-1220 appliances
1301	ib550a	Infoblox-550-A appliances
1302	ib1050a	Infoblox-1050-A appliances
1303	ib1550a	Infoblox-1550-A appliances
1304	ib1552a	Infoblox-1552-A appliances
1305	ib1852a	Infoblox-1852-A appliances
1306	ib250a	Infoblox-250-A appliances
1307	ib2000a	Infoblox-2000-A appliances
1401	ib810	Trinzic 810 appliances
1402	ib820	Trinzic 820 appliances
1403	ib1410	Trinzic 1410 appliances
1404	ib1420	Trinzic 1420 appliances
1405	ib1400	Trinzic Reporting 1400 appliances
1406	ib800	Trinzic Reporting 800 appliances
1411	ib2200	Trinzic Reporting 2200 appliances
1412	ib2210	Trinzic 2210 appliances
1413	ib2220	Trinzic 2220 appliances
1421	ib4010	Infoblox-4010 appliances

ID	Description	Definition
1422	ib4030	Infoblox-4030 appliances
1423	ib4000	Infoblox-4000 appliances

INFOBLOX MIBS

You can configure a NIOS appliance as an SNMP-managed device so that an SNMP management station can send queries to the appliance and retrieve information from its MIBs. Perform the following tasks to access the Infoblox MIBs:

1. Configure a NIOS appliance to accept queries, as described in [Configuring SNMPv3 Users](#) on page 1040.
2. Load the MIB files onto the management system. To obtain the latest Infoblox MIB files:
 - a. From the **Data Management** tab, select the **Grid** tab -> **Grid Manager** tab, and then select **Download** -> **SNMP MIBs** from the Toolbar.
 - b. In the **Save As** dialog box, navigate to a directory to which you want to save the MIBs.
 - c. Click **Save**.
3. Use a MIB browser or SNMP management application to query the objects in each MIB.

The NIOS appliance allows read-only access to the MIBs. This is equivalent to the **Get** and **Get Next** operations in SNMP.

Loading the Infoblox MIBs

If you are using an SNMP manager toolkit with strict dependency checking, you must download the following Infoblox MIBs in the order they are listed:

1. IB-SMI-MIB.txt
2. IB-TRAP-MIB.txt
3. IB-PLATFORMONE-MIB.txt
4. IB-DNSONE-MIB.txt
5. IB-DHCPONE-MIB.txt
6. IB-DNSSERV-MIB.txt
7. IB-DHCPSEV-MIB.txt
8. IB-DHCPV6ONE-MIB.txt

In addition, if the SNMP manager toolkit you use requires a different MIB file naming convention, you can rename the MIB files accordingly.

NET-SNMP MIBs

NIOS appliances support NET-SNMP (formerly UCD-SNMP), a collection of applications used to implement the SNMP protocol. The NET-SNMP MIBs provide the top-level infrastructure for the SNMP MIB tree. They define, among other things, the objects in the SNMP traps that the agent sends when the SNMP engine starts and stops. For information about NET-SNMP and the MIB files distributed with NET-SNMP, refer to <http://net-snmp.sourceforge.net/>.

For SNMP traps to function properly, you must download the following NET-SNMP MIBs directly from <http://net-snmp.sourceforge.net/docs/mibs/>:

- NET-SNMP-MIB
- UCD-SNMP-MIB

Note: Ensure that you save the MIBs as text files in the directory to which you save all the other MIB files.

BGP4 MIB

Infoblox supports BGP4 (Border Gateway Protocol) for DNS anycast addressing. BGP is configured to send SNMP traps to neighboring routers, as defined in *RFC4273 Definitions of Managed Objects for BGP-4*. You must enable and configure the SNMP trap receiver on the Grid member for the member to send SNMP traps. For information, see [SNMP MIB Hierarchy](#) on page 1048.

The BGP protocol service is configured to send SNMP queries about BGP runtime data. The information is returned using the following OIDs and definitions:

OID	Definition
1.3.6.1.2.1.15.900.1.1	Number of peers
1.3.6.1.2.1.15.900.1.2	Number of active peers
1.3.6.1.2.1.15.900.1.3	Number of AS path entries
1.3.6.1.2.1.15.900.1.4	Number of BGP community entries
1.3.6.1.2.1.15.900.1.5	Total number of prefixes

For each configured BGP peer (a, b, c, d), the information is returned using the following OIDs and definitions:

OID	Definition
1.3.6.1.2.1.15.900.1.9.a.b.c.d.1	IP address: same as a.b.c.d
1.3.6.1.2.1.15.900.1.9.a.b.c.d.2	State: 0=down, 1=up
1.3.6.1.2.1.15.900.1.9.a.b.c.d.3	ASN
1.3.6.1.2.1.15.900.1.9.a.b.c.d.4	Prefixes
1.3.6.1.2.1.15.900.1.9.a.b.c.d.5	Up/Down time

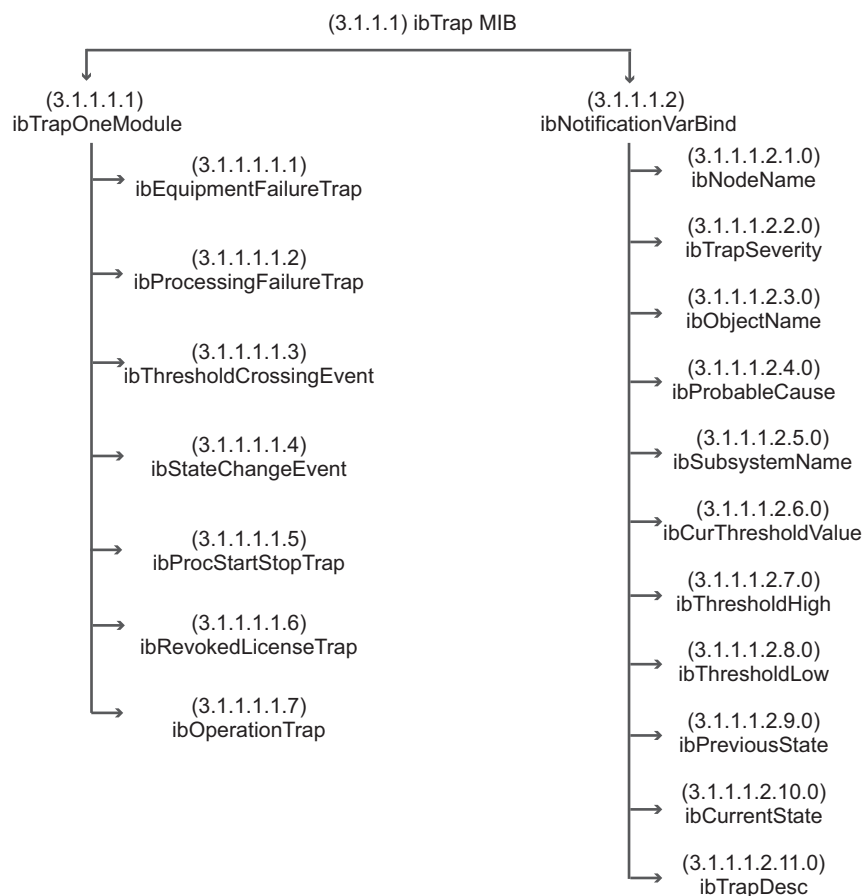
ibTrap MIB

NIOS appliances send SNMP traps when events, internal process failures, or critical service failures occur. The ibTrap MIB defines the types of traps that a NIOS appliance sends and the value that each MIB object represents. The Infoblox SNMP traps report objects which the ibTrap MIB defines. [Figure 37.3](#) illustrates the ibTrap MIB structure. It provides the OID and textual description for each object.

Note: OIDs shown in the illustrations and tables in this section do not include the prefix .1.3.6.1.4.1.7779.

The ibTrap MIB comprises two trees, ibTrapOneModule and ibNotificationVarBind. The ibTrapOneModule tree contains objects for the types of traps that a NIOS appliance sends. The ibNotificationVarBind tree contains objects that the Infoblox SNMP traps report. You cannot send queries for the objects in this MIB module. The objects are used only in the SNMP traps.

Figure 37.3 ibTrapOne MIB Structure



Interpreting Infoblox SNMP Traps

Depending on the SNMP management application your management system uses, the SNMP traps you receive may list the OIDs for all relevant MIB objects from both the `ibTrapOneModule` and `ibNotificationVarBind` trees. For OIDs that have string values, the trap lists the text. For OIDs that contain integers, you can use the tables in this section to find out the values. Some SNMP management applications list only the object names and their corresponding values in the SNMP traps. Whether or not your SNMP management application lists OIDs, you can use the tables in this section to find out the corresponding value and definition for each MIB object.

The following is a sample trap a NIOS appliance sends:

```
418:Jan 31 18:52:26 (none) snmptrapd[6087]: 2008-01-31 18:52:26 10.35.1.156 [UDP:
[10.35.1.156]:32772]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (1080)
0:00:10.80 SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.4.0
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "10.35.1.156"
SNMPv2-SMI::enterprises.
7779.3.1.1.1.2.3.0 = STRING: "ntp_sync" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.9.0 =
INTEGER: 15 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.10.0 = INTEGER: 16
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "The NTP service is out of
synchronization."
```

The sample trap lists the OIDs and their corresponding values that can help you identify the cause of an event or problem. To identify the possible cause and recommended actions for the trap, use the `ibTrapDesc` tables. For information, see [ibTrapDesc \(OID 3.1.1.1.2.11.0\)](#) on page 1070.

You can interpret the sample trap as follows:

Using the [ibTrapOneModule](#) table, you find out OID 7779.3.1.1.1.4.0 represents an Object State Change trap. This trap includes the following objects: `ibNodeName`, `ibObjectName`, `ibPreviousState`, `ibCurrentState`, and `ibtrapDesc`. For each object, the trap displays the OID and its corresponding value. The following is how you can interpret the rest of the trap:

- `ibNodeName` (OID 7779.3.1.1.1.2.1.0)
 - Using the [ibNotificationVarBind \(OID 3.1.1.1.2\)](#) table, you find out OID 7779.3.1.1.1.2.1.0 represents the MIB object `ibNodeName`, which is the IP address of the appliance on which the trap occurred. Therefore, the statement “7779.3.1.1.1.2.1.0 = STRING: “10.35.1.156” SNMPv2-SMI::enterprises.” tells you the IP address of the appliance on which the trap occurred.
- `ibObjectName` (OID 7779.3.1.1.1.2.3.0)
 - The statement “7779.3.1.1.1.2.3.0 = STRING: “ntp_sync” SNMPv2-SMI::enterprises.” tells you the MIB object `ibObjectName`, which is the name of the object for which the trap was generated, has a value of “ntp_sync” that indicates NTP synchronization issues.
- `ibPreviousState` (OID 7779.3.1.1.1.2.9.0)
 - The statement “7779.3.1.1.1.2.9.0 = INTEGER: 15 SNMPv2-SMI::enterprises.” tells you the MIB object `ibPreviousState`, which indicates the previous state of the appliance, has a value of 15. Using the [ibPreviousState and ibCurrentState Values](#) table, you know that 15 represents “ntp-sync-up”, which means the NTP server was up and running.
- `ibCurrentState` (OID 7779.3.1.1.1.2.10.0)
 - The statement “7779.3.1.1.1.2.10.0 = INTEGER: 16 SNMPv2-SMI::enterprises.” tells you the MIB object `ibCurrentState`, which indicates the current state of the appliance, has a value of 16. Using the [ibPreviousState and ibCurrentState Values](#) table, you know that 16 represents “ntp-sync-down”, which means the NTP server is now out of sync.
- `ibTrapDesc` (OID 7779.3.1.1.1.2.11.0)
 - The last statement “7779.3.1.1.1.2.11.0 = STRING: “The NTP service is out of synchronization.” states the description of the trap. Using the [Object State Change Traps](#) table for `ibTrapDesc`, you can find out the trap description and recommended actions for this problem.

Types of Traps (OID 3.1.1.1.1)

ibTrapOneModule defines the types of traps that the NIOS appliance can send. There are five types of SNMP traps. [Table 37.4](#) describe the types of traps and their objects in the ibTrapOneModule tree.

Table 37.4 *ibTrapOneModule*

OID	Trap Type	MIB Object	Description
3.1.1.1.1.1	Equipment Failure	ibEquipmentFailureTrap	<p>The NIOS appliance generates this trap when a hardware failure occurs. This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibTrapSevertiy • ibObjectName (equipment name) • ibProbableCause • ibTrapDesc <p>For a list of trap descriptions, see Equipment Failure Traps on page 1070.</p>
3.1.1.1.1.2	Processing and Software Failure	ibProcessingFailureTrap	<p>The NIOS appliance generates this trap when a failure occurs in one of the software processes. This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibTrapSeverity • ibSubsystemName • ibProbableCause • ibTrapDesc <p>For a list of trap descriptions, see Processing and Software Failure Traps on page 1071.</p>

OID	Trap Type	MIB Object	Description
3.1.1.1.1.3	Threshold Crossing	ibThresholdCrossingEvent	<p>The NIOS appliance generates this trap when any of the following events occur:</p> <ul style="list-style-type: none"> • System memory or disk usage exceeds 90%. • CPU usage exceeds the trigger value for 15 seconds. • A problem occurs when the Grid Master replicates its database to its Grid members. • DHCP address usage crosses a watermark threshold. For more information about tracking IP address usage, see Threshold Crossing Traps on page 1077. • The number or percentage of the DNS security alerts exceeds the thresholds of the DNS security alert triggers. <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibTrapSeverity • ibObjectName (threshold name) • ibCurThresholdvalue • ibThresholdHigh • ibThresholdLow • ibTrapDesc <p>For a list of trap descriptions, see Threshold Crossing Traps on page 1077.</p>
3.1.1.1.1.4	Object State Change	ibStateChangeEvent	<p>The NIOS appliance generates this trap when there is a change in its state, such as:</p> <ul style="list-style-type: none"> • The link to one of the configured ports goes down, and then goes back up again. • A failover occurs in an HA (high availability) pair configuration. • A member connects to the Grid Master. • An appliance in a Grid goes offline. <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibTrapSeverity • ibObjectName • ibPreviousState • ibCurrentState • ibTrapDesc <p>For a list of possible trap descriptions, see Object State Change Traps on page 1083.</p>

OID	Trap Type	MIB Object	Description
3.1.1.1.1.5	Process Started and Stopped	ibProcStartStopTrap	<p>The NIOS appliance generates this type of trap when any of the following events occur:</p> <ul style="list-style-type: none"> • When you enable HTTP redirection. • When you change the HTTP access setting. • When you change the HTTP session time out setting. • When a failover occurs in an HA pair configuration. <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibSubsystemName • ibTrapDesc <p>For a list of possible trap descriptions, see Process Started and Stopped Traps on page 1085.</p>
3.1.1.1.1.6		ibRevokedLicenseTrap	<p>The NIOS appliance generates this trap when a license is revoked.</p> <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibTrapSeverity • ibSubsystemName • ibTrapDesc
3.1.1.1.1.7		ibOperationTrap	<p>The NIOS appliance generates this trap when a software operation is noteworthy.</p> <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> • ibNodeName • ibTrapSeverity • ibSubsystemName • ibProbableCause • ibTrapDesc

Trap Binding Variables (OID 3.1.1.1.2)

Each SNMP trap contains information about the event or the problem. The Infoblox SNMP traps include MIB objects and their corresponding values from the `ibNotificationVarBind` module. [Table 37.5](#) describes the objects in the `ibNotificationVarBind` module.

Table 37.5 ibNotificationVarBind (OID 3.1.1.1.2)

Note: The OIDs shown in the following table do not include the prefix “.1.3.6.1.4.1.7779.”.

OID	MIB Object (Type)	Description
3.1.1.1.2.1.0	ibNodeName (DisplayString)	The IP address of the appliance on which the trap occurs. This may or may not be the same as the appliance that sends the trap. This object is used in all types of traps.
3.1.1.1.2.2.0	ibTrapSeverity (Integer)	The severity of the trap. There are five levels of severity. See Trap Severity (OID 3.1.1.1.2.2.0) on page 1059 for details.
3.1.1.1.2.3.0	ibObjectName (DisplayString)	The name of the object for which the trap was generated. This is used in the Equipment Failure traps, Threshold Crossing Event traps, and the Object State Change traps. The following shows what this object represents depending on the type of traps: <ul style="list-style-type: none"> Equipment Failure traps: The equipment name. Threshold Crossing Event traps: The object name of the trap. State Change traps: The object that changes state.
3.1.1.1.2.4.0	ibProbableCause (Integer)	The probable cause of the trap. See ibProbableCause Values on page 1059 for the definitions of each value.
3.1.1.1.2.5.0	ibSubsystemName (DisplayString)	The subsystem for which the trap was generated, such as NTP or SNMP. This object is used in the Processing and Software Failure traps and the Process Start and Stop traps. See ibSubsystemName Values (OID 3.1.1.1.2.9.0) on page 1066 for definitions.
3.1.1.1.2.6.0	ibCurThresholdValue (Integer)	The current value of the threshold counter. This object is used in the Threshold Crossing traps.
3.1.1.1.2.7.0	ibThresholdHigh (Integer)	This object is used in Threshold Crossing traps. For CPU usage, this is the trigger value of the SNMP trap. For DHCP address usage, this is the value of the high watermark. This only applies when the appliance sends a trap to indicate that DHCP address usage is above the configured high watermark value for a DHCP address range. For more information, see Threshold Crossing Traps on page 1077.
3.1.1.1.2.8.0	ibThresholdLow (Integer)	This object is used in Threshold Crossing traps. For CPU usage, this is the reset value of the SNMP trap. For DHCP address usage, this is the value for the low watermark. This only applies when the appliance sends a trap to indicate that DHCP address usage went below the configured low watermark value for a DHCP address range. For more information, see Threshold Crossing Traps on page 1077.
3.1.1.1.2.9.0	ibPreviousState (Integer)	The previous state of the appliance. This object is used in the Object State Change traps. See ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068 for definitions of each value.

OID	MIB Object (Type)	Description
3.1.1.1.2.10.0	ibCurrentState (Integer)	The current state of the appliance. This object is used in the Object State Change traps. See ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) on page 1068 for the definition of each value.
3.1.1.1.2.11.0	ibTrapDesc (DisplayString)	The description of the trap. This object is used in all types of traps. See ibTrapDesc (OID 3.1.1.1.2.11.0) on page 1070 for the description, possible cause, and recommended actions for each Infoblox SNMP trap.

Trap Severity (OID 3.1.1.1.2.2.0)

The object ibTrapSeverity defines the severity level for each Infoblox SNMP trap. There are five levels of severity.

Value	Description
1	Undetermined
2	Informational: Event that requires no further action.
3	Minor: Event that does not require user intervention.
4	Major: Event that requires user intervention and assistance from Infoblox Technical Support.
5	Critical: Problem that affects services and system operations, and requires assistance from Infoblox Technical Support.

ibProbableCause Values (OID 3.1.1.1.2.4.0)

[Table 37.6](#) lists the values that are associated with the object ibProbableCause (OID 3.1.1.1.2.4.0). These values provide information about the events, such as software failures, that trigger traps.

Table 37.6 ibProbableCause Values

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
0	ibClear	NA	SNMP Trap is cleared.
1	ibUnknown	NA	An unknown failure has occurred.
2	ibPrimaryDiskFailure	Primary drive is full.	NA
3	ibFanFailure-old	NA	Unused.
4	ibPowerSupplyFailure	A power supply failure has occurred.	NA

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
5	ibDBFailure	NA	A database daemon monitoring failure has occurred.
6	ibApacheSoftwareFailure	NA	An apache software failure has occurred.
7	ibSerialConsoleFailure	NA	An Infoblox serial console software failure has occurred.
11	ibControldSoftwareFailure	NA	A controld failure has occurred.
12	ibUpgradeFailure	NA	A system upgrade failure has occurred.
13	ibSNMPDFailure	NA	SNMP Server failure has occurred.
15	ibSSHDSOFTWAREFailure	NA	An SSH daemon failure has occurred.
16	ibNTPDSOFTWAREFailure	NA	An NTP daemon failure has occurred.
17	ibClusterdSoftwareFailure	NA	A cluster daemon failure has occurred.
18	ibLCDSoftwareFailure	NA	An LCD daemon failure has occurred.
19	ibDHCPdSoftwareFailure	NA	A DHCP daemon monitoring failure has occurred.
20	ibNamedSoftwareFailure	NA	A named daemon monitoring failure has occurred.
21	ibAuthServerGroupDown	NA	NAC Authentication server group is down.
22	ibAuthServerGroupUp	NA	NAC Authentication server group is up.
24	ibNTLMSOFTWAREFailure	NA	An NTLM monitoring failure has occurred.
25	ibNetBIOSDaemonFailure	NA	A NetBIOS daemon monitoring failure has occurred.
26	ibWindowBindDaemonFailure	NA	An NT domain service monitoring failure has occurred.

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
27	ibTFTPDSOFTWAREFAILURE	NA	A TFTP daemon failure has occurred.
28	ibUNUSED28 (28)	NA	Unused.
29	ibBackupSoftwareFailure	NA	Backup failed.
30	ibBackupDatabaseSoftwareFailure	NA	Database backup failed.
31	ibBackupModuleSoftwareFailure	NA	Module backup failed.
32	ibBackupSizeSoftwareFailure	NA	File size exceeded the quota. Backup failed.
33	ibBackupLockSoftwareFailure	NA	Another backup is in progress. Backup will not be performed.
34	ibHTTPFileDistSoftwareFailure	NA	An HTTP file distribution daemon failure has occurred.
35	ibOSPFSOFTWAREFAILURE	NA	An OSPF routing daemon failure has occurred.
36	ibAuthDHCPNamedSoftwareFailure	NA	An auth named server failure has occurred.
37	ibFan1Failure	Fan <n> failure has occurred.	NA
38	ibFan2Failure	Fan <n> failure has occurred.	NA
39	ibFan3Failure	Fan <n> failure has occurred.	NA
40	ibFan1OK	Fan <n> is OK.	NA
41	ibFan2OK	Fan <n> is OK.	NA
42	ibFan3OK	Fan <n> is OK.	NA
44	ibFTPDSoftwareFailure	NA	An FTP daemon failure has occurred.
46	ibPowerSupplyOK	NA	The power supply is OK.
47	ibWebUISoftwareFailure	NA	A WebUI software failure has occurred.

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
48	ibUNUSED	NA	Unused.
49	ibADAgentSyncFailure	NA	An AD agent client synchronizing domain data failure has occurred.
50	ibIFMAPSoftwareFailure	NA	An IF-MAP server failure has occurred.
51	ibCaptivePortalSoftwareFailure	NA	A Captive Portal service failure has occurred.
52	ibDuplicateIPAddressFailure	NA	A Duplicate IP Address has been detected.
53	ibBGPSoftwareFailure	NA	An BGP routing daemon failure has occurred.
54	ibRevokedLicense	NA	A license has been revoked.
58	ibGUILoginFailure	NA	An admin failed to log in to the GUI.
59	ibSerialConsoleLoginFailure	NA	An admin failed to log in to the serial console.
60	ibSystemReboot	NA	A system reboot was initiated.
61	ibSystemRestart	NA	A system restart was initiated.
62	ibZoneTransferFailure	NA	A zone transfer failure occurred.
63	ibDHCPLeaseConflict	NA	DHCP address conflicts with an existing lease.
64	ibDHCPAddressConflict	NA	DHCP address conflicts with an existing fixed address.
65	ibDHCPRangeConflict	NA	DHCP address conflicts with an existing range.
66	ibDHCPHostConflict	NA	DHCP address conflicts with an existing host.
67	ibSyslogFailure	NA	A syslog daemon failure occurred.

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
68	ibPowerSupply1Failure	NA	Power supply 1 failure has occurred.
69	ibPowerSupply2Failure	NA	Power supply 2 failure has occurred.
70	ibPowerSupply1OK	NA	Power supply 1 is OK.
71	ibPowerSupply2OK	NA	Power supply 2 is OK.
72	ibReportingTaskSwFailure	NA	A reporting task monitoring failure has occurred.
73	ibReportingDbBackupFailure	NA	A reporting db backup/restore operation failure has occurred.
74	ibFan4Failure	NA	Fan 4 failure has occurred.
75	ibFan5Failure	NA	Fan 5 failure has occurred.
76	ibFan6Failure	NA	Fan 6 failure has occurred.
77	ibFan7Failure	NA	Fan 7 failure has occurred.
78	ibFan8Failure	NA	Fan 8 failure has occurred.
79	ibFan4OK	NA	Fan 4 is OK.
80	ibFan5OK	NA	Fan 5 is OK.
81	ibFan6OK	NA	Fan 6 is OK.
82	ibFan7OK	NA	Fan 7 is OK.
83	ibFan8OK	NA	Fan 8 is OK.
2029	ibHSMGroupFailure	NA	A HSM operation failure has occurred in BIND.
2030	ibHSMGroupOK	NA	HSM operation has succeeded in BIND.
3001	ibRAIDIsOptimal	The system's RAID array is now running in an optimal state.	NA
3002	ibRAIDIsDegraded	The system's RAID array is in a degraded state.	NA
3003	ibRAIDIsRebuilding	The system's RAID array is rebuilding.	NA
3004	ibRAIDStatusUnknown	Unable to retrieve RAID array state!	NA
3005	ibRAIDBatteryIsOK	The system's RAID battery is OK.	NA

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
3006	ibRAIDBatteryFailed	A RAID battery failure has occurred.	NA
3007	ibRAIDOptimalMismatch	NA	The system's RAID array is now running in an optimal state (Mismatched disk(s) found).
3008	ibRAIDDegradeMismatch	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found).
3009	ibRAIDRebuildingMismatch	NA	The system's RAID array is rebuilding (Mismatched disk(s) found).
3010	ibRAIDBatteryWeak	NA	Please replace the system's RAID battery soon.
3011	ibRAIDIsDegradedDisk1	NA	The system's RAID array is in a degraded state. RAID Disk1 is EMPTY.
3012	ibRAIDIsDegradedDisk2	NA	The system's RAID array is in a degraded state. RAID Disk2 is EMPTY.
3013	ibRAIDIsDegradedDisk3	NA	The system's RAID array is in a degraded state. RAID Disk3 is EMPTY.
3014	ibRAIDIsDegradedDisk4	NA	The system's RAID array is in a degraded state. RAID Disk4 is EMPTY.
3015	ibRAIDIsDegradedDisk5	NA	The system's RAID array is in a degraded state. RAID Disk5 is EMPTY.
3016	ibRAIDIsDegradedDisk6	NA	The system's RAID array is in a degraded state. RAID Disk6 is EMPTY.
3017	ibRAIDIsDegradedDisk7	NA	The system's RAID array is in a degraded state. RAID Disk7 is EMPTY.
3018	ibRAIDIsDegradedDisk8	NA	The system's RAID array is in a degraded state. RAID Disk8 is EMPTY.
3019	ibRAIDIsRebuildingDisk1	NA	The system's RAID array is rebuilding. RAID Disk1 is OFFLINE.
3020	ibRAIDIsRebuildingDisk2	NA	The system's RAID array is rebuilding. RAID Disk2 is OFFLINE.

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
3021	ibRAIDIsRebuildingDisk3	NA	The system's RAID array is rebuilding. RAID Disk3 is OFFLINE.
3022	ibRAIDIsRebuildingDisk4	NA	The system's RAID array is rebuilding. RAID Disk4 is OFFLINE.
3023	ibRAIDIsRebuildingDisk5	NA	The system's RAID array is rebuilding. RAID Disk5 is OFFLINE.
3024	ibRAIDIsRebuildingDisk6	NA	The system's RAID array is rebuilding. RAID Disk6 is OFFLINE.
3025	ibRAIDIsRebuildingDisk7	NA	The system's RAID array is rebuilding. RAID Disk7 is OFFLINE.
3026	ibRAIDIsRebuildingDisk8	NA	The system's RAID array is rebuilding. RAID Disk8 is OFFLINE.
3027	ibRAIDDegradedMismatchDisk 1	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk1 is EMPTY.
3028	ibRAIDDegradedMismatchDisk 2	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk2 is EMPTY.
3029	ibRAIDDegradedMismatchDisk 3	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk3 is EMPTY.
3030	ibRAIDDegradedMismatchDisk 4	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk4 is EMPTY.
3031	ibRAIDDegradedMismatchDisk 5	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk5 is EMPTY.
3032	ibRAIDDegradedMismatchDisk 6	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk6 is EMPTY.

Value	OID 3.1.1.2.4.0 ibProbableCause	Equipment Failure Traps	Processing and Software Failure Traps
3033	ibRAIDDegradedMismatchDisk 7	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk7 is EMPTY.
3034	ibRAIDDegradedMismatchDisk 8	NA	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk8 is EMPTY.
3035	ibRAIDRebuildingMismatchDisk 1	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk1 is OFFLINE.
3036	ibRAIDRebuildingMismatchDisk 2	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk2 is OFFLINE.
3037	ibRAIDRebuildingMismatchDisk 3	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk3 is OFFLINE.
3038	ibRAIDRebuildingMismatchDisk 4	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk4 is OFFLINE.
3039	ibRAIDRebuildingMismatchDisk 5	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk5 is OFFLINE.
3040	ibRAIDRebuildingMismatchDisk 6	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk6 is OFFLINE.
3041	ibRAIDRebuildingMismatchDisk 7	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk7 is OFFLINE.
3042	ibRAIDRebuildingMismatchDisk 8	NA	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk8 is OFFLINE.

ibSubsystemName Values (OID 3.1.1.1.2.9.0)

[Table 37.7](#) lists the values that are associated with the object `ibSubsystemName` (OID 3.1.1.1.2.9.0). These values provide information about the subsystems that trigger the traps.

Table 37.7 *ibSubsystemName Values*

Value	OID 3.1.1.1.2.9.0 ibSubsystemName
0	Uses the original ibObjectName and ibSubsystemName when the trap is cleared. The process failure trap is appended to the CLEAR trap descriptions.
1	N/A
2	N/A
3	N/A
4	N/A
5	Db_jnlId
6	httpd
7	serial_console
11	controld
12	N/A
13	Snmpd
15	Sshd
16	Ntpd
17	Clusterd
18	Lcd
19	Dhcpd
20	Named
24	NTLM
25	Netbiosd
26	Winbindd
27	Tftpd
29	N/A
30	db_dump
31	N/A
32	Scheduled_backups
33	N/A
34	HTTPd
35	OSPF

ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0)

The ibPreviousState object indicates the state of the appliance before the event triggered the trap. The ibCurrentState object indicates the current state of the appliance. [Table 37.8](#) shows the message and description for each state.

Table 37.8 ibPreviousState and ibCurrentState Values

Value	Description	Definition
0	None	No previous state.
1	ha-active	The HA pair is in ACTIVE state.
2	ha-passive	The HA pair is in PASSIVE state.
3	ha-initial	The HA pair is in INITIAL state.
4	Grid-connected	The Grid member is connected to the Grid Master.
5	Grid-disconnected	The Grid member is not connected to the Grid Master.
6	enet-link-up	The ethernet port link is active.
7	enet-link-down	The ethernet port link is inactive.
8	replication-online	The replication is online.
9	replication-offline	The replication is offline.
10	replication-snapshotting	The replication is snapshotting.
11	service-up	The service is up.
12	service-down	The service is down.
13	ha-replication-online	The HA pair replication is online.
14	ha-replication-offline	The HA pair replication is offline.
15	ntp-syn-up	The NTP server is synchronizing.
16	ntp-syn-down	The NTP server is out of synchronization.
17	ms-server-up	Microsoft server is up.
18	ms-server-down	Microsoft server is down.
19	ms-service-up	Microsoft service connection is active.
20	ms-service-down	Microsoft service connection is inactive.
21	nac-server-group-down	NAC Authentication server group is down.
22	nac-server-group-up	NAC Authentication server group is up.
23	mgm-service-up	MGM service is active.
24	mgm-service-down	MGM service is inactive.
25	ha-active-active	HA Pair is in Dual Active state.
26	ftp-service-working	FTP service is working.
27	ftp-service-failed	FTP service failed.
28	ftp-service-inactive	FTP service is inactive.
29	tftp-service-working	TFTP service is working.

Value	Description	Definition
30	tftp-service-failed	TFTP service failed.
31	tftp-service-inactive	TFTP service is inactive.
32	dns-service-working	DNS service is working.
33	dns-service-failed	DNS service failed.
34	dns-service-inactive	DNS service is inactive.
35	ntp-service-working	NTP service is working.
36	ntp-service-failed	NTP service failed.
37	ntp-service-inactive	NTP service is inactive.
38	http-file-dist-service-working	HTTP File Dist service is working.
39	http-file-dist-service-failed	HTTP File Dist service failed.
40	http-file-dist-service-inactive	HTTP File Dist service is inactive.
41	bloxtools-service-working	bloxTools service is working.
42	bloxtools-service-warnin	bloxTools service is in warning state.
43	bloxtools-service-failed	bloxTools service failed.
44	bloxtools-service-inactive	bloxTools service is inactive.
45	dhcp-service-working	DHCP service is working.
46	dhcp-service-warning	DHCP service is in warning state.
47	dhcp-service-failed	DHCP service failed.
48	dhcp-service-inactive	DHCP service is inactive.
49	captive-portal-service-working	Captive portal service is working.
50	captive-portal-service-failed	Captive portal service failed.
51	captive-portal-service-inactive	Captive portal service inactive.
52	ifmap-service-working	IF-MAP service is working.
53	ifmap-service-failed	IF-MAP service failed.
54	ifmap-service-inactive	IF-MAP service inactive.
56	hsm-group-down	HSM operation failed.
57	hsm-group-up	HSM operation succeeded.
59	reporting-service-working	Reporting service is working.
60	reporting-service-failed	Reporting service failed.
61	reporting-service-inactive	Reporting service inactive.

ibTrapDesc (OID 3.1.1.1.2.11.0)

The ibTrapDesc object lists the trap messages of all Infoblox SNMP traps. This section lists all the SNMP traps by their trap types. Each trap table describes the trap message, severity, cause, and recommended actions.

Note: Contact Infoblox Technical Support for assistance when the recommended actions do not resolve the problems.

Equipment Failure Traps

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
Primary Drive Full			
Primary drive is full.	Major	The primary disk drive reached 100% of usage.	Review the syslog file to identify the possible cause of this problem.
Fan Monitoring			
Fan <n> failure has occurred.	Minor	The specified fan failed. The fan number <n> can be 1, 2, or 3.	Inspect the specified fan for mechanical or electrical problems.
Fan <n> is OK.	Informational	The specified fan is functioning properly. The fan number <n> can be 1, 2, or 3.	No action is required.
Power Supply Failure: monitored at 1 minute			
A power supply failure has occurred.	Major	The power supply failed.	Inspect the power supply for the possible cause of the failure.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
RAID monitoring, at 1 minute interval			
A RAID battery failure has occurred.	Major	The system RAID battery failed. The alert light is red.	Inspect the battery for the possible cause of the failure.
The system's RAID battery is OK.	Informational	The system RAID battery is charging and functioning properly. The alert light changed from red to green.	No action is required.
Unable to retrieve RAID array state!	Undetermined	The appliance failed to retrieve the RAID array state. The alert light is red.	Review the syslog file to identify the possible cause of this problem.
The system's RAID array is now running in an optimal state.	Informational	The RAID system is functioning at an optimal state.	No action is required.
The system's RAID array is in a degraded state.	Major	The RAID system is degrading.	Review the syslog file to identify the possible cause of this problem.
The system's RAID array is rebuilding.	Minor	The RAID system is rebuilding.	No action is required.

Processing and Software Failure Traps

Note: The `ibSubsystemName` object is associated with certain traps of the Processing and Software Failure traps. Therefore, you cannot map all the traps of the Processing and Software Failure traps with the `ibSubsystemName`. If there is no value in the `ibSubsystemName`, then it belongs to the N/A category. For more information on the values for the `ibSubsystemName`, see the [Table 37.7](#) on page 1067.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
Named Daemon Failure			
A named daemon monitoring failure has occurred.	Critical	The named process failed.	Review the syslog file to identify the possible cause of this problem.
DHCP Daemon Failure			
A DHCP daemon monitoring failure has occurred.	Critical	The dhcpd process failed.	Review the syslog file to identify the possible cause of this problem.
SSH Daemon Failure			
An SSH daemon failure has occurred.	Major	The sshd process failed.	Review the syslog file to identify the possible cause of this problem.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
NTP Daemon Failure, monitored every 10 minutes			
An NTP daemon failure has occurred.	Major	The ntpd process failed.	Review the syslog file to identify the possible cause of this problem.
Cluster Daemon Failure			
A cluster daemon failure has occurred.	Critical	The clusterd process failed.	Review the syslog file to identify the possible cause of this problem.
LCD Daemon Failure			
An LCD daemon failure has occurred.	Major	The LCD process failed. The alert light is yellow.	<ol style="list-style-type: none"> 1. Inspect the LCD panel for the possible cause of this problem. 2. Review the syslog file to identify the possible cause of this problem.
Apache Software httpd failure, monitored every 2 minutes			
An Apache software failure has occurred.	Critical	The request to monitor the Apache server failed.	Review the syslog file to identify the possible cause of this problem.
Serial Console Failure			
An Infoblox serial console software failure has occurred.	Major	The Infoblox serial console failed.	Review the syslog file to identify the possible cause of this problem.
Controld Software Failure			
A controld failure has occurred.	Critical	The controld process failed.	Review the syslog file to identify the possible cause of this problem.
SNMP Sub-agent Failure			
An SNMP server failure has occurred.	Major	The one-subagent process failed.	Review the syslog file to identify the possible cause of this problem.
TFTPD and FTPD Failure			
A TFTP daemon failure has occurred.	Critical	The tftpd process failed.	Review the syslog file to identify the possible cause of this problem.
An FTP daemon failure has occurred.	Critical	The ftpd process failed.	Review the syslog file to identify the possible cause of this problem.
HTTP File Distribution, monitored at 10 second intervals			
An HTTP file distribution daemon failure has occurred.	Critical	The HTTP file distribution process failed.	Review the syslog file to identify the possible cause of this problem.
DNS ONE quagga Processes (zebra & ospfd)			
An OSPF routing daemon failure has occurred.	Critical	Either the zebra process or the ospfd process failed. Both the zebra and ospfd process belongs to ospf subsystem.	Review the syslog file to identify the possible cause of this problem.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
Backup Failure			
Backup failed.	Minor	<p>The backup failed. One of the following could be the cause of the failure:</p> <ul style="list-style-type: none"> • The appliance could not access a backup directory. • The backup was interrupted by one of the following signals: SIGINT, SIGHUP, or SIGTERM. • Incorrect login or connection failure in an FTP backup. • The backup failed to create temporary files. 	Review the syslog file to identify the possible cause of this problem.
Database Backup Failure			
Database backup failed.	Not implemented	The db_dump process failed.	Review the syslog file to identify the possible cause of this problem.
Backup Module Failure			
Module backup failed.	Not implemented	The backup of product-specific files failed.	Review the syslog file to identify the possible cause of this problem.
Backup File Size Exceeded			
File size exceeded the quota. Backup failed.	Not implemented	The backup failed because the file size exceeded the limit of 5GB.	Limit the size of the backup file to less than 5GB.
Another backup is in progress. Backup will not be performed.	Not implemented	The backup failed because of an attempt to back up or merge files while another backup or restore was in progress.	Wait until the backup or restore is complete before starting another backup.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
Watchdog Process Monitoring			
WATCHDOG: ‹registered client name‹ failed on ‹server IP address‹	Critical	The watchdog process detected a registered client failure on a specific server. The ‹registered client name‹ could be one of the following: <ul style="list-style-type: none"> • Clusterd_timeout • DB_Sentinel • Process_Manager • Clusterd_monitor • Disk_monitor 	Review the syslog file to identify the possible cause of this problem.
Microsoft Server			
Microsoft server <i>hostname</i> has failed.	Major	The Microsoft server could not be reached.	Check that the Microsoft server is connected to the network and configured properly.
Microsoft server <i>hostname</i> is OK.	Informational	The Microsoft server can be reached and is functioning properly.	No action is required.
Microsoft DNS/DHCP Service			
Service connection to Microsoft DNS server <i>hostname</i> has failed.	Major	The Microsoft DNS service is not responding.	Check that the DNS service is configured and running on the Microsoft server.
Service connection to Microsoft DHCP server <i>hostname</i> has failed.	Major	The Microsoft DHCP service is not responding.	Check that the DHCP service is configured and running on the Microsoft server.
Service connection to Microsoft DNS server <i>hostname</i> is OK.	Informational	The Microsoft DNS service is responding.	No action is required.
Service connection to Microsoft DHCP server <i>hostname</i> is OK.	Informational	The Microsoft DHCP service is responding.	No action is required.
NAC Authentication Server Group			
NAC Authentication server group is down	Major	None of the servers in the NAC authentication server group can be reached.	Review the syslog.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
NAC Authentication server group is up	Informational	The NAC authentication server group is responding.	No action is required.
GUI Login			
An admin failed to log in to the GUI.	Major	An admin failed to log in to the GUI.	Check the credentials of the admin.
Serial Console Login			
An admin failed to log in to the serial console.	Major	An admin failed to log in through the serial console.	Check the credentials and permissions, and check that the serial console is enabled.
Reboot			
A system reboot was initiated.	Informational	A system reboot command was sent.	No action is required.
DHCP Lease Conflict			
DHCP address conflicts with an existing lease.	Major	The discovery process found a DHCP lease conflict.	In the IP Map or List panel, select a conflicting address, and then click Resolve Conflict . For more information, see Resolving DHCP Lease Conflicts on page 514.
DHCP Fixed Address Conflict			
DHCP address conflicts with an existing fixed address.	Major	The discovery process found a fixed address conflict.	In the IP Map or List panel, select a conflicting address, and then click Resolve Conflict . For more information, see Resolving Fixed Address Conflicts on page 514.
DHCP Range Conflict			
DHCP address conflicts with an existing range.	Major	The discovery process found a conflict with an existing range.	In the IP Map or List panel, select a conflicting address, and then click Resolve Conflict . For more information, see Resolving DHCP Range Conflicts on page 514.
DHCP Host Conflict			
DHCP address conflicts with an existing host.	Major	The discovery process found a conflict with an existing host address.	In the IP Map or List panel, select a conflicting address, and then click Resolve Conflict . For more information, see Resolving Host Conflicts on page 515.
Syslog Daemon Failure			
A syslog daemon failure occurred.	Critical	Syslog process stopped.	Review the syslog file to identify the possible cause of this problem.
Process Stop/Start			

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
The system stopped and started a process.	Major	The system restarted a process.	Review the syslog file to identify the possible cause of this problem.
Zone Transfer Failed			
A zone transfer failure occurred.	Major	A zone transfer failed.	Review the syslog file to identify the possible cause of this problem.

Threshold Crossing Traps

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	ibObjectName OID 3.1.1.1.2.3.0	Description/Cause	Recommended Actions
System Memory Usage				
System has run out of memory.	Major	memory	<p>The appliance ran out of memory. The appliance encountered this problem when one of the following occurred:</p> <ul style="list-style-type: none"> • The total free memory on the appliance was less than or equal to 0%. • The total physical memory was less than the total free memory. • The percentage of free memory compared to the total physical memory was less than 5%, and the free swap percentage was less than 80%. • The percentage of free memory compared to the total physical memory was less than 5%, plus the numbers of both swap INs and swap OUTs were greater than or equal to 3,200. • The percentage of free memory compared to the total physical memory was between 5% and 10%, the free swap percentage was greater than or equal to 80%, plus the numbers of both swap INs and swap OUTs were greater than or equal to 3,200. • The percentage of free memory compared to the total physical memory was greater than 10%, the free swap percentage was less than 80%, plus the numbers of both swap INs and swap OUTs were greater than or equal to 3,200. <hr/> <p>Note: Free memory = free physical RAM + free cache buffers. The high threshold for swap pages is 3,200.</p> <hr/>	Review the syslog file to identify the possible cause of this problem.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	ibObjectName OID 3.1.1.1.2.3.0	Description/Cause	Recommended Actions
System memory usage is over the configured Trigger value.	Minor	memory	<p>The memory usage on the appliance exceeded the configured Trigger value. For more information, see Defining Thresholds for Traps on page 1043.</p> <p>The appliance encountered this problem when one of the following occurred:</p> <ul style="list-style-type: none"> • The percentage of free memory compared to the total physical memory was less than 5%, and the free swap percentage was less than 90%. • The percentage of free memory compared to the total physical memory was less than 5%, plus the number of swap INs was less than 3,200 and the number of swap OUTs was greater than or equal to 3,200. • The percentage of free memory compared to the total physical memory was between 5% and 10%, and the free swap percentage was less than 80%. • The percentage of free memory compared to the total physical memory was greater than 5%, plus the number of swap INs was less than 3,200 and the number of swap OUTs was greater than or equal to 3,200. <hr/> <p>Note: Free memory = free physical RAM + free cache buffers. The high threshold for swap pages is 3,200.</p> <hr/>	Review the syslog file to identify the possible cause of this problem.
System memory is OK.	Minor	memory	The memory usage on the system is at or below the Reset value after it went above the Trigger value.	No action is required.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	ibObjectName OID 3.1.1.1.2.3.0	Description/Cause	Recommended Actions
Primary Hard Drive Usage (monitored every 30 seconds)				
System primary hard disk usage is over the configured Trigger value. The default value is 85.	Minor	disk usage	The appliance sends this trap when primary hard disk usage first exceeds the configured Trigger value. The alert light is yellow. For more information, see Defining Thresholds for Traps on page 1043.	Review the syslog file to identify the possible cause of this problem.
Primary drive is full.	Major	disk usage	The primary hard disk usage exceeded 95%. The alert light is red.	Review the syslog file to identify the possible cause of this problem.
Primary drive usage is OK.	Minor	disk usage	The appliance sends this trap when the primary hard disk usage first moves at or below the configured Reset value after it exceeded the Trigger value. The default is 70. The alert light is green.	No action is required.
CPU Usage				
CPU usage above threshold value.	Major	cpu usage	CPU usage exceeded the trigger value for 15 seconds. For more information, see Defining Thresholds for Traps on page 1043.	Monitor CPU usage.
CPU usage OK.	Minor	cpu usage	CPU usage dipped below the reset value after the “CPU usage above threshold value” trap was sent.	No action is required.
Note: Use the CLI command <code>set thresholdtrap</code> to enable the CPU usage trap and configure the trigger and reset values. For information, refer to the <i>Infoblox CLI Guide</i> .				
Replication Statistics Monitoring				
Grid queue replication problem.	NA	For send trap: Cluster_Send_Queue For receive trap: Cluster_Recv_Queue	The system encountered this problem when all of the following conditions occurred: <ul style="list-style-type: none"> • The node was online. • The number of the replication queue being sent from the master column was greater than 0, or the number of the queue received was greater than 0. • It was more than 10 minutes since the last replication queue was sent and monitored. 	Review the syslog file to identify the possible cause of this problem.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	ibObjectName OID 3.1.1.1.2.3.0	Description/Cause	Recommended Actions
DHCP Range Threshold Crossing				
DHCP threshold crossed: Member: <i><DHCP server node VIP></i> Network: <i><network>/<network view></i> Range: <i><DHCP range>/<network view></i> High: <i><high percentage></i> (95% by default) Low : <i><low percentage></i> (0% by default) Current Usage: <i><current usage percentage></i> Active Leases: <i><number of active leases></i> Available Leases: <i><number of available leases></i> Total Addresses: <i><total addresses></i>	N/A	Threshold	The system encountered this problem when one of the following conditions occurred: <ul style="list-style-type: none"> The address usage in the DHCP range is greater than the configured High Trigger value and when it first dips below the Reset value after it hit the Trigger value. The address usage in the DHCP range goes below the Low Trigger value and when it first goes above the Reset value after it hit the Trigger value. 	Review the syslog file to identify the possible cause of this problem.
DHCP DDNS Updates Deferred				
DHCP DNS updates deferred: Retried at least once: <i><number of retries></i> Maximum number of deferred updates since start of problem episode (or restart): <i><max number></i>	N/A	Threshold	The DNS updates were deferred because of DDNS update errors.	Review the syslog file to identify the possible cause of this problem.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	ibObjectName OID 3.1.1.1.2.3.0	Description/Cause	Recommended Actions
Database Capacity Usage				
Over 85% database capacity used.	Minor	db_usage	The appliance database usage exceeded 85%.	Increase the database capacity.
Database capacity used is OK.	Minor	db_usage	The appliance database usage is less than 85%.	No action is required.
DNS Monitor				
DNS Monitor	Major	For invalid ports: "dns_security_port" For invalid TXIDs: "dns_security_txid"	<p>DNS security alert. There were <i>actual</i> DNS responses to {invalid ports with invalid TXID} in the last minute, comprising <i>percent%</i> of all responses. Primary sources: <i>ip_address</i> sent <i>count</i>, <i>ip_address</i> sent <i>count</i>. where</p> <ul style="list-style-type: none"> <i>actual</i> is the total number of DNS responses arrive on invalid ports or have invalid TXIDs. <i>percent%</i> is the percentage of invalid DNS responses over the total number of DNS responses. <i>ip_address</i> is the IP address of the primary source that generated the invalid DNS responses. <i>count</i> is the number of invalid responses generated by the specified IP address. <p>Example: DNS security alert. There were 1072 DNS responses to invalid ports in the last minute, comprising 92% of all responses. Primary sources: 10.0.0.0 sent 1058, 2.2.2.2 sent 14.</p>	<ol style="list-style-type: none"> Review the following: <ul style="list-style-type: none"> DNS alert status syslog file Limit access or block connections from the primary sources. For information, see Configuring Rate Limiting Rules on page 1028.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	ibObjectName OID 3.1.1.1.2.3.0	Description/Cause	Recommended Actions
RootFS Partition Monitor				
Root file system is full.	Major	Root filesystem	The Root filesystem usage exceeded the maximum.	Review the syslog file to identify the possible cause of this problem.
Root file system disk usage is over the configured Trigger value.	Minor	Root filesystem	The appliance sends this trap when the Root filesystem usage first exceeds the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.
Root file system disk usage is OK.	Minor	Root filesystem	The appliance sends this trap when the Root filesystem disk usage first moves at or below the configured Reset value after it exceeded the Trigger value. For information on setting the Trigger and Reset values, see Defining Thresholds for Traps on page 1043.	No action
Reporting				
Reporting drive is full.	Major	Reporting	Reporting drive reached the maximum capacity.	Review the syslog file to identify the possible cause of this problem.
Reporting drive usage is over the configured Trigger value.	Minor	Reporting	The appliance sends this trap when the Reporting volume first exceeds the configured Trigger value. The default Trigger value is 80.	Review the syslog file to identify the possible cause of this problem.
Reporting drive usage is OK.	Minor	Reporting	The appliance sends this trap when the Reporting volume first moves at or below the configured Reset value after it exceeded the Trigger value. The default Reset value is 71. For information on setting the Trigger and Reset values, see Defining Thresholds for Traps on page 1043.	No action

Object State Change Traps

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
Service Shutdown			
Shutting down services due to database snapshot.	Major	The appliance is shutting down its services while synchronizing the database with the Grid Master.	No action is required.
Network Interfaces Monitoring			
LAN port link is down. Please check the connection.	Major	The LAN port is up, but the link is down.	Check the LAN link connection.
HA port link is down. Please check the connection.	Major	The HA port is up, but the link is down.	Check the HA link connection.
MGMT port link is down. Please check the connection.	Major	The MGMT port is enabled, but the link is down.	Check the MGMT link connection.
LAN port link is up.	Major	The LAN port link is up and running.	No action is required.
HA port link is up.	Major	The HA port link is up and running.	No action is required.
MGMT port link is up.	Major	The MGMT port link is up and running.	No action is required.
HA State Change from Initial to Active			
The node has become ACTIVE.	Informational	A node in an HA pair becomes active. The HA pair starts up.	No action is required.
HA State Change from Passive to Active			
The node has become ACTIVE.	Informational	The node changed from a passive to an active node.	No action is required.
HA State Change to Active-Active			
The node is in an ACTIVE-ACTIVE state.	Informational	The node is in the active state.	No action is required.
HA State Change from Initial to Passive			
The node has become PASSIVE.	Informational	A node in an HA pair becomes passive. The HA pair starts up, and the node is not a Grid Master candidate.	No action is required.
Node Connected to Grid			
The Grid member is connected to the Grid Master.	Informational	The Grid member joined the Grid, and it is not a Grid Master candidate.	No action is required.
Node Disconnected from Grid			

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
The Grid member is not connected to the Grid Master.	Informational	The Grid member lost its connection to the Grid Master.	No action is required.
Replication State Monitoring			
ha-replication-online (13)	Informational	The HA replication is online.	No action is required.
ha-replication-offline (14)	Informational	The HA replication is offline.	No action is required.
NTP is out of sync, monitored every 30 seconds			
The NTP server is out of synchronization.	Major	The Infoblox NTP server and the external NTP server are not synchronized.	Review the syslog file to identify the possible cause of this problem.
NTP service returned to working state.	Informational	The NTP service started working again.	No action is required.
DHCP service state change			
DHCP service returned to working state.	Informational	The DHCP service started working again.	No action is required.
DHCP service is in a warning state.	Informational	The DHCP service is in a warning state.	Review the syslog file
DHCP service became inactive.	Informational	The DHCP service became inactive.	Check if an admin disabled the service.
DNS service state change			
DNS service returned to working state.	Informational	The DNS service started working again.	No action is required.
DNS service is in a warning state.	Informational	The DNS service is in a warning state.	Review the syslog file
DNS service became inactive.	Informational	The DNS service became inactive.	Check if an admin disabled the service.
NTP service state change			
NTP service resumed synchronization.	Informational	The NTP service started working again.	No action is required.
NTP service became inactive.	Informational	The NTP service became inactive.	Check if an admin disabled the service.
TFTP service state change			
TFTP service returned to working state.	Informational	The TFTP service started working again.	No action is required.
TFTP service became inactive.	Informational	The TFTP service became inactive.	Check if an admin disabled the service.
FTP service state change			
FTP service returned to working state.	Informational	The FTP service started working again.	No action is required.

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
FTP service became inactive.	Informational	The FTP service became inactive.	Check if an admin disabled the service.
HTTP service state change			
HTTP service returned to working state.	Informational	The HTTP service started working again.	No action is required.
bloxTools service state change			
bloxTools service returned to working state.	Informational	The bloxTools service started working again.	No action is required.
bloxTools service is in a warning state	Informational	The bloxTools service is in a warning state.	Review the syslog file
bloxTools service became inactive.	Informational	The bloxTools service became inactive.	Check if an admin disabled the service.
bloxTools service failed.	Critical	The bloxTools daemon failed.	Review the syslog file

Process Started and Stopped Traps

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
Httpd Start			
The process started normally.	Informational	The httpd process started.	No action is required.
Httpd Stop			
The process stopped normally.	Informational	The httpd process stopped.	No action is required.
Process Stop/Start			
The system stopped and started a process.	Major	The system restarted a process.	No action is required.
Zone Transfer Failed			
A zone transfer failure occurred.	Critical	A zone transfer failed.	Review the syslog file

Revoked License Trap

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
Revoked License			
This trap is generated when a license is revoked	Critical	A license was revoked.	Obtain and install new license

ibPlatformOne MIB

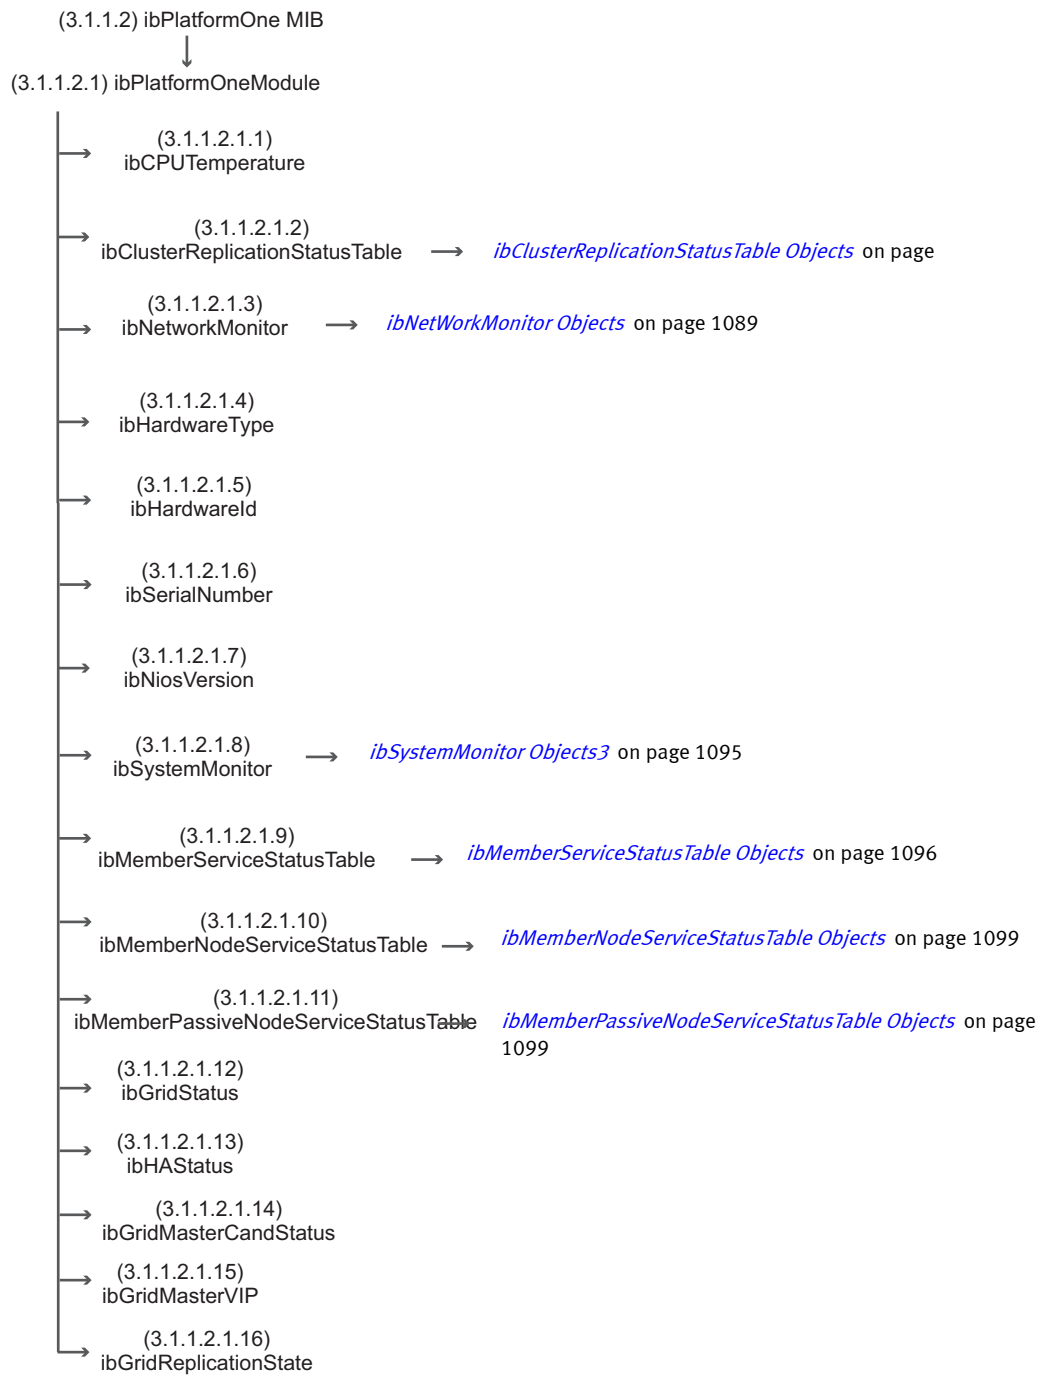
The ibPlatformOne MIB provides information about the CPU temperature of the appliance, the replication status, the average latency of DNS requests, DNS security alerts, CPU and memory utilization of the appliance, and the Infoblox service status. [Figure 37.4](#) illustrates the structure of the PlatformOne MIB. (Note that the OIDs in the illustration do not include the prefix .1.3.6.1.4.1.7779.)

The ibPlatformOne MIB contains the following objects:

- **ibCPUTemperature** (IbString) tracks the CPU temperature of the appliance.
- **ibClusterReplicationStatusTable** provides information in tabular format about the replication status of the appliance. For information, see [ibClusterReplicationStatusTable](#) on page 1088.
- **ibNetworkMonitor** provides information about the average latency of authoritative and nonauthoritative replies to DNS queries for different time intervals. It also provides information about invalid DNS responses that arrive on invalid ports or have invalid DNS transaction IDs. For information, see [ibNetwork Monitor](#) on page 1089.
- **ibHardwareType** (IbString) provides information about the hardware platform. For an Infoblox appliance, it provides the model number of the Infoblox hardware platform. For vNIOS appliances, it identifies whether the hardware platform is Riverbed or VMware.
- **ibHardwareId** (IbString) provides the hardware ID of the NIOS appliance.
- **ibSerialNumber** (IbString) provides the serial number of the Infoblox hardware platform.
- **ibNiosVersion** (IbString) provides the version of the NIOS software.
- **ibSystemMonitor** provides information about the CPU and memory utilization of the appliance. For information, see [ibSystemMonitor](#) on page 1095.
- **ibGridStatus** provides information about an appliance. It indicates whether the appliance is a Grid Master, member, or an independent appliance.
- **ibHAStatus** provides information about the HA status of a member. It indicates if the member is part of an HA configuration, and if it is the active or passive node.
- **ibGridMasterCandStatus** indicates if a member is a Grid Master candidate.
- **ibGridMasterVIP** provides the Grid Master virtual IP address.
- **ibGridReplicationState** provides information about the replication status.

The ibPlatformOne MIB also contains the following tables that provide status of the Infoblox services as well as system and hardware services on the appliance you query:

- **ibMemberServiceStatusTable** provides status of the Infoblox services, such as the DNS and DHCP services, on a queried appliance. For information, see [ibMemberServiceStatusTable](#) on page 1095.
- **ibMemberNodeServiceStatusTable** provides status of the system and hardware services on a queried appliance. For information, see [ibMemberNodeServiceStatusTable](#) on page 1098.
- **ibMemberPassiveNodeServiceStatusTable** provides status of the system and hardware services on the passive node of an HA pair if the queried appliance is the VIP or the active node of an HA pair. For independent appliances and the passive nodes of HA pairs, this table does not display any status. For information, see [ibMemberPassiveNodeServiceStatusTable](#) on page 1099.

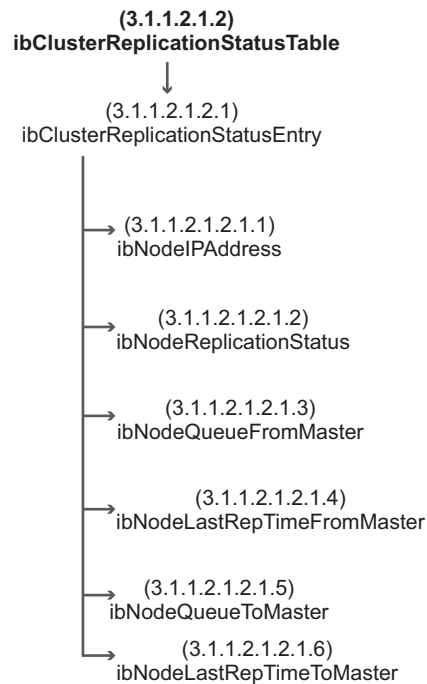
Figure 37.4 *ibPlatformOne MIB Structure*

ibClusterReplicationStatusTable

ibClusterReplicationStatusTable (object ID 3.1.1.2.1.2.1) provides information about the Grid replication status. For information about Infoblox SNMP traps, see [ibTrapDesc \(OID 3.1.1.2.11.0\)](#) on page 1070.

[Figure 37.5](#) shows the sub branches of ibClusterReplicationStatusTable.

Figure 37.5 ibClusterReplicationStatusTable Objects



[Table 37.9](#) provides information about the ibClusterReplicationStatusTable objects.

Table 37.9 ibClusterReplicationStatusTable Objects

Object (Type)	Description
ibClusterReplicationStatusEntry	A conceptual row that provides information about the Grid replication status. The status indicates whether the appliance is sending replication queues, receiving queues, or having problems with the replication.
ibNodeIPAddress (IpAddress)	IP address of a Grid member.
ibNodeReplicationStatus (String)	Replication status of the Grid member. The replication status can be one of the following: online, offline, or snapshotting.
ibNodeQueueFromMaster (Integer)	“Sent” queue size from master.
ibNodeLastRepTimeFromMaster (String)	Last sent time from master.
ibNodeQueueToMaster (Integer)	“Receive” queue size from master.
ibNodeLastRepTimeToMaster (String)	Last receive time from master.

ibNetwork Monitor

As shown in [Figure 37.6](#), the ibNetwork Monitor has one subtree, ibNetworkMonitorDNS, that branches out into the following:

- ibNetworkMonitorDNSActive (Integer) reports on whether DNS latency monitoring is enabled. This is the only object in this branch. When you send a query for this object, the appliance responds with either “active” (1) or “nonactive” (0).
- ibNetworkMonitorDNSNonAA provides information about the average latency of nonauthoritative replies to DNS queries for 1-, 5-, 15-, and 60-minute intervals. For information, see [ibNetworkMonitorDNSNonAA Objects](#) on page 1091.
- ibNetworkMonitorDNSAA provides information about the average latency of authoritative replies to DNS queries for 1-, 5-, 15-, and 60-minute intervals. For information, see [ibNetworkMonitorDNSAA Objects](#) on page 1092.
- ibNetworkMonitorDNSSecurity provides information about the invalid DNS responses that arrive on invalid ports or have invalid DNS transaction IDs. ibNetworkMonitorDNSSecurity branches out into the following:
 - ibNetworkMonitorDNSSecurityInvalidPort
 - ibNetworkMonitorDNSSecurityInvalidTxid
 - ibNetworkMonitorDNSSecurityInvalidPortOnly (Counter)
 - ibNetworkMonitorDNSSecurityInvalidPortCount (Counter)
 - ibNetworkMonitorDNSSecurityInvalidTxidOnly (Counter)
 - ibNetworkMonitorDNSSecurityInvalidTxidCount (Counter)
 - ibNetworkMonitorDNSSecurityInvalidTxidAndPort (Counter)

For information, see [Table 37.12](#) on page 1093.

Figure 37.6 ibNetWorkMonitor Objects

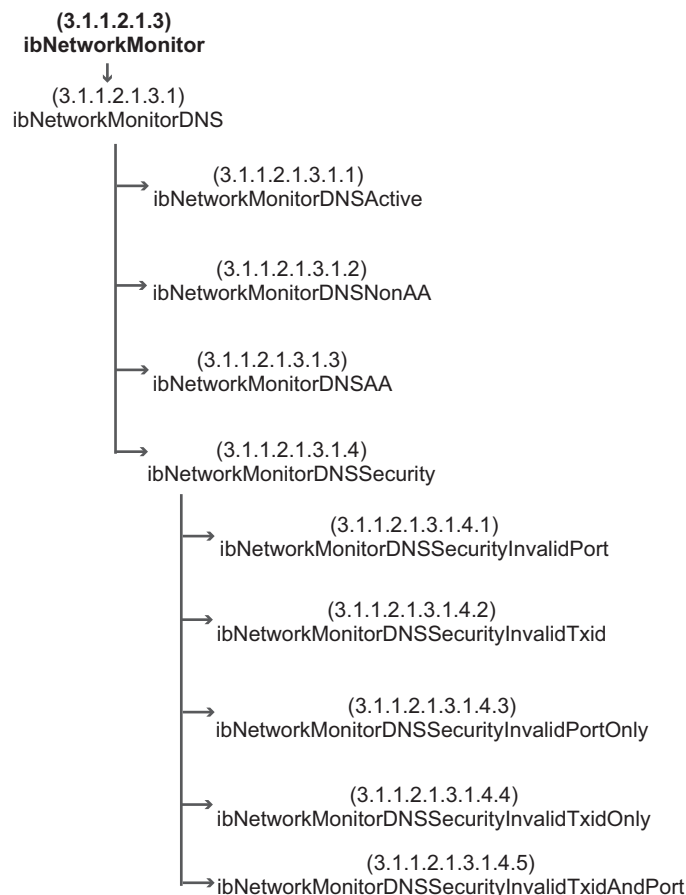
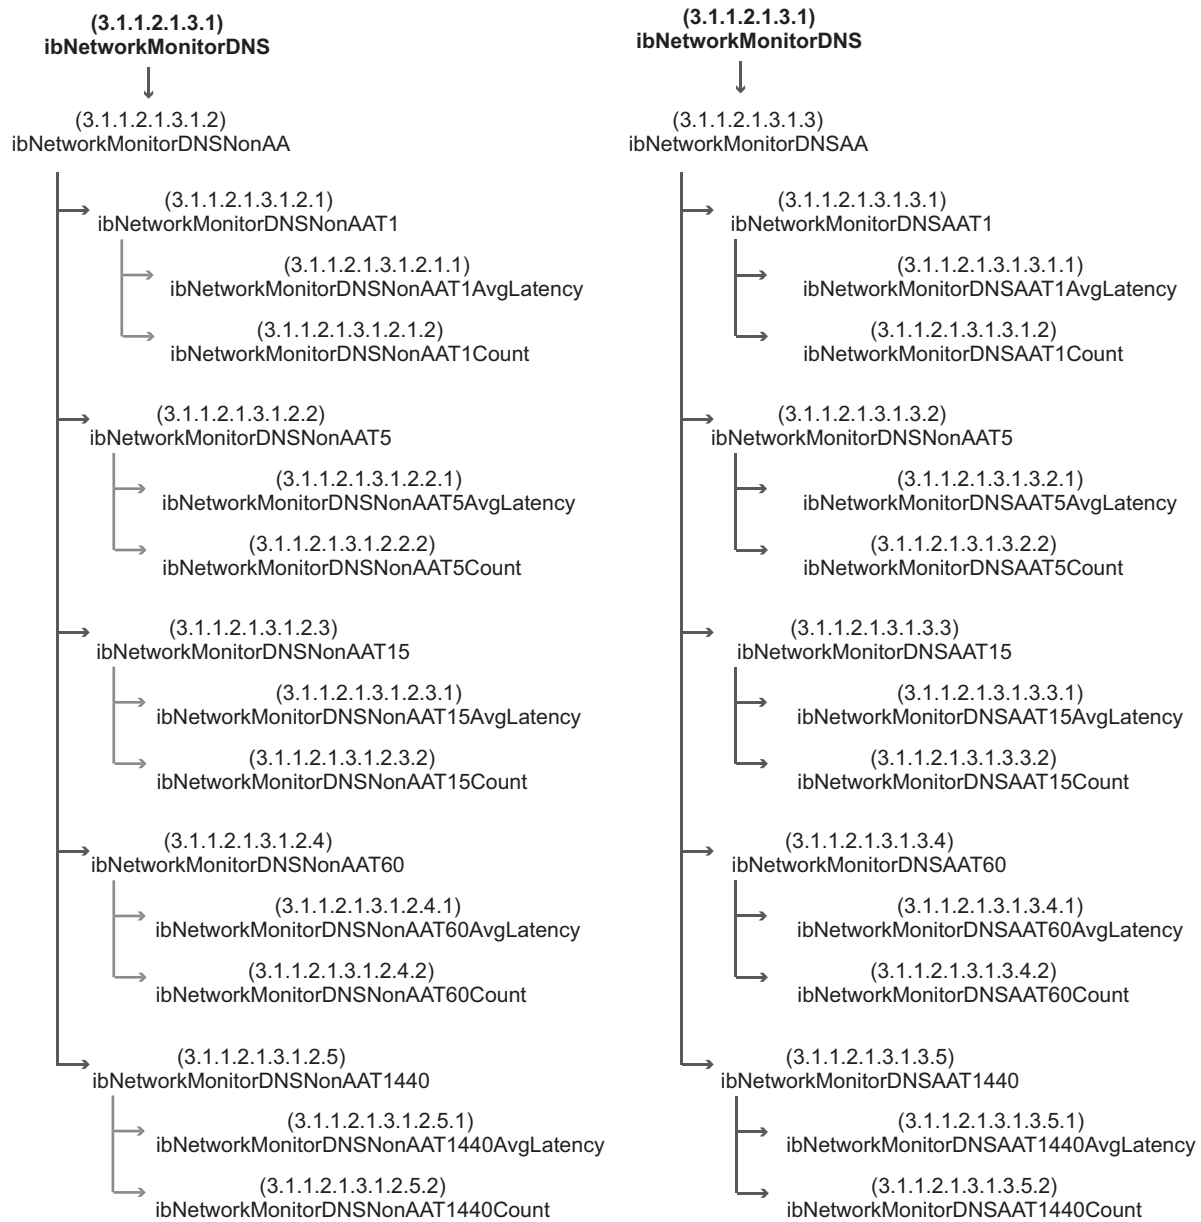


Figure 37.7 *ibNetworkMonitorDNSNonAA and ibNetworkMonitorDNSAA Subtrees*

[Table 37.10](#) describes the objects in `ibNetworkMonitorDNSNonAA`. You can send queries to retrieve values for these objects.

Table 37.10 ibNetworkMonitorDNSNonAA Objects

Object (Type)	Description
<code>ibNetworkMonitorDNSNonAAT1</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last minute.
<code>ibNetworkMonitorDNSNonAAT1AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last minute.
<code>ibNetworkMonitorDNSNonAAT1Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last minute.
<code>ibNetworkMonitorDNSNonAAT5</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last five minutes.
<code>ibNetworkMonitorDNSNonAAT5AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last five minutes.
<code>ibNetworkMonitorDNSNonAAT5Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last five minutes.
<code>ibNetworkMonitorDNSNonAAT15</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last 15 minutes.
<code>ibNetworkMonitorDNSNonAAT15AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last 15 minutes.
<code>ibNetworkMonitorDNSNonAAT15Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last 15 minutes.
<code>ibNetworkMonitorDNSNonAAT60</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last 60 minutes.
<code>ibNetworkMonitorDNSNonAAT60AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last 60 minutes.
<code>ibNetworkMonitorDNSNonAAT60Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last 60 minutes.
<code>ibNetworkMonitorDNSNonAAT1440</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last 24 hours.
<code>ibNetworkMonitorDNSNonAAT1440AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last 24 hours.
<code>ibNetworkMonitorDNSNonAAT1440Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last 24 hours.

[Table 37.11](#) describes the objects in `ibNetworkMonitorDNSAA`. You can send queries to retrieve values for these objects.

Table 37.11 ibNetworkMonitorDNSAA Objects

Object (Type)	Description
<code>ibNetworkMonitorDNSAAT1</code>	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last minute.
<code>ibNetworkMonitorDNSAAT1AvgLatency</code> (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last minute.
<code>ibNetworkMonitorDNSAAT1Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last minute.
<code>ibNetworkMonitorDNSAAT5</code>	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last five minutes.
<code>ibNetworkMonitorDNSAAT5AvgLatency</code> (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last five minutes.
<code>ibNetworkMonitorDNSAAT5Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last five minutes.
<code>ibNetworkMonitorDNSAAT15</code>	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last 15 minutes.
<code>ibNetworkMonitorDNSAAT15AvgLatency</code> (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last 15 minutes.
<code>ibNetworkMonitorDNSAAT15Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last 15 minutes.
<code>ibNetworkMonitorDNSAAT60</code>	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last 60 minutes.
<code>ibNetworkMonitorDNSAAT60AvgLatency</code> (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last 60 minutes.
<code>ibNetworkMonitorDNSAAT60Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last 60 minutes.
<code>ibNetworkMonitorDNSAAT1440</code>	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last 24 hours.
<code>ibNetworkMonitorDNSAAT1440AvgLatency</code> (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last 24 hours.
<code>ibNetworkMonitorDNSAAT1440Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last 24 hours.

[Table 37.12](#) describes the objects in `ibNetworkMonitorDNSSecurity`. You receive SNMP traps with these objects when you enable the following:

- SNMP traps
- DNS network monitoring
- DNS alert monitoring

Table 37.12 ibNetworkMonitorDNSSecurity Objects

Object (Type)	Description
<code>ibNetworkMonitorDNSSecurityInvalidPort</code>	Tracks the number of invalid DNS responses that arrive on invalid ports. For information about invalid ports, see Monitoring DNS Transactions on page 1024. This object contains a subtree with six objects that track invalid ports within a certain time interval. For information, see Table 37.13 .
<code>ibNetworkMonitorDNSSecurityInvalidTxid</code>	Tracks the number of invalid TXIDs (DNS transaction IDs). For information about invalid TXIDs, see Monitoring DNS Transactions on page 1024. This object contains a subtree with six objects that track invalid TXIDs within a certain time interval. For information, see Table 37.14 .
<code>ibNetworkMonitorDNSSecurityInvalidPortOnly (Counter)</code>	Tracks the number of DNS responses with both of the following conditions: <ul style="list-style-type: none"> • Arrive on invalid ports • Have valid TXIDs
<code>ibNetworkMonitorDNSSecurityInvalidTxidOnly (Counter)</code>	Tracks the number of DNS responses with both of the following conditions: <ul style="list-style-type: none"> • Arrive on valid ports • Have Invalid TXIDs
<code>ibNetworkMonitorDNSSecurityInvalidPortCount (Counter)</code>	Tracks the total number of invalid DNS responses that arrive on invalid ports.
<code>ibNetworkMonitorDNSSecurityInvalidTxidCount (Counter)</code>	Tracks the total number of DNS responses that have invalid DNS transaction IDs.
<code>ibNetworkMonitorDNSSecurityInvalidTxidAndPort (Counter)</code>	Tracks the number of DNS responses with both of the following conditions: <ul style="list-style-type: none"> • Arrive on invalid ports • Have invalid TXIDs

[Table 37.13](#) describes the objects in `ibNetworkMonitorDNSSecurityInvalidPort`.

Table 37.13 ibNetworkMonitorDNSSecurityInvalidPort Objects

Object (Type)	Description
<code>ibNetworkMonitorDNSSecurityInvalidPort1</code> (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last one minute.
<code>ibNetworkMonitorDNSSecurityInvalidPort5</code> (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last five minutes.
<code>ibNetworkMonitorDNSSecurityInvalidPort15</code> (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last 15 minutes.
<code>ibNetworkMonitorDNSSecurityInvalidPort60</code> (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last 60 minutes.
<code>ibNetworkMonitorDNSSecurityInvalidPort1440</code> (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last 24 hours.
<code>ibNetworkMonitorDNSSecurityInvalidPortCount</code> (Counter)	Tracks the total number of invalid DNS responses that arrive on invalid ports.

[Table 37.14](#) describes the objects in `ibNetworkMonitorDNSSecurityInvalidTxid`.

Table 37.14 ibNetworkMonitorDNSSecurityInvalidTxid Objects

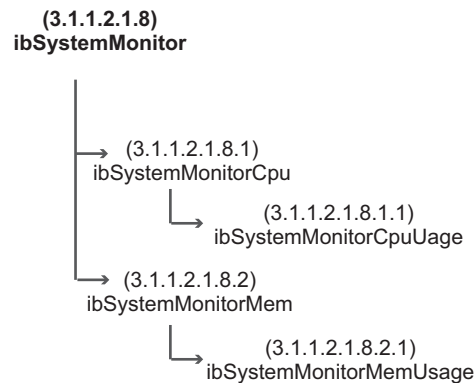
Object (Type)	Description
<code>ibNetworkMonitorDNSSecurityInvalidTxid1</code> (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last one minute.
<code>ibNetworkMonitorDNSSecurityInvalidTxid5</code> (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last five minutes.
<code>ibNetworkMonitorDNSSecurityInvalidTxid15</code> (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last 15 minutes.
<code>ibNetworkMonitorDNSSecurityInvalidTxid60</code> (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last 60 minutes.
<code>ibNetworkMonitorDNSSecurityInvalidTxid1440</code> (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last 24 hours.
<code>ibNetworkMonitorDNSSecurityInvalidTxidCount</code> (Counter)	Tracks the total number of DNS responses that have invalid DNS transaction IDs.

ibSystemMonitor

As shown in [Figure 37.4](#), ibSystemMonitor (object ID 3.1.1.2.1.2.8) has the following subtrees:

- ibSystemMonitorCpu: Contains ibSystemMonitorCpuUsage (Integer) that reports the CPU usage of the appliance.
- ibSystemMonitorMem: Contains ibSystemMonitorMemUsage (Integer) that reports the memory usage of the appliance.

Figure 37.8 *ibSystemMonitor Objects3*




ibMemberServiceStatusTable

As shown in [Figure 37.9](#), ibMemberServiceStatusTable (object ID 3.1.1.2.1.2.9) has one subtree, ibMemberServiceStatusEntry, which contains the following objects:

- ibServiceName (String) reports the names of the Infoblox services. For a list of Infoblox services, see [Infoblox Services for ibMemberServiceStatusTable](#).
- ibServiceStatus (Integer) reports the status of the Infoblox services. For a list of service status, see [Service Status](#) on page 1098.
- ibServiceDesc (String) describes the details of the status.

ibMemberServiceStatusTable displays the current status of the Infoblox services on the appliance that you query. For an HA pair, this table displays the service status of the active node. If the appliance you query is the passive node of an HA pair, this table reflects the service status of the passive node, which can be “inactive” or “unknown.”

You can also query ibMemberNodeServiceStatusTable and ibMemberPassiveNodeServiceStatusTable that display system and hardware status on the queried appliance. For information, see [ibMemberNodeServiceStatusTable](#) on page 1098 and [ibMemberPassiveNodeServiceStatusTable](#) on page 1099.

Figure 37.9 *ibMemberServiceStatusTable* Objects

Infoblox Services for *ibMemberServiceStatusTable*

[Table 37.15](#) lists the values and descriptions of the Infoblox services that appear in *ibMemberServiceStatusTable*.

Table 37.15 *ibServiceName* Values for *ibMemberServiceStatusTable*

Value	Description	Definition
1	dhcp	DHCP service
2	dns	DNS service
3	ntp	NTP service
4	tftp	File distribution using the TFTP service
5	http-file-dist	File distribution using the HTTP service
6	ftp	File distribution using the FTP service
7	bloxtools-move	Moving the bloxTools service
8	bloxtools	The bloxTools environment
9	node-status	Member status
10	disk-usage	Disk usage
11	enet-lan	LAN port
12	enet-lan2	LAN2 port
13	enet-ha	HA port
14	enet-mgmt	MGMT port
15	lcd	LCD
16	memory	Memory
17	replication	Replication service
18	db-object	Database object
19	raid-summary	RAID array

Value	Description	Definition
20	raid-disk1	RAID Disk 1 (For appliances with RAID arrays)
21	raid-disk2	RAID Disk 2 (For appliances with RAID arrays)
22	raid-disk3	RAID Disk 3 (For appliances with RAID arrays)
23	raid-disk4	RAID Disk 4 (For appliances with RAID arrays)
24	raid-disk5	RAID Disk 5 (For appliances with RAID arrays)
25	raid-disk6	RAID Disk 6 (For appliances with RAID arrays)
26	raid-disk7	RAID Disk 7 (For appliances with RAID arrays)
27	raid-disk8	RAID Disk 8 (For appliances with RAID arrays)
28	fan1	Fan 1
29	fan2	Fan 2
30	fan3	Fan 3
31	fan4	Fan 4
32	fan5	Fan 5
33	fan6	Fan 6
34	fan7	Fan 7
35	fan8	Fan 8
36	power-supply1	Power supply 1
37	power-supply2	Power supply 2
38	ntp-sync	NTP synchronization
39	cpu1-temp	CPU temperature
40	cpu2-temp	CPU temperature
41	sys-temp	System temperature
42	raid-battery	RAID battery
43	cpu-usage	CPU usage
44	ospf	OSPF
45	bgp	BGP
46	mgm-service	Multi-Grid management
47	subGrid-conn	Grid in Master Grid
48	network-capacity	Network capacity
49	reporting	Reporting service
50	dns-cache-acceleration	DNS Cache Acceleration services
51	ospf6	OSPF6

Service Status

When you query the service status on an appliance, the response includes the status of the services. [Table 37.16](#) shows the values and descriptions of the status. Note that for internal Grid operations, the NTP service is always in the “working” state even if it has been disabled through the Infoblox GUI.

Table 37.16 ibServiceStates Values

Value	Description	Definition
1	working	The service is functioning properly.
2	warning	The service is having some issues. Check the service or hardware function and the syslog to identify the problem.
3	failed	The service failed. Review the syslog to identify the problem.
4	inactive	The service is disabled or out of service.
5	unknown	The appliance cannot detect the current status of the service.


ibMemberNodeServiceStatusTable

As shown in [Figure 37.10](#), `ibMemberNodeServiceStatusTable` (object ID 3.1.1.2.1.10) has one subtree, `ibMemberNodeServiceStatusEntry`, which contains the following objects:

- `ibMemberNodeServiceName` (String) reports the names of the system and hardware services. For a list of service names, see [Infoblox Services for ibMemberServiceStatusTable](#) on page 1096.
- `ibMemberNodeServiceStatus` (Integer) reports the status of the services. For a list of service status, see [Service Status](#) on page 1098.
- `ibMemberNodeServiceDesc` (String) describes the details of the status.

`ibMemberNodeServiceStatusTable` displays the current status of the system and hardware services on the appliance that you query. For example, when you query an independent appliance, this table shows the information about the independent appliance. When you query the VIP of an HA pair, this table shows the information about the active node. For the active node of the HA pair, you can also query `ibMemberPassiveNodeStatusTable` to get the status of the passive node. For information, see [ibMemberPassiveNodeServiceStatusTable](#) on page 1099.

Note: For an independent appliance and the passive node of an HA pair, no information is returned when you query `ibMemberPassiveNodeServiceStatusTable`.

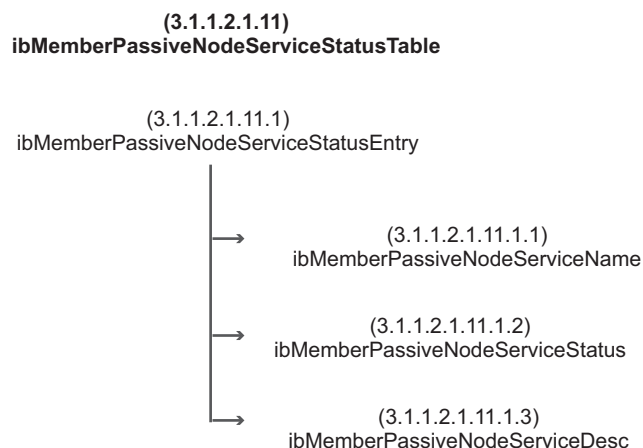
Figure 37.10 *ibMemberNodeServiceStatusTable Objects*

ibMemberPassiveNodeServiceStatusTable

As shown in [Figure 37.11](#), `ibMemberPassiveNodeServiceStatusTable` (object ID 3.1.1.2.1.2.11) has one subtree, `ibMemberPassiveNodeServiceStatusEntry`, which contains the following objects:

- `ibMemberPassiveNodeServiceName` (String) reports the names of the system and hardware services. For a list of service names, see [Infoblox Services for ibMemberServiceStatusTable](#) on page 1096.
- `ibMemberPassiveNodeServiceStatus` (Integer) reports the status of the services. For a list of possible service status, see [Service Status](#) on page 1098.
- `ibMemberPassiveNodeServiceDesc` (String) describes details of the status.

`ibMemberPassiveNodeServiceStatusTable` displays the current status of the system and hardware services on the passive node of an HA pair when you query the VIP of the HA pair. For independent appliances and the passive nodes of HA pairs, this table does not display any status.

Figure 37.11 *ibMemberPassiveNodeServiceStatusTable Objects*

ibDHCPOne MIB

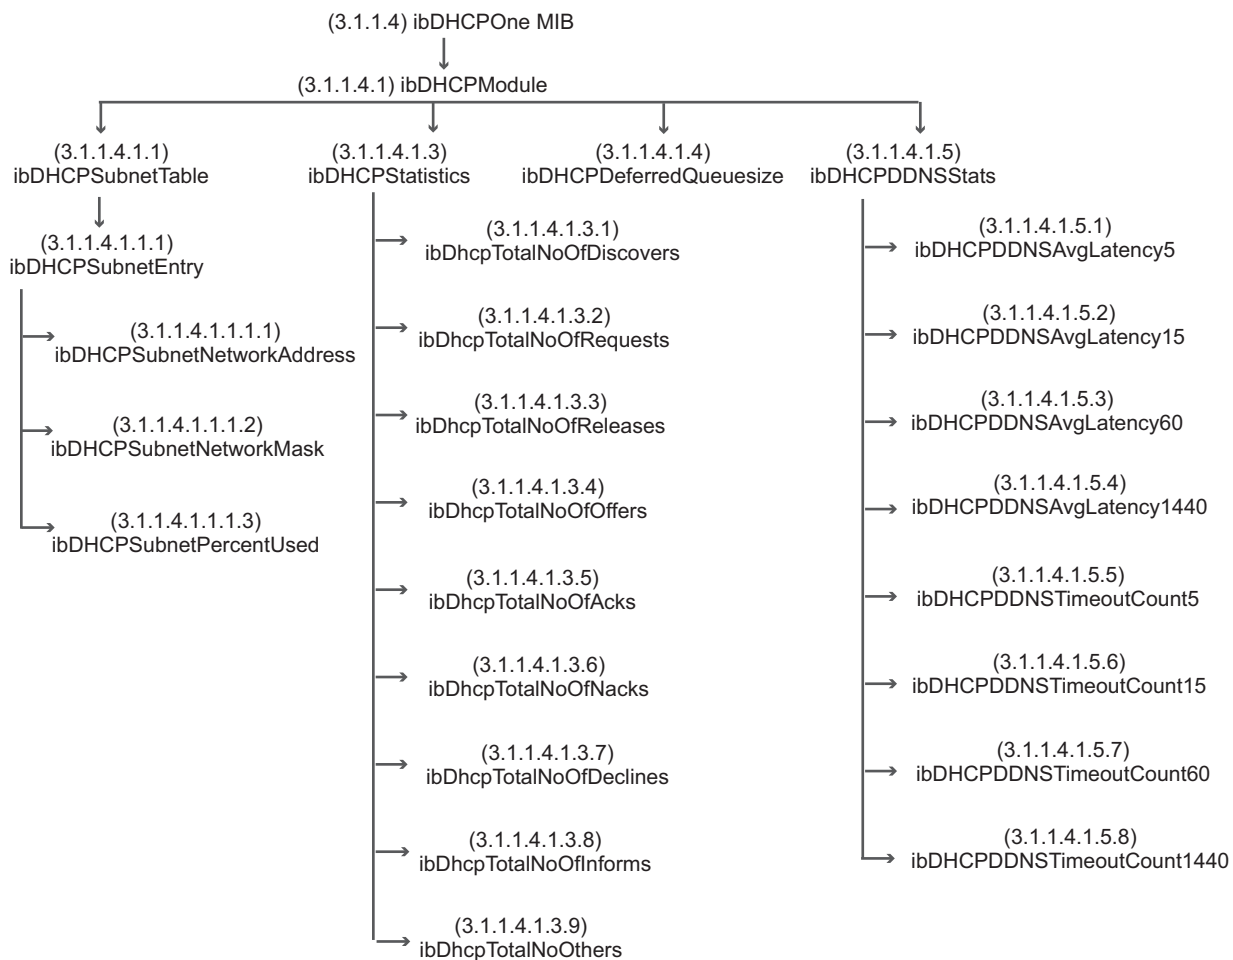
The ibDHCPOne MIB provides information about address usage within a subnet, DHCP lease statistics, and DHCP packet counts. It includes two modules, ibDHCPModule for IPv4 data and ibDHCPv6Module for IPv6 data.

ibDHCPModule

Figure 37.12 illustrates the structure of the ibDHCPModule. (Note that the OIDs shown in the illustration do not include the prefix .1.3.6.1.4.1.7779.) ibDHCPModule contains the following objects:

- ibDHCPSubnetTable provides statistical data about the DHCP operations of the appliance. For information, see [ibDHCPSubnetTable](#) on page 1101.
- ibDHCPStatistics maintains counters for different types of packets. For information, see [ibDHCPStatistics](#) on page 1102.
- ibDHCPDeferredQueueSize tracks the total number of deferred DDNS updates that are currently in the queue to be retried. When DDNS updates are deferred due to timeout or server issues, the DHCP server puts these updates in this queue.
- ibDHCPDDNSStats monitors the average latency for the DDNS updates in microseconds and the number of timeouts during different time intervals. For information, see [ibDHCPDDNSStats](#) on page 1103.

Figure 37.12 ibDHCPModule



ibDHCPSubnetTable

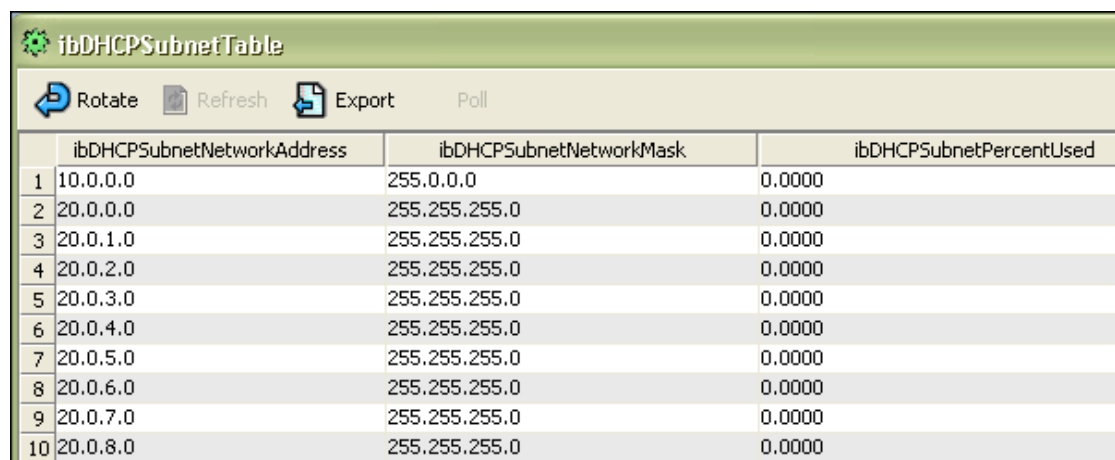
ibDHCPSubnetTable provides statistical data about the DHCP operations of the appliance. It contains the following objects:

Table 37.17 ibDHCPSubnetTable

Object (Type)	Description
ibDHCPSubnet Entry	File that contains the objects for monitoring DHCP operations on the appliance.
ibDHCPSubnetNetworkAddress (IbIpAddr)	The subnetworks, in IP address format, that have IP addresses for lease. A subnetwork may have many address ranges for lease.
ibDHCPSubnetNetworkMask (IbIpAddr)	The subnet mask in dotted decimal format.
ibDHCPSubnetPercentUsed (Integer)	The percentage of dynamic DHCP addresses leased out at this time for each subnet. Fixed addresses are always counted as leased for this calculation, if the fixed addresses are within a leased address range.

Following is an example of the table as viewed through a MIB browser:

Figure 37.13 MIB Browser View 1



	ibDHCPSubnetNetworkAddress	ibDHCPSubnetNetworkMask	ibDHCPSubnetPercentUsed
1	10.0.0.0	255.0.0.0	0.0000
2	20.0.0.0	255.255.255.0	0.0000
3	20.0.1.0	255.255.255.0	0.0000
4	20.0.2.0	255.255.255.0	0.0000
5	20.0.3.0	255.255.255.0	0.0000
6	20.0.4.0	255.255.255.0	0.0000
7	20.0.5.0	255.255.255.0	0.0000
8	20.0.6.0	255.255.255.0	0.0000
9	20.0.7.0	255.255.255.0	0.0000
10	20.0.8.0	255.255.255.0	0.0000

ibDHCPStatistics

ibDHCPStatistics maintains counters for different types of packets. The counters always start with zero when the DHCP service is restarted. Therefore, the numbers reflect the total number of packets received since the DHCP service was last restarted on the appliance. The ibDHCPStatistics module contains the following objects:

Table 37.18 ibDHCPStatistics

Object (Type)	Description
ibDhcpTotalNoOfDiscovers (Counter)	The number of DHCPDISCOVER messages that the appliance received. Clients broadcast DHCPDISCOVER messages when they need an IP address and network configuration information.
ibDhcpTotalNoOfRequests (Counter)	The number of DHCPREQUEST messages that the appliance received. A client sends a DHCPREQUEST message requesting configuration information, after it receives the DHCPOFFER message.
ibDhcpTotalNoOfReleases (Counter)	The number of DHCPRELEASE messages that the appliance received from its clients. A client sends a DHCP release when it terminates its lease on an IP address.
ibDhcpTotalNoOfOffers (Counter)	The number of DHCPOFFER messages that the appliance has sent to clients. The appliance sends a DHCPOFFER message to a client. It contains an IP address and configuration information.
ibDhcpTotalNoOfAcks (Counter)	The number of DHCPACK messages that the appliance sent to clients. It sends a DHCPACK message to a client to confirm that the IP address offered is still available.
ibDhcpTotalNoOfNacks (Counter)	The number of DHCPNACK messages that the appliance sent to clients. It sends a DHCPNACK message to withdraw its offer of an IP address.
ibDhcpTotalNoOfDeclines (Counter)	The number of DHCPDECLINE messages that the appliance received. A client sends a DHCPDECLINE message if it determines that an offered IP address is already in use.
ibDhcpTotalNoOfInforms (Counter)	The number of DHCPINFORM messages that the appliance received. A client sends a DHCPINFORM message when it has an IP address but needs information about the network.
ibDhcpTotalNoOfOthers (Counter)	The total number of DHCP messages other than those used in negotiation, such as DHCPFORCERENEW, DHCPKNOWN, and DHCPLEASEQUERY.

ibDHCPDDNSStats

ibDHCPDDNSStats monitors the average latency for the DHCP DDNS updates in microseconds and the number of timeouts during different time intervals. The ibDHCPDDNSStats module contains the following objects:

Table 37.19 ibDHCPStatistics

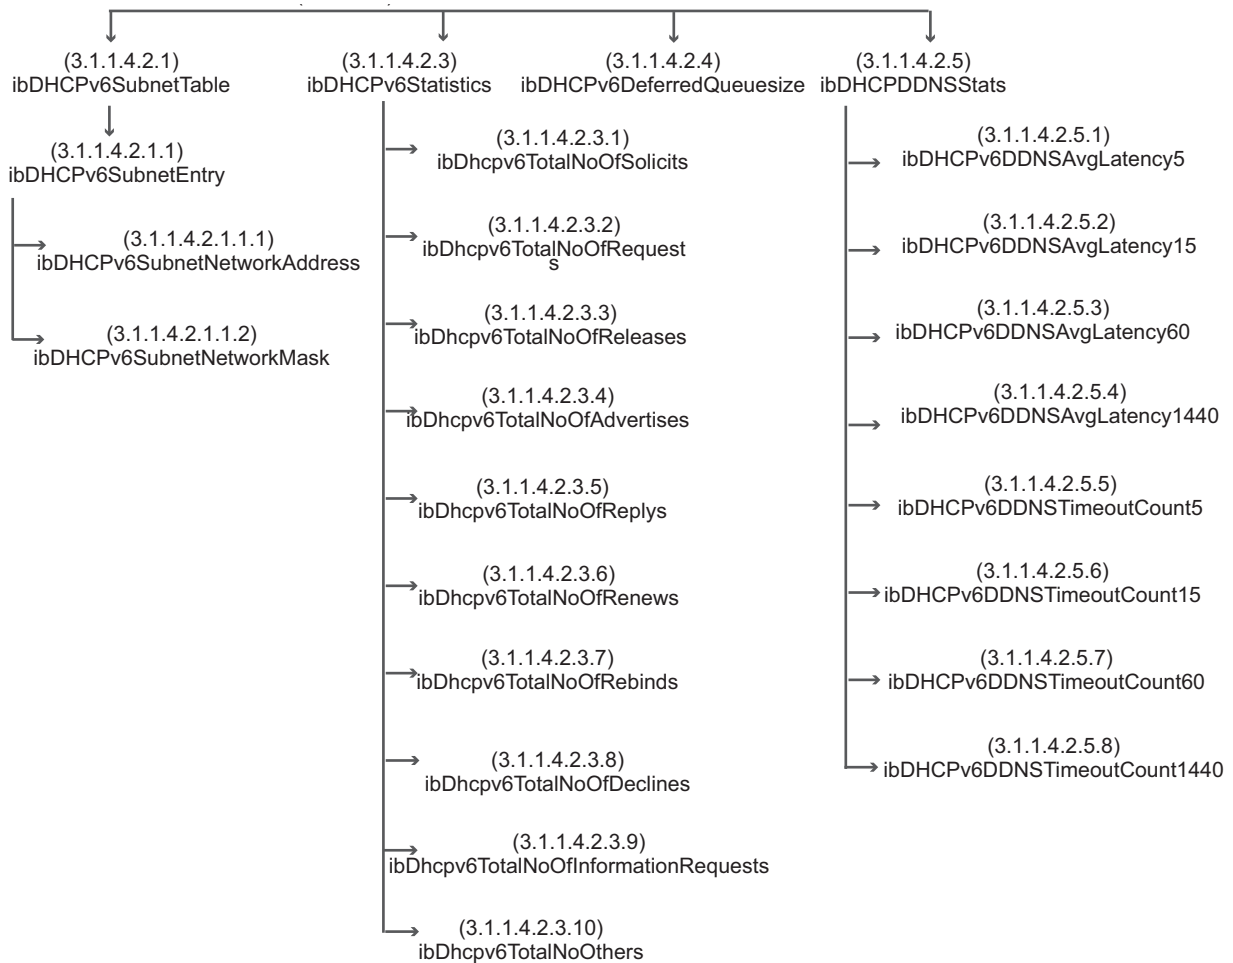
Object (Type)	Description
ibDHCPDDNSAvgLatency5 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last five minutes.
ibDHCPDDNSAvgLatency15 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last 15 minutes.
ibDHCPDDNSAvgLatency60 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last 60 minutes.
ibDHCPDDNSAvgLatency1440 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last 24 hours.
ibDHCPDDNSTimeoutCount5 (Integer)	The number of timeouts for the DHCP DDNS updates in the last five minutes.
ibDHCPDDNSTimeoutCount15 (Integer)	The number of timeouts for the DHCP DDNS updates in the last 15 minutes.
ibDHCPDDNSTimeoutCount60 (Integer)	The number of timeouts for the DHCP DDNS updates in the last 60 minutes.
ibDHCPDDNSTimeoutCount1440 (Integer)	The number of timeouts for the DHCP DDNS updates in the last 24 hours.

ibDHCpV6Module

Figure 37.14 illustrates the structure of the ibDHCpV6Module, which contains the following objects:

- ibDHCpV6SubnetTable provides statistical data about the DHCPv6 operations of the appliance. For information, see [ibDHCpV6SubnetTable](#) on page 1105.
- ibDHCpV6Statistics maintains counters for different types of packets. For information, see [ibDHCpV6Statistics](#) on page 1105.
- ibDHCpV6DeferredQueueSize tracks the total number of deferred DDNS updates that are currently in the queue to be retried. When DDNS updates are deferred due to timeout or server issues, the DHCP server puts these updates in this queue.
- ibDHCpV6DDNSStats monitors the average latency for the DDNS updates in microseconds and the number of timeouts during different time intervals. For information, see [ibDHCpV6DDNSStats](#) on page 1106.

Figure 37.14 *ibDHCpV6Module*



ibDHCPv6SubnetTable

ibDHCPSubnetTable provides statistical data about the DHCPv6 operations of the appliance. It contains the following objects:

Table 37.20 ibDHCPSubnetTable

Object (Type)	Description
ibDHCPv6Subnet Entry	File that contains the objects for monitoring DHCPv6 operations on the appliance.
ibDHCPv6SubnetNetworkAddress (IbIpAddr)	The subnetworks, in IPv6 address format, that have IPv6 addresses for lease. A subnetwork may have many address ranges for lease.
ibDHCPv6SubnetNetworkMask (IbIpAddr)	The subnet mask in CIDR notation format.

ibDHCPv6Statistics

ibDHCPv6Statistics maintains counters for different types of packets. The counters always start with zero when the DHCP service is restarted. Therefore, the numbers reflect the total number of packets received since the DHCP service was last restarted on the appliance. The ibDHCPv6Statistics module contains the following objects:

Table 37.21 ibDHCPv6Statistics

Object (Type)	Description
ibDhcpv6TotalNoOfSolicits (Counter)	The number of Solicit messages that the Grid member received, including Solicit messages embedded in Relay-Forward messages. A DHCP client sends a Solicit message to locate DHCP servers.
ibDhcpv6TotalNoOfRequests (Counter)	The number of Request messages that the Grid member received. A DHCP client sends a Request message to request one or more IP addresses and configuration parameters from a DHCP server.
ibDhcpv6TotalNoOfReleases (Counter)	The number of Release messages that the Grid member received. A DHCP client sends a Release message when it terminates its lease and releases its IP address.
ibDhcpv6TotalNoOfAdvertises (Counter)	The number of Advertise messages that the Grid member sent. When a DHCP server receives a Solicit message, it can respond with an Advertise message to indicate that the server is available for DHCP service.
ibDhcpv6TotalNoOfReplies (Counter)	The number of Reply messages that the Grid member sent. A DHCP server sends a Reply message that includes IP addresses and configuration parameters when it responds to Solicit, Request, Renew or Rebind message. It sends a Reply message with configuration parameters only when it responds to an Information-Request message.
ibDhcpv6TotalNoOfRenews (Counter)	The number of Renew messages that the Grid member received. A DHCP client sends a Renew message to a DHCP server to extend the lifetimes on the leases granted by the DHCP server and to update other properties.

Object (Type)	Description
ibDhcpv6TotalNoOfRebinds (Counter)	The number of Rebind messages that the Grid member received. A DHCP client sends a Rebind message to extend the lifetime of its lease and to update configuration parameters.
ibDhcpv6TotalNoOfDeclines (Counter)	The number of Decline messages that the Grid member received. A DHCP client sends a Decline message to a DHCP server when it discovers that the IP address offered by a DHCP server is already in use.
ibDhcpv6TotalNoOfInformationRequests (Counter)	The number of Information-Request messages that the Grid member received. A client sends an Information-Request message to retrieve configuration parameters, such as the IP addresses of DNS servers in the network.
ibDhcpv6TotalNoOfOthers (Counter)	The total number of DHCP messages other than those used in negotiation.

ibDHCPv6DDNSStats

ibDHCPv6DDNSStats monitors the average latency for the DHCPv6 DDNS updates in microseconds and the number of timeouts during different time intervals. The ibDHCPv6DDNSStats module contains the following objects:

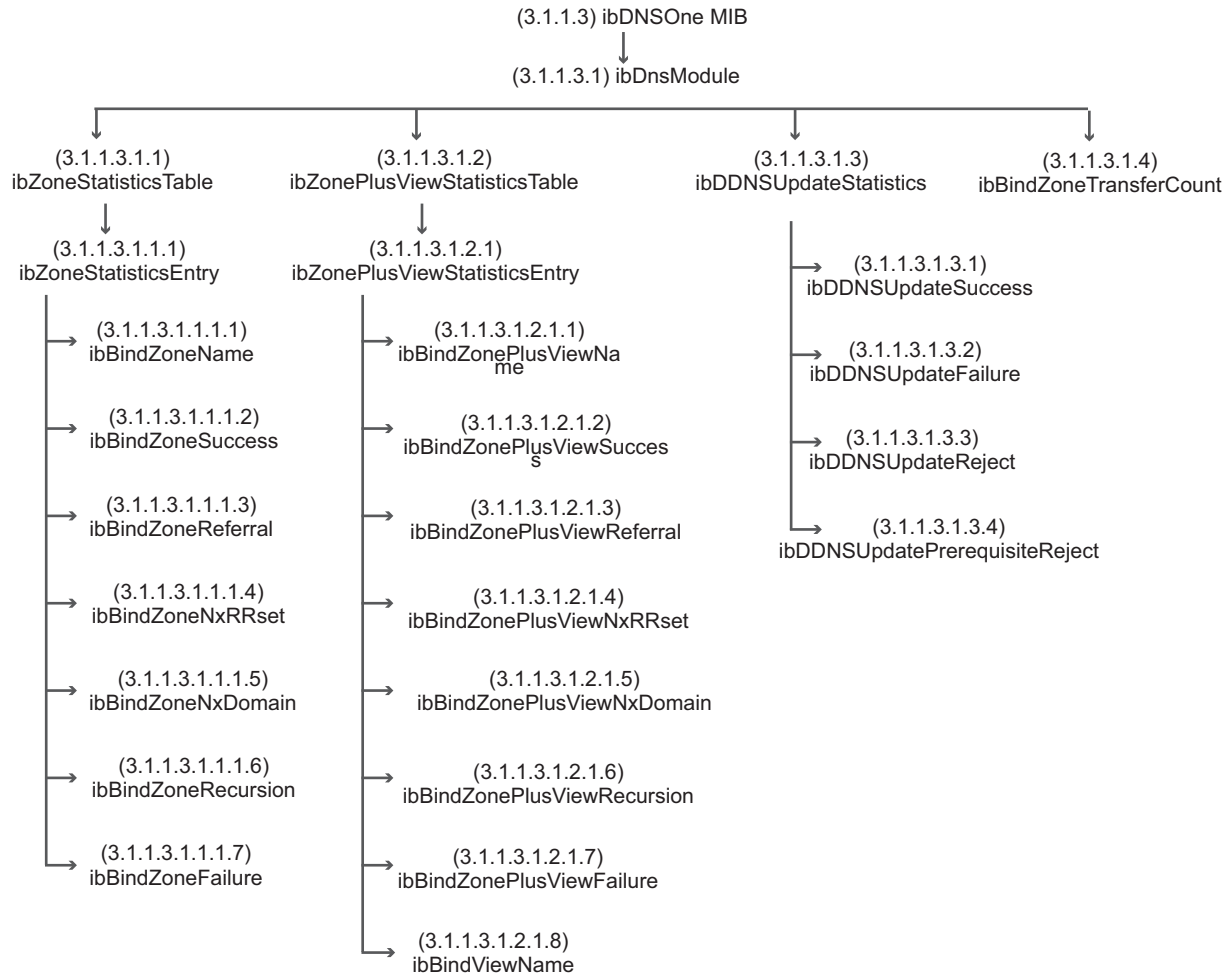
Table 37.22 ibDHCPStatistics

Object (Type)	Description
ibDHCPv6DDNSAvgLatency5 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last five minutes.
ibDHCPv6DDNSAvgLatency15 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last 15 minutes.
ibDHCPv6DDNSAvgLatency60 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last 60 minutes.
ibDHCPv6DDNSAvgLatency1440 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last 24 hours.
ibDHCPv6DDNSTimeoutCount5 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last five minutes.
ibDHCPv6DDNSTimeoutCount15 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last 15 minutes.
ibDHCPv6DDNSTimeoutCount60 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last 60 minutes.
ibDHCPv6DDNSTimeoutCount1440 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last 24 hours.

ibDNSOne MIB

The ibDNSOne MIB provides DNS statistics about all zones in all views. [Figure 37.15](#) illustrates the structure of the ibDNSOne MIB. (Note that the OIDs shown in the illustration do not include the prefix 1.3.6.1.4.1.7779.) The ibDNSOne MIB contains four subtrees: ibZoneStatisticsTable (Counter64), ibZonePlusViewStatisticsTable (Counter64), ibDDNSUpdateStatistics (Counter64), and ibBindZoneTransferCount (Counter64).

Figure 37.15 *ibDNSOne MIB*



Using the DNS Zone Statistics Tables

ibZoneStatisticsTable and ibZonePlusViewStatisticsTable provide DNS statistics for all zones in all DNS views, including the default and all user-defined DNS views. You can use the information in these tables to calculate the total number of recursive queries on the DNS server. Depending on whether your DNS server is an authoritative or a caching-only server, you calculate the total number of recursive queries differently. For information, see [Calculating Recursive DNS Queries](#) on page 1110.

ibZoneStatisticsTable

ibZoneStatisticsTable contains DNS statistics of all zones in the default DNS view. DNS statistics of user-defined DNS views are captured in ibZonePlusViewStatisticsTable. For information, see [ibZonePlusViewStatisticsTable](#) on page 1109.

ibZoneStatisticsTable includes a “summary” zone that provides global statistics for the DNS server, including statistics for all zones in the default and user-defined DNS views.

The syntax of the objects in ibZoneStatisticsTable uses a Counter64 format. In some cases, the counter format may not be compatible with SNMP toolkits that use a 32-bit counter. Ensure that you reconfigure or update these tools to use the Counter64 format. ibZoneStatisticsTable contains the following objects:

Table 37.23 *ibZoneStatisticsTable*

Object (Type)	Description
ibBindZoneName (IbString)	DNS zone name. The index name for global statistics is “summary.”
ibBindZoneSuccess (Counter64)	The number of successful responses since the DNS process started.
ibBindZoneReferral (Counter64)	The number of DNS referrals since the DNS process started.
ibBindZoneNxRRset (Counter64)	The number of DNS queries received for non-existent records.
ibBindZoneNxDomain (Counter64)	The number of DNS queries received for non-existent domains.
ibBindZoneRecursion (Counter64)	The number recursive queries received since the DNS process started.
ibBindZoneFailure (Counter64)	The number of failed queries since the DNS process started.

Following is an example of the table as viewed through a MIB browser:

Figure 37.16 MIB Browser View

ibZoneStatisticsTable						
Rotate Refresh Export Poll						
	ibBindZoneName	ibBindZoneSuccess	ibBindZoneReferral	ibBindZoneNxRRset	ibBindZoneNxDomain	ibBindZoneRecursion
1	parent	0	0	0	0	0
2	digzone	0	0	0	0	0
3	summary	1	3	0	0	0
4	ddnszone	0	0	0	0	0
5	reboottestzone	1	0	0	0	0
6	10.in-addr.arpa	0	0	0	0	0
7	20.in-addr.arpa	0	0	0	0	0
8	40.in-addr.arpa	0	0	0	0	0
9	zone-1-1-a.parent	0	0	0	0	0
10	zone-1-2-a.parent	0	0	0	0	0
11	zone-1-3-a.parent	0	0	0	0	0
12	zone-1-4-a.parent	0	0	0	0	0
13	zone-1-5-a.parent	0	0	0	0	0
14	zone-1-6-a.parent	0	0	0	0	0
15	zone-1-7-a.parent	0	0	0	0	0
16	zone-1-8-a.parent	0	0	0	0	0
17	zone-1-9-a.parent	0	0	0	0	0
18	zone-1-10-a.parent	0	0	0	0	0

ibZonePlusViewStatisticsTable

ibZonePlusViewStatisticsTable provides DNS statistics about all zones in user-defined DNS views. DNS statistics about zones in the default view are captured in ibZoneStatisticsTable. Note that information in ibZonePlusViewStatisticsTable is rolled up to the “summary” zone in ibZoneStatisticsTable. For information, see [ibZoneStatisticsTable](#) on page 1108.

The syntax of the objects in ibZonePlusViewStatisticsTable uses a Counter64 format. In some cases, the counter format may not be compatible with SNMP toolkits that use a 32-bit counter. Ensure that you reconfigure or update these tools to use the Counter64 format. ibZonePlusViewStatisticsTable contains the following objects:

Table 37.24 ibZonePlusViewStatisticsTable

Object (Type)	Description
ibBindZonePlusViewName (lString)	The zone name.
ibBindZonePlusViewSuccess (Counter64)	The number of successful responses since the DNS process started.
ibBindZonePlusViewReferral (Counter64)	The number of DNS referrals since the DNS process started.
ibBindZonePlusViewNxRRset (Counter64)	The number of DNS queries received for non-existent records.
ibBindZonePlusViewNxDomain (Counter64)	The number of DNS queries received for non-existent domains.

Object (Type)	Description
ibBindZonePlusViewRecursion (Counter64)	The number of recursive queries received since the DNS process started.
ibBindZonePlusViewFailure (Counter64)	The number of failed queries since the DNS process started.
ibBindViewName (lString)	The DNS view name.

Calculating Recursive DNS Queries

You can use the information in `ibZoneStatisticsTable` and `ibZonePlusViewStatisticsTable` to calculate the total number of recursive queries.

Following is an example of `ibZoneStatisticsTable` indexed by zone names in the default view:

index	ibBindZoneName ibBindZoneNxDomain	ibBindZoneSuccess ibBindZoneRecursion	ibBindZoneReferral ibBindZoneFailure	ibBindZoneNxRRset
=====	=====	=====	=====	=====
"abc.com"	abc.com	0	0	0
	0	0	0	
"summary"	summary	5	0	0
	0	0	0	
"internal.com"	internal.com	1	0	0
	0	0	0	

Following is an example of `ibZonePlusViewStatisticsTable` indexed by zone names in all user-defined views:

index	ibBindZonePlusViewName ibBindZonePlusViewNxDomain	ibBindZonePlusViewSuccess ibBindZonePlusViewRecursion	ibBindZonePlusViewReferral ibBindZonePlusViewFailure	ibBindZonePlusViewNxRRset ibBindViewName
=====	=====	=====	=====	=====
"ext1.com"	ext1.com	1	0	0
	0	0	0	DNS1
"ext2.com"	ext2.com	2	0	0
	0	0	0	DNS1
"ext3.com"	ext3.com	0	0	0
	0	0	0	DNS2

Use the `ibBindZoneSuccess` object in both tables to determine the total number of recursive queries. If your DNS server is a caching-only server, the total number of recursive queries is the number indicated in the `ibBindZoneSuccess` object of the "summary" zone. In this example, for a caching-only server, the total number of recursive queries is 5.

If your DNS server is an authoritative server, add all the numbers in `ibBindZoneSuccess` for all zones in both tables, excluding the "summary" zone. In this example, the total is 4. You then subtract this number from the number in `ibBindZoneSuccess` of the "summary" zone. In this case, the total number of recursive queries is 1 for an authoritative DNS server.

ibDDNSUpdateStatistics

ibDDNSUpdateStatistics provides statistical data about DDNS updates. The counters always start with zero when the DNS service is restarted. They report the total numbers since the DNS service was last restarted.

ibDDNSUpdateStatistics contains the following objects:

Table 37.25 ibDDNSUpdateStatistics

Object (Type)	Description
ibDDNSUpdateSuccess (Counter64)	The number of successful dynamic DNS updates.
ibDDNSUpdateFailure (Counter64)	The number of all failed dynamic DNS updates, excluding those reported by the ibDDNSUpdateReject object.
ibDDNSUpdateReject (Counter64)	The number of dynamic DNS updates that failed because they were denied by the DNS server.
ibDDNSUpdatePrerequisiteReject (Counter64)	The number of dynamic DNS updates that failed because the prerequisites were not satisfied. This is also included in the total number of failures reported by the ibDDNSUpdateFailure object.

ibBindZoneTransferCount

ibBindZoneTransferCount (Counter64) provides the total number of successful zone transfers from an Infoblox primary or secondary DNS server to a DNS client, since the DNS service was last restarted. Note that this counter tracks the number of successful full zone transfers (AXFRs) and incremental zone transfers (IXFRs).

IB-DNSSERV-MIB

The IB-DNSSERV-MIB contains one object, ibDnsServConfig, which reports the DNS BIND version implemented by the NIOS software.

IB-DNSHITRATIO-MIB

The IB-DNSHITRATIO-MIB contains one object, ibDnsHitRatio, which provides information about the DNS cache hit ratio. Note that the MIB variable (ibDNSOne) for cache hit rate has an OID of 1.3.6.1.4.1.7779.3.1.1.3.1.5, where 1.3.6.1.4.1.7779 is the prefix.

IB-DNSQUERYRATE-MIB

The IB-DNSQUERYRATE-MIB contains one object, ibDnsQueryRate, which provides information about the DNS queries per second. Note that the MIB variable (ibDNSOne) for DNS query rate has an OID of 1.3.6.1.4.1.7779.3.1.1.3.1.6, where 1.3.6.1.4.1.7779 is the prefix.

IB-DHCPSErv-MIB

The IB-DHCPSErv-MIB contains one object, ibDhcpv4ServerSystemDescr, which provides the DHCP server name and its DHCP version.



Chapter 38 Infoblox Reporting Solution

This chapter describes the Infoblox reporting solution and its features. It explains how to view predefined reports and create user-defined reports and searches. It also provides best practices for customizing searches. It contains the following sections:

- [*Infoblox Reporting Solution*](#) on page 1116
 - [*Supported Platforms for Reporting*](#) on page 1119
 - [*Infoblox-4030 Supported Reports*](#) on page 1120
- [*Grid Reporting Properties*](#) on page 1121
 - [*Configuring General Grid Reporting Properties*](#) on page 1121
 - [*Setting Network Port for Reporting*](#) on page 1122
 - [*Setting Email Properties*](#) on page 1122
 - [*Defining PDF Settings*](#) on page 1122
 - [*Defining DNS Query Settings*](#) on page 1122
 - [*Configuring the Capture of DNS Queries and Responses*](#) on page 1123
- [*Scheduling Report Deliveries*](#) on page 1129
- [*About Reports*](#) on page 1130
 - [*Adding New Reports*](#) on page 1131
 - [*Adding New Panels to Reports*](#) on page 1131
 - [*Cloning Predefined Reports*](#) on page 1132
 - [*Modifying User-Defined Reports*](#) on page 1132
 - [*Deleting User-Defined Reports*](#) on page 1132
- [*About Searches*](#) on page 1132
 - [*Guidelines for Customizing Searches*](#) on page 1133
 - [*Reporting Indexes and Update Time Intervals*](#) on page 1134
 - [*Cloning Searches*](#) on page 1137
 - [*Modifying Searches*](#) on page 1138
 - [*Scheduling Searches*](#) on page 1138
 - [*Exporting Searches*](#) on page 1139
 - [*Scheduling Exports of Search Results*](#) on page 1139
- [*About Alerts*](#) on page 1140
 - [*Alerting Logic*](#) on page 1140
 - [*Defining Alerts*](#) on page 1141

- [Previewing Alerting Logic](#) on page 1141
- [Defining Advanced Alert Settings](#) on page 1142
- [About IP Blocks and IP Block Groups](#) on page 1142
 - [Adding IP Block Groups](#) on page 1143
 - [Modifying IP Block Groups](#) on page 1143
 - [Deleting IP Block Groups and IP Blocks](#) on page 1144
 - [Exporting IP Block Groups and IP Blocks](#) on page 1144
 - [Printing IP Block Groups and IP Blocks](#) on page 1144
- [Predefined Reports](#) on page 1145
 - [Predefined Report Categories](#) on page 1145
 - [Changing Report Formats](#) on page 1147
 - [DDNS Update Rate Trend](#) on page 1147
 - [DHCP Lease History](#) on page 1148
 - [DHCP Top Lease Clients](#) on page 1148
 - [Top Devices Identified](#) on page 1149
 - [Device Trend](#) on page 1150
 - [Device Class Trend](#) on page 1150
 - [Top Device Classes](#) on page 1151
 - [Top Devices Denied an IP Address](#) on page 1151
 - [Device Fingerprint Change Detected](#) on page 1152
 - [DHCPv4 Usage Statistics](#) on page 1153
 - [DHCPv4 Range Utilization Trend](#) on page 1154
 - [DHCPv4 Usage Trend](#) on page 1154
 - [DHCP Message Rate Trend](#) on page 1154
 - [DNS Query Rate by Query Type](#) on page 1155
 - [DNS Query Rate by Server](#) on page 1155
 - [DNS Daily Query Rate by Server](#) on page 1155
 - [DNS Daily Peak Hour Query Rate by Server](#) on page 1156
 - [DNS Response Latency Trend](#) on page 1156
 - [DNS Top NXDOMAIN / NOERROR \(no data\)](#) on page 1156
 - [DNS Top Clients Per Domain](#) on page 1157
 - [DNS Top SERVFAIL Errors Received](#) on page 1157
 - [DNS Top SERVFAIL Errors Sent](#) on page 1158
 - [DNS Top Timeout Recursive Queries](#) on page 1158
 - [DNS Top Requested Domain Names](#) on page 1159
 - [DNS Cache Hit Rate Trend](#) on page 1159
 - [DNS Top RPZ Hits](#) on page 1166
 - [DNS Top RPZ Hits by Clients](#) on page 1167
 - [FireEye Alerts](#) on page 1168
 - [Threat Protection Event Count By Severity Trend](#) on page 1169
 - [Threat Protection Event Count By Member Trend](#) on page 1169
 - [Threat Protection Event Count By Rule](#) on page 1170
 - [Threat Protection Event Count By Time](#) on page 1170
 - [Threat Protection Event Count By Category](#) on page 1171
 - [Threat Protection Event Count By Member](#) on page 1171

-
- [*DNS Top Clients*](#) on page 1159
 - [*DNS Replies Trend*](#) on page 1159
 - [*DNS Zones Last Queried*](#) on page 1160
 - [*DNS Resource Records Last Queried*](#) on page 1160
 - [*DNS Query Trend per IP Block Group*](#) on page 1162
 - [*IPAMv4 Network Usage Statistics*](#) on page 1162
 - [*DNS Statistics per DNS View*](#) on page 1163
 - [*DNS Statistics per Zone*](#) on page 1163
 - [*IPAMv4 Top Utilized Networks*](#) on page 1164
 - [*DHCPv4 Top Utilized Networks*](#) on page 1165
 - [*CPU Utilization Trend*](#) on page 1165
 - [*Memory Utilization Trend*](#) on page 1165
 - [*Traffic Rate*](#) on page 1166
 - [*Managing Reports*](#) on page 1172
 - [*Printing Reports*](#) on page 1172
 - [*Backing Up Reporting Data*](#) on page 1172
 - [*Scheduling the Backup of the Reporting Database*](#) on page 1173
 - [*Restoring the Reporting Database*](#) on page 1174

INFOBLOX REPORTING SOLUTION

The Infoblox reporting solution automates the collection, analysis, and presentation of core network service data that assists you in planning and mitigating network outage risks so you can manage your networks more efficiently.

You can add any of the supported Trinzic Reporting platforms as a member to the Grid and configure it as a reporting appliance. The reporting appliance collects data from Infoblox members, stores the data in the database, and generates reports that provide statistical data about IPAM, DNS, DHCP, and system activities and performance. For information about how an Infoblox reporting appliance works with your Grid, see [Introduction to Grids](#) on page 223. For information about Infoblox platforms that support reporting, see [Supported Platforms for Reporting](#).

You can set up a reporting appliance solely for reporting purposes. You cannot add licenses to run other services, such as DNS and DHCP, on a reporting appliance. When you set up a reporting appliance with valid licenses in the Grid, the reporting server acts as an indexer that receives data from Grid members while the members are forwarders that transmit information to the reporting server. Depending on your needs, you can enable certain Grid members as forwarders and disable others so the reporting server receives only the information you need from specific members. Note that the reporting service is disabled by default.

Note: Before removing a reporting license or reporting member, make sure that you disable the DNS Resource Records Last Queried feature, DNS Zones Last Queried feature, and Data collection for all DNS Queries to a Zone feature at the Grid level or at the member level. You cannot disable the Reporting tab at a later time as this tab will be removed when you remove the reporting license.

You must complete the following before you can view and manage reports in the Grid:

- Configure Grid reporting properties, as described in [Grid Reporting Properties](#) on page 1121.

When you enable the Grid reporting service, all members transmits data to the reporting server. You can disable data transmission from specific members to the reporting server.

After you set up and configure the reporting server and enable reporting service on specific members, you can view and manage reports through the **Reporting** tab of Grid Manager. Infoblox provides predefined reports and searches that capture useful information about the activities and performance of core network services (IPAM, DNS, DHCP, and system) in your Grid. You can also create your own reports and searches based on your organization's needs.

Note: You must enable reporting service on the Grid or members before you can use the reporting functions. Also, ensure that its host name has only alphanumeric characters, underscores, dots, and dashes. For information about how to enable the reporting service, see [Configuring General Grid Reporting Properties](#) on page 1121.

Reporting member uses two types of data sources: file based data source and script based data source to generate reports. When the reporting member is down or unreachable, the file based data sources are queued until the reporting member is up and running. However, the script based data sources are lost if the size of the queued data exceeds 500 KB.

The amount of data in the queue is managed as explained below:

- Rotates the reporting syslog files (extracted from /var/log/messages) at 120 MB retaining one older file. The data in the queue depends on the file size when the reporting member becomes unreachable.
- The CSV files overwrites the oldest data with the new data at regular intervals. So, the CSV file contains only the latest events.

The following table lists the report, report category, source type, data source type (file or script based), and frequency for updating the queue data:

Report Category	Reports	Source Type	Data Source (File Based or Script Based)	Update Frequency
DHCP Performance	DHCP Message Rate Trend	ib:dhcp:message	File based (csv)	Overwritten every 1 minute
	DHCPv4 Usage Trend DHCPv4 Range Utilization Trend	ib:dhcp:range	File based (csv)	Overwritten every 1 hour
DHCP Lease History	DHCP Lease History	ib:dhcp:lease_history	File based (syslog)	Rotates at 120 MB; retains one older copy; queued data is between 120 MB and 240 MB
	DHCP Top Lease Clients			
DHCP Fingerprint	Top Devices Identified Device Trend Device Class Trend Top Device Classes	ib:dhcp:lease_history	File based (syslog)	Based on summary search report, which is updated during the 16 th and 46 th minutes of each hour
	Top Devices Denied an IP Address	ib:dhcp:lease_history	File based (syslog)	Based on summary search report, which is updated during the 19 th and 49 th minutes of each hour
	Device Fingerprint Change Detected	ib:dhcp:lease_history	File based (syslog)	Executed every 24 hours
DNS Performance	DNS Response Latency Trend	ib:dns:perf	Script based	Executed every 1 minute
DDNS	DDNS Update Rate Trend	ib:ddns	File based (syslog)	Rotates at 120MB; retains one older copy; queued data is between 120MB and 240MB.
DDI Utilization	DHCPv4 Usage Statistics DHCPv4 Top Utilized Networks	ib:dhcp:network	File based (csv)	Overwritten every 1 hour
	IPAM Network Usage IPAM Top Networks	ib:ipam:network	File based (csv)	Overwritten every 1 hour

Report Category	Reports	Source Type	Data Source (File Based or Script Based)	Update Frequency
	DNS Zone Statistics Per DNS View	ib:dns:view	File based (csv)	Overwritten every 24 hours
	DNS Statistics per Zone	ib:dns:zone	File based (csv)	Overwritten every 24 hours
System Utilization	CPU Utilization Trend Memory Utilization Trend Traffic rate	ib:system	Script based	Executed every 1 minute
DNS Query	DNS Replies Trend	ib:dns:stats	Script based	Executed every 1 minute
	DNS Cache Hit Rate Trend	ib:dns:query:cache_hit_rate	Script based	Executed every 1 minutes
	DNS Query Rate by Query Type	ib:dns:query:qps	Script based	Executed every 1 minute
	DNS Query Rate by Server DNS Daily Query Rate by Server DNS Daily Peak Hour Query Rate by Server	ib:dns:query:by_member	Script based	Executed every 1 minute
	DNS Top Clients	ib:dns:query:top_clients	Script based	Executed every 10 minutes
	DNS Top Requested Domain Names	ib:dns:query:top_requested_domain_names	Script based	Executed every 10 minutes
	DNS Top Clients Per Domain DNS Top NXDOMAIN / NOERROR (no data) DNS Top SERVFAIL Errors Received DNS Top SERVFAIL Errors Sent DNS Top Timed-Out Recursive Queries	ib:dns:reserved	Script based	Executed every 10 minutes
	DNS Query Trend per IP Block Group	ib:dns:reserved	Script based	Executed every 5 minutes
Security	DNS Top RPZ Hits	ib:dns:reserved	Script based	Executed every 10 minutes
	DNS Top RPZ Hits by Clients	ib:dns:reserved	Script based	Executed every 10 minutes

Report Category	Reports	Source Type	Data Source (File Based or Script Based)	Update Frequency
	FireEye Alerts	ib:syslog	Script based	Updated immediately when alerts are logged in the syslog.
	Threat Protection Event Count By Severity Trend	ib:reserved1	File based (csv)	Overwritten every 5 minutes.
	Threat Protection Event Count By Member Trend			
	Threat Protection Event Count By Rule			
	Threat Protection Event Count By Time			
	Threat Protection Event Count By Category			
	Threat Protection Event Count By Member			

You can do the following in the **Reporting** tab:

- Configure Grid reporting settings, as described in [Grid Reporting Properties](#) on page 1121.
- Schedule the delivery of reports in PDF format to specified email addresses, as described in [Scheduling Report Deliveries](#) on page 1129.
- Create new reports, as described in [Adding New Reports](#) on page 1131.
- Create, modify, and delete user-defined searches, as described in [About Searches](#) on page 1132.
- View predefined reports, as described in [Predefined Reports](#) on page 1145.
- Create user defined reports, as described in [Managing Reports](#) on page 1172.
- Print a list of reports, as described in [Printing Reports](#) on page 1172.
- Back up and restore the reporting database, as described in [Backing Up Reporting Data](#) on page 1172 and [Restoring the Reporting Database](#) on page 1174.

Supported Platforms for Reporting

Infoblox provides a few reporting appliances that address your organization needs. [Table 38.1](#) lists the supported Trinzie Reporting platforms based on IP capacities and average DHCP leases and DNS queries per second:

Table 38.1 Trinzie Reporting Platforms

Enterprise Model	Supported Infoblox Appliance	Daily Maximum Data Consumption*
Very Large enterprises Service providers	Trinzie Reporting 4000 Appliance	20 GB
Large enterprises	Trinzie Reporting 2000 Appliance	10 GB
Mid-size enterprises	Trinzie Reporting 1400 Appliance	5 GB

Enterprise Model	Supported Infoblox Appliance	Daily Maximum Data Consumption*
Mid-size enterprises	Trinzic Reporting VM-800 (virtual appliance)	1 GB

Note: * The daily maximum data consumption includes all DNS, DDNS, IPAM, DHCP, and system traffic or events from all members with data transmission enabled within your Grid. When data traffic exceeds the daily maximum, the reporting server sends an SNMP trap and email notification, if configured. After five (5) daily maximum warnings in a rolling period of 30 days, you cannot view reports or perform any report related functions. For information about how to avoid this problem, see [Guidelines for Customizing Searches](#) on page 1133. Note that the reporting server continues to process incoming data during the violation state. However, you cannot view any reports or manage any reporting related functions until you fix the violation issue.

For information about the Trinzic Reporting platforms, their specifications, and how to install them as reporting appliances, refer to the following:

- *Infoblox Installation Guide for the Trinzic Reporting 4000 Appliance*
- *Infoblox Installation Guide for the Trinzic Reporting 2000 Appliance*
- *Infoblox Installation Guide for the Trinzic Reporting 1400 Appliance*
- *Infoblox Installation Guide for the Trinzic Reporting VM-800 Appliance*

Infoblox-4030 Supported Reports

The IB-4030 appliance supports the following reports:

DNS Reports	Security (DNS) Reports	System Reports
<ul style="list-style-type: none"> • DNS Replies Trend • DNS Cache Hit Rate Trend • DNS Query Rate by Query Type • DNS Response Latency Trend • DNS Query Rate by Server • DNS Daily Query Rate by Server • DNS Daily Peak Hour Query Rate by Server • DNS Top SERVFAIL Errors Sent/Received • DNS Top Timed-Out Recursive Queried • DNS Query Capture for DNS Domain 	<ul style="list-style-type: none"> • DNS Top RPZ Hits • DNS Top RPZ Hits by Client • FireEye Alerts Report • Threat Protection Event Count by Time • Threat Protection Event Count by Severity Trend • Threat Protection Event Count by Rule • Threat Protection Event Count by Member • Threat Protection Event Count by Member Trend • Threat Protection Event Count by Category 	<ul style="list-style-type: none"> • CPU Utilization Trend • Memory Utilization Trend • Traffic Rate Trend

GRID REPORTING PROPERTIES

After you set up a dedicated reporting appliance in your Grid, you can configure Grid settings to communicate with the reporting appliance and retrieve report data through the Grid Master. You can select specific report categories and configure report settings. You must be a superuser to view and configure Grid reporting properties.

Note: The more zones that are monitored, the longer it takes to bring up the “Grid Reporting Properties” editor.

Complete the following to set up your reporting solution:

- Configure general reporting properties, including the selection of report categories, as described in [Configuring General Grid Reporting Properties](#) on page 1121.
- Specify the network port for reporting, as described in [Setting Network Port for Reporting](#) on page 1122.
- Define email properties for messages related to reporting, as described in [Setting Email Properties](#) on page 1122.
- Configure the PDF settings of the reports that you schedule for delivery, as described in [Defining PDF Settings](#) on page 1122.
- Configure the reporting appliance to capture DNS queries and responses, as described in [Configuring the Capture of DNS Queries and Responses](#) on page 1123.

Note: Make sure that you click **Restart** to restart services each time after you change the reporting configuration.

Configuring General Grid Reporting Properties

To configure general Grid reporting properties:

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **General** -> **Basic** tab
3. After you complete the following, click **Save & Close**.
 - **Reporting Server:** Grid Manager displays the name of the reporting server.
 - **Enable Data Indexing:** Data transmission is disabled by default. When you select this check box, all Grid members transmit data to the reporting appliance. Enabling data transmissions for all members can affect the overall data consumption on the reporting server. For information about the daily maximum data consumption per day for your reporting appliance, see [Supported Platforms for Reporting](#) on page 1119.
 - **Report Category:** Select the reports you want the reporting server to generate. The reporting server automatically configures data sources and configurations required to generate the reports you select here. The required data is stored in the reporting server database. No report categories are selected by default. For a list of report categories, see [Predefined Report Categories](#) on page 1145.
 - **Report Settings:**
 - **Total Custom Reports per user:** Displays the total number of user-defined reports each user can create. The maximum number per superuser is 300 and per limited-access user is 5.
 - **Total Custom Reports for Grid:** Displays the total number of user-defined reports that are allowed in the Grid. The maximum number is 300. You cannot modify this field.

Setting Network Port for Reporting

All Grid members use port 9997 for reporting service by default. This port is used for data transmissions between the reporting member and other members. Ensure that you configure your firewall rules to allow traffic on this port. You can designate another network port for reporting purposes.

To set the network port for reporting:

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **General** -> **Advanced** tab and complete the following:
 - **Port:** Enter the port number you want to use for reporting purposes. The default port is 9997.
3. Save the configuration.

Setting Email Properties

You can enable and configure the settings of email messages that the appliance sends when there are changes to the link status of the network port and when certain events that affect the reporting feature arise.

To configure email properties:

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Email** tab and complete the following:
 - **Email As:** Enter the name of the admin or organization that sends the message.
 - **Email Subject Prefix:** Enter the subject prefix for the email messages.
 - **Email Format:** From the drop-down list, select the format for the email messages. The default is **HTML**.
3. Click Save & Close.

Defining PDF Settings

You can schedule to send reports in PDF format to email addresses that you define. For information about scheduling the delivery of reports, see [Scheduling Report Deliveries](#) on page 1129.

To define the PDF settings for the reports:

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **PDF** tab and complete the following:
 - **Paper Size:** From the drop-down list, select the paper size.
 - **Report Orientation:** From the drop-down list, select the page orientation for the report.
3. Click Save & Close.

Defining DNS Query Settings

You can configure up to 1000 zones to be monitored. You configure the zones you want to monitor by adding a zone to the table using the *Selector Dialog* or *Bulk Add Dialog* at **Grid Reporting Properties** -> **Basic** tab -> **DNS**. You may also monitor clients and capture DNS queries.

Note: Infoblox recommends that you keep the number of zones below 1000 because more than that may adversely affect performance.

Selecting Zones for Monitoring

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Basic** tab -> **DNS**.

- Under **DNS Resource Records Last Queried**, select the check box for **Monitor queries for resource records in the following zones**.
- Click the Add icon to add zones. From the drop-down list, click **Select Zones** to select zones in the *Zone Selector* dialog box, or click **Bulk Add Zones** to add zones in bulk.
- Under **DNS Zones Last Queried**, select the **Monitor queries made to the following zones** check box.

Note: If you disable **Monitor queries made to the following zones**, the value in the "Records Monitored" continues to display "Yes" until you toggle the flag.

- Save the configuration.
- To verify zones are being monitored, in the **DNS** tab -> **Zones** tab, make sure that the Records Monitored column displays "Yes" for the selected zones.

Configuring the Capture of DNS Queries and Responses

You can capture DNS queries and responses for later analysis. A capture file for logging DNS queries and responses is compressed and sent every 10 minutes or when it reaches 100 MB in size, whichever comes sooner. The capture file is automatically exported to an FTP or SCP server that you specify. Note that capturing DNS queries and responses will affect system performance. Infoblox recommends that you constantly monitor the FTP or SCP server to ensure that it has sufficient disk space. DNS queries and responses are stored on the appliance if the FTP or SCP server becomes unreachable. The maximum storage capacity varies based on the appliance model. After reaching the maximum limit, the appliance overwrites the old data with the new one. For information about the maximum hard drive space, see [Maximum Hard Drive Space used for DNS queries and Responses](#) on page 1126. The amount of data captured depends on the DNS query rate and the domains that are included in or excluded from the capture. For information about how to exclude domains, see [Excluding Domains From Query and Response Capture](#) on page 1127.

Capturing DNS Queries

You can capture queries to all domains or limit the capture to specific domains. You can apply the Bulk Add Domains feature to tailor query capture to a desired subset of domains or zones. While performing query captures, NIOS matches the specified domain name(s) and everything that belongs to the domain. For example, when you specify 'foo.com' as the domain, NIOS captures queries sent to 'foo.com,' 'mail.foo.com,' and 'ftp.foo.com.' NIOS captures queries to domains for which a name server is authoritative; it also captures recursive queries. Note that this feature does not support wildcard characters or regular expressions.

The DNS query generates a query message in the following format:

```
<dd-mm-YYYY HH:MM:SS:uuu> <client IP>#<port> query: <query_Domain name> <class name> <type name> <- or +>[SETDC] <(name server ip)>
```

Sample DNS query message:

```
30-Apr-2013 13:35:02.187 client 10.120.20.32#42386: query: foo.com IN A + (100.90.80.102)
```

Capturing DNS Responses

You can capture DNS responses for the DNS queries sent to the server. The amount of data captured depends on the domains that are included in or excluded from the capture. A DNS response is based on a query generated for a domain. In the response message, NIOS captures the TTL value of a resource record, the resource record type, and resource data.

Following are characteristics of the response messages:

- They log only the answer section and do not include the authority and additional sections.
- Responses to all queries are logged, including queries with the type "ANY."
- The RR (resource record) list is not available at the end of a response message if rcode has a value other than NOERROR or if the response is NOERROR (nodata).
- Responses to all RR types, including those records not managed by NIOS such as HINFO records, are logged. However, there are few exceptions for some of the scenarios with DNSSEC records.

- Responses containing DNSSEC RRs (DNSKEY, DS, NSEC, NSEC3, NSEC3PARAM, RRSIG) when queried for non-DNSSEC RRs are not logged. However, responses are logged if a DNSSEC RR is explicitly queried.
- DNS updates are not logged in responses.

DNS Response Message Format and Examples

The DNS query generates a response message in the following format:

```
<dd-mm-YYYY HH:MM:SS:uuu> client <client ip>#port <UDP or TCP>: [view: DNS view] query:
<queried domain name> <class name> <type name> response: <rcode> <flags> [<RR in text
format>; [<RR in text format>;] ...]
Flags = <- or +>[ATED]
```

where,

```
- = recursion not available
+ = recursion available (from DNS message header)
A = authoritative answer (from DNS message header)
t = truncated response (from DNS message header)
E = EDNS OPT record present (from DNS message header)
D = DNSSEC OK (from EDNS OPT RR)
V = responding server has validated DNSSEC records
```

Following are some DNS response samples:

Example 1: When querying an A record

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a2.foo.com IN A response:
NOERROR +AED a2.foo.com. 28800 IN A 1.1.1.2;
```

Example 2: When querying an AAAA record

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a4.foo.com IN AAAA response:
NOERROR +AED a4.foo.com. 28800 IN AAAA ab::a;
```

Example 3: When querying an A record over IPv6

```
07-Apr-2013 20:16:49.083 client 2001::2#57398 UDP: query: a2.foo.com IN A response: NOERROR
+AED a2.foo.com. 28800 IN A 1.1.1.2;
```

Example 4: When querying an A record over TCP

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 TCP: query: a2.foo.com IN A response:
NOERROR +ED a2.foo.com. 28800 IN A 1.1.1.2;
```

Example 5: When querying ANY record

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a2.foo.com IN ANY response:
NOERROR +ED a2.foo.com. 28800 IN A 1.1.1.2;
```

Example 6: When querying an A record with multiple addresses

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a1.foo.com IN A response:
NOERROR +ED a1.foo.com. 28800 IN A 1.1.1.1; a1.foo.com. 28800 IN A 11.1.1.1;
```

Example 7: When querying an aliased A record

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: c2.foo.com IN A response:
NOERROR +ED c2.foo.com. 28800 IN CNAME a2.foo.com.; a2.foo.com. 28800 IN A 1.1.1.2;
```

Example 8: When querying an NXDOMAIN

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: non-exist.foo.com IN A
response: NXDOMAIN +ED
```

Example 9: Response message for NOERROR/nodata

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: al.foo.com IN SRV response:
NOERROR +ED
```

Example 10: Response message for refused query

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: refused.com IN A response:
REFUSED +ED
```

Example 11: Response message when server fails

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#12345 UDP: query: servfail.com IN A response:
SERVFAIL +E
```

Example 12: Response message when query A record in a signed zone

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: al.signed.com IN A response:
NOERROR +ED al.signed.com. 28800 IN A 1.1.1.1;
```

Example 13: Response message for explicit query to DNSSEC RRs

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: al.signed.com IN RRSIG
response: NOERROR +ED al.signed.com. 28800 IN RRSIG A 5 3 28800 20130616004903 20130611234903
4521 signed.com. evROke7RbnkjFTsumT3JJg76bduFLfdEEnsziTXHQCbVYBS5rDy+qbUI
HCQuN/ldCNTJbZQ8MEhuatzfms+2Y5K2sU67P9Yg6GkOMxst2LcJiBm/
YqrYizBWGKpLF6J0PdX05133Xwq8XxUstUEJxKfuzcKSY6jaSduQIdFL v6A=; al.signed.com. 900 IN
RRSIG NSEC 5 3 900 20130616004903 20130611234903 4521 signed.com.
CnFmXMx9D+ZkDsztQbW2xx8XCROGNMBp0baxFXS/Pxxhg4PQcq58laI97y2XgqswN/wKNhY8p9hkes5+6t/ihCOIbw
FryxtdivPfYFFf3jafedFN ymZu05K9bYUfCUzzTGIRzoJYhxBM7xFT8fMvxni9ngsbLym82Tqv3Nua 6wU=;
```

To configure DNS queries and responses:

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Basic** tab -> **DNS** tab.
3. Under **Data Collection for all DNS Queries/Responses to a Domain**, complete the following:
 - Select the **Capture DNS Queries** check box to enable the capturing of DNS queries. This enables the feature set for configuration.
 - Select the **Capture DNS Responses** check box to enable the capturing of DNS responses. This enables the feature set for configuration.
4. Select **Capture queries to all domains** to capture queries to all domains and zones. This is applicable only for DNS queries.
5. Select **Limit capture to these domains** to capture DNS queries and responses to domains and zones one at a time.
6. Specify domains for DNS capture operations in the Domain table by clicking the Add icon, and choosing **Add Domain** or **Bulk Add Domains** from the menu.

To define the destination for capture files, do the following:

- Choose the transfer protocol from the **Export to** drop-down list: **FTP** or **SCP**.
 - In the **Directory Path** field, enter the directory to which the capture file will be saved on the server.
 - In the **Server Address** field, enter the IP address of the FTP or SCP server to which the capture files will be saved.
 - Enter the file server account **Username** and **Password** values.
7. Save the configuration.

[Table 38.2](#) lists the maximum hard drive space required for capturing DNS queries and responses for supported Infoblox appliance models.

Table 38.2 Maximum Hard Drive Space used for DNS queries and Responses

Supported Infoblox Appliances	Maximum Hard Drive Space for DNS Query /Response Capture (MB)
IB-250-A	700
IB-550-A	2500
IB-1050-A	4800
IB-1550-A, IB-1552-A	7000
IB-1852-A	21000
IB-2000-A	15000
Trinzic 100	400
Trinzic 810	900
Trinzic 820	3100
Trinzic 1410	6000
Trinzic 1420	10000
Trinzic 2210	12000
Trinzic 2220	28000
Infoblox-4010	40000
IB-VM-100	400
IB- VM-250 (50G and 120G)	700
IB- VM-550 (50G and 120G)	2500
IB- VM-1050 (50G and 120G)	5000
IB-VM-1550 (120G)	7000
IB-VM-1850 (120G)	22000
IB-VM-2000 (120G)	15000
IB-VM-810 (120G)	900
IB-VM-820 (120G)	3000
IB-VM-1410 (120G)	6000
IB-VM-1420 (120G)	10000
IB-VM-2210 (120G)	12000
IB-VM-2220 (120G)	28000
IB-VM-4010 (120G)	40000

Excluding Domains From Query and Response Capture

You can exclude individual domains and their subdomains from DNS query and response capturing. You can also use the Bulk Add Domains feature to a subset of domains to exclude them from query and response capturing. Subdomains can also be specified for exclusion. NIOS matches the specified domain names and their subdomains while filtering them in the Exclusion list. For example, when you specify 'foo.com' as the domain to be excluded, NIOS filters queries for 'foo.com,' 'mail.foo.com,' and 'ftp.foo.com.'

Note: IDNs are not supported for the domains that are added to the Inclusion list and Exclusion list. You can use the punycode representation of an IDN in these lists.

To exclude a domain from query and response capturing, do the following:

1. From the **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Basic** tab -> **DNS** tab.
3. Under **Data Collection for all DNS Queries/Responses to a Domain**, select the **Exclude the following domains** check box.
4. Click the Add icon and select **Add Domain** or **Bulk Add Domains** and specify domains in the Domain table.

Note: NIOS first matches the domains in the Exclusion list and then matches the domains in the Inclusion list. NIOS does not capture queries and responses for the subdomains in the **Capture DNS Queries/Responses** list (Inclusion list) if their domains are added to the **Exclude the following domains** list (Exclusion list).

The following table provides examples of domains and subdomains added to the Inclusion and Exclusion lists and the corresponding effects on the query and response capture operations:

Capture DNS Queries/Responses (Inclusion List)	Exclude the following domains (Exclusion List)	Queried Domain	Queries/Responses Captured	Results
foo.com	--	<ul style="list-style-type: none"> foo.com finance.foo.com 	Yes	Exclusion list is empty and therefore matches the Inclusion list. NIOS captures queries/responses made to foo.com and finance.foo.com
	--	<ul style="list-style-type: none"> corp1.com 	No	NIOS does not capture queries/responses made to corp1.com as this domain is not mentioned in the Inclusion list.
Capture All	foo.com	<ul style="list-style-type: none"> foo.com 	No	Matches the Exclusion list and NIOS does not capture queries made to foo.com.
		<ul style="list-style-type: none"> finance.foo.com 	No	Subdomain matches the Exclusion list and NIOS does not capture queries/responses made to finance.foo.com.
		<ul style="list-style-type: none"> corp1.com 	Yes	Does not match the Exclusion list. Matches the Inclusion list and therefore NIOS captures queries/responses made to corp1.com.

Capture DNS Queries/Responses (Inclusion List)	Exclude the following domains (Exclusion List)	Queried Domain	Queries/Responses Captured	Results
foo.com	it.foo.com	<ul style="list-style-type: none"> foo.com finance.foo.com 	Yes	Does not match the Exclusion list and therefore NIOS captures queries/responses made to foo.com and finance.foo.com.
		<ul style="list-style-type: none"> it.foo.com ms.it.foo.com 	No	Matches the Exclusion list and excludes their subdomains. NIOS does not capture queries/responses made to it.foo.com and ms.it.foo.com.
it.foo.com	foo.com			Domain is added to the Exclusion list and its subdomain is added to the Inclusion list. Therefore, this is not a valid configuration as queries/responses are not captured. The appliance displays a warning message for such invalid configuration.
it.foo.com	it.foo.com			Domain is added to both the Exclusion and Inclusion lists. This is not a valid configuration as queries/responses are not captured. The appliance displays a warning message for such invalid configuration.
foo.com	corp1.com			Domain added to the Inclusion list is not the subdomain of the domain added to the Exclusion list. This is a redundant configuration as the outcome is the same even if the domain is removed from the Exclusion list. The appliance displays a warning message for such invalid configuration.

Monitoring Client Queries

You can view the presence of clients in the network that are sending large numbers of queries to DNS zones or DNS domains. To monitor the top clients querying DNS zones, do the following:

Under **DNS Top Clients Per Domain**, select the **Monitor Queries made to the following zones** check box. Only authoritative zones are supported, to a limit of 1000 zones.

1. To select zones one at a time, choose individual check boxes.
2. Click the Add icon and select **Add Domain** or **Bulk Add Domains** to add new zone information for excluding.
3. To specify the number of clients to be listed, choose the **Top N Limit** value. The default value is 10.

Monitoring IP Block Group Queries

You can view the user defined IP block groups that are querying DNS domains. To monitor the IP Block Groups, do the following:

Under **DNS Query trend per IP Block**, select the **Monitor Queries made from the following groups** check box.

- Click the *Add* icon to add a group to the group table. From the drop-down list, click **Select Group** to select groups in the *Group Selector* dialog box, or click **Bulk Add Groups** to add multiple groups.

- To select all groups, select the **Group** check box. Or, select individual check box to select the group one at a time.
- To delete a group, select the group and click the *Delete* icon.

Defining Security

You can specify a limit to display the number of top clients, who receive re-written responses through the RPZ, in **DNS Top RPZ Hits**. You can also specify the total number of RPZ entries for each client. To specify the details under **Security**:

1. Enter a value for **Top N Limit** to specify the maximum number of top clients that can be listed in the report.
2. Specify the **Total RPZ Entries per Client**. This indicates the number of entries for each client in RPZ.

Note: You have to select the **Security** check box before you define values here. To select the check box, **Reporting** tab -> **Grid Reporting Properties** -> **General** tab -> **Basic** tab -> select the check box **Security** under **Report Category**.

SCHEDULING REPORT DELIVERIES

1. From the **Reporting** tab -> **Reports** tab, select a report dashboard.
2. Click **Schedule** from the Toolbar.
3. In the *Schedule Report Settings* editor, complete the following:
 - **Schedule Report Settings:** Select this check box to enable report delivery.
 - **Run Report Every:** From the drop-down list, select **Hour**, **Day**, **Week**, or **Month**. Depending on your selection, complete the following to define the day and time when the appliance runs the report:
 - **Minutes Past Hour:** When you select **Hour**, enter the number of minutes past the hour.
 - **Time:** When you select **Day**, enter the time of the day.
 - **Time** and **Weekday:** When you select **Week**, enter the time of the day in hh:mm:ss format and select the day of the week from the drop-down list.
 - **Day of Month** and **Time:** When you select **Month**, enter the day of the month and the time in hh:mm:ss format.
 - **Report Format:** All reports are sent in PDF format. You cannot modify this field.
 - **Send to owner:** Select this to send the report to the report owner. This is displayed only for custom reports.
 - **Email List:** Click the Add icon to add an email address to which the report is delivered. Grid Manager adds a row to the table. Select the row and enter the email address. Click the Add icon again to add another email address. You can also select an email address and click the Delete icon to delete it.
4. Click **OK**.
5. Click **Save & Close**.

Note: You cannot schedule a predefined report. GUI will throw error if you attempt to schedule a predefined report. However, you can configure a predefined report that contains charts to send in PDF format to email addresses that you define. Also, you cannot schedule a report that has multiple panels.

ABOUT REPORTS

Infoblox provides predefined reports that are categorized by core network service functions, such as DNS query and system utilization. Predefined reports contain predefined search criteria that retrieve specific data from the reporting database. Each predefined report is associated with a search. You cannot modify predefined searches. For information about searches, see [About Searches](#) on page 1132.

Predefined reports provide summary views for most of the data and trends in your Grid. They are tailored for optimal performance so your reporting server can run them in a reasonable speed under normal circumstances. Though you cannot modify the search criteria for these reports, you can define filters to further refine the report data. For more information about predefined reports, see [Predefined Reports](#). When you select a predefined report from the **Reporting** -> **Reports** tab, Grid Manager displays it as a single-panel report.

Note: Ensure that you have Flash plug-in installed to view reports.

You can also create a user-defined report by cloning a predefined report and its search or by adding a new report. You can then modify the search criteria for the new report and add new panels to the report. A user-defined report may contain more than one panel, and each panel contains a search, which generates a corresponding report. For example, you can add a report called “DHCP Activities,” and then add DHCP searches, such as DHCP Top Lease Clients and DHCP Lease History, to the new panels. When you save the “DHCP Activities” report, the reporting server saves all the searches in the panels and displays reports with updated data in the “DHCP Activities” report. User-defined reports can provide you with a single point of access when you want to review multiple reports that are relevant to the activities you want to monitor.

Note: IDNs are not supported on the reporting server. The reporting server manages IDNs in punycode. The reports generated by collecting reporting data from the DNS server displays all the data in punycode only. However, the DNS Zones Last Queried report and DNS Resource Records Last Queried report support IDNs because these reports collect data from the NIOS database.

Each report comes with a set of available filters that you can use to further refine the report data. Note that filters of different type use the AND logic and filters of the same type use the OR logic. For example, the appliance uses the AND logic when you apply both the “Time” and “A record” filters. It uses the OR logic when you apply the “Member equals marketing.corp100.com” and “Member equals hr.corp100.com” filters.

You can also use quick filters to narrow down the list of reports displayed in the **Reports** tab. Grid Manager provides the following quick filters: Local Reports, Global Reports, and System Reports. For information about how to use quick filters, see [Using Quick Filters](#) on page 68.

You can do the following in reports:

- Use available filters to refine the report data. Each report comes with a set of available filters.
- Change the report format, as described in [Changing Report Formats](#).
- Click **Reports** to go back to the **Reports** tab.
- In reports that are in table format, you can sort data based on the following:
 - You can only sort by certain columns, such as **Timestamp**.
 - If all data in a column are numeric, you can sort numerically.
 - If data in a column are in string format, the appliance sort by string.
- Use the navigation buttons to page through reports that contain multiple pages.
- You can mouse over a graph to display the coordinates of any point in the graph.

Note: All timestamps displayed in reports and all start and end times you select for a report are based on the time zone you configure on the reporting server, not the Grid Master.

You can do the following to create a user-defined report:

1. Add a new report, as described in [Adding New Reports](#) on page 1131.

2. Add new panels and searches to the dashboard, as described in [Adding New Panels to Reports](#) on page 1131.
- or
1. Clone an existing predefined report, as described in [Adding New Reports](#) on page 1131.
 2. Modify the search of the new report, if you have cloned the search, as described in [Modifying Searches](#) on page 1138.

Adding New Reports

When you add a new report, Grid Manager displays it in the **Reports** tab. You can add multiple panels and searches to the new report. For information, see [Adding New Panels to Reports](#) on page 1131.

To add a new report:

1. From the **Reporting** tab -> **Reports** tab, click **Add** from the Toolbar.
 2. In the *Reporting Dashboard* wizard, complete the following:
 - **Name:** Enter the name of the new report. Use ASCII characters.
 - **Description:** A brief description about this report. You cannot modify this field.
 - **Scope:** Select the **Set as Global Report** check box if you want to make this report globally available to all users. Only superusers can see this field when creating a new report. Limited-access users cannot see this field. They can only create personal reports.
 - **Comment:** Enter useful information about this dashboard.
 3. Click **Next** to enter extensible attributes.
- or
- Click **Save & Close**.
- Grid Manager saves the new report and displays it in the **Reports** tab.

Adding New Panels to Reports

You can add panels to a user-defined report, and then add a search to the newly created panel. When you add a search to the panel, Grid Manager generates the corresponding report in the panel. When you save the report, Grid Manager updates the searches in each panel.

To add a new panel to a new report:

1. From the **Reporting** tab -> **Reports** tab, select a user-defined report.
 2. In the Reporting Dashboard, click the Add Panel icon.
 3. In the *Reporting Search Selector* dialog box, select a search that you want to place in the panel.
 4. Grid Manager displays the following in the new panel:
 - **Panel Title:** Enter the report title here.
 - **Search Name:** Displays the search you have selected. You can click **Select** to select a new search.
 - **Search Type:** Displays the search type. You cannot modify this.
 - **Search Category:** Displays the search category. You cannot modify this.
 - **Panel Type:** Select the report type from the drop-down list. Depending on the search category, you can select different types of format, such as Table, Line Chart, or Stacked Area.
 5. Click **Save**.
- Grid Manager places the panel in the new report.

Note: You can rearrange panels by dragging and dropping them to their desired locations within the report.

Cloning Predefined Reports

To create a user-defined report:

1. From the **Reporting** tab -> **Reports** tab, select a report.
2. Click **Clone** from the Toolbar.
3. From the *Clone a Report* wizard, complete the following:
 - **Clone Report:** Display the name of the selected report.
 - **Name:** Enter the name of the new report. Use ASCII characters only.
 - **Set as Global Report:** Select the check box to make the new report globally available to all users. Only superusers can see this field when cloning a report. Limited-access users cannot see this field. They can only create personal reports.
 - **Clone Reporting Searches as well:** Select this to clone the search criteria of the report. If you do not select this, the appliance does not clone the search, and you cannot modify the search criteria for the report or add the search to a report panel.

The new report contains the same data as the original report. You can modify the search criteria for the new report if you clone its search.
4. Click **OK**.

Modifying User-Defined Reports

1. From the **Reporting** tab -> **Reports** tab, select a user-defined report.
2. Click the Edit icon.
3. The *Clone a Report* editor provides the following tabs from which you can modify data:
 - **General** tab: This tab displays the report name, its description and scope. You cannot modify these fields. You can only modify the comments you entered.
 - **Delivery Schedule** tab: Modify the delivery schedule of the report. For information, see [Scheduling Report Deliveries](#) on page 1129.
 - **Extensible Attribute** tab: Add or modify extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions** tab: Add and modify administrative permissions. For information, see [About Administrative Permissions](#) on page 160.
4. Click **Save & Close**.

Deleting User-Defined Reports

1. From the **Reporting** tab -> **Reports** tab, select a user-defined report.
2. Click the Delete icon.
3. The *Delete Confirmation* dialog box, click **Yes**.

ABOUT SEARCHES

Searches are criteria the reporting server uses to generate reports. Each predefined report has an associated search. You cannot modify or delete searches for predefined reports. You can however clone a predefined search when you clone its corresponding report. You can also create a new search by cloning an existing search, and then modify the search criteria. When you modify search criteria for a report, ensure that you follow some best practices to minimize impact on performance of the reporting server. For example, define specific and restrictive search criteria for a new

report so the reporting server can generate the report in a reasonable speed. Reports with search criteria that involve a large amount of data take longer to generate. Infoblox recommends that you follow some guidelines when modifying searches for new reports. For information about best practices for custom searches, see [Guidelines for Customizing Searches](#) on page 1133.

You can use the same search in multiple reports. You can also use quick filters to narrow down the list of searches displayed in the **Searches** tab. Grid Manager provides the following quick filters: Local Searches, Global Searches, System Searches, and Scheduled Searches. For information about how to use quick filters, see [Using Quick Filters](#) on page 68.

To create a new search, complete the following:

1. Clone an existing search, as described in [Cloning Searches](#) on page 1137.
2. Modify the settings of the new search, as described in [Modifying Searches](#) on page 1138.
3. Schedule how often the reporting server runs the search. By default, the reporting server runs a predefined search only when you open a report. For DHCP Lease History report however, the search is scheduled hourly by default. For information about scheduling a search, see [Scheduling Searches](#) on page 1138.
4. Define alerting settings to notify administrators when a match occurs for a Search.

You can also do the following in the **Searches** tab:

- Export search data in XML or CSV format, as described in [Exporting Searches](#) on page 1139.
- Schedule the export of search data, as described in [Scheduling Exports of Search Results](#) on page 1139.

Guidelines for Customizing Searches

When you follow some best practices for searching, you can optimize the performance of your reporting server and be able to view and manage your reports more efficiently. Depending on the type of search and the data you want to search for, Infoblox highly recommends that you use the following guidelines:

- Specify shorter start and end times when running detailed reports. When you define a long time duration, more data is included in the search, and the server takes longer to process the data.
- Faster searches depend on how specifically you define your search criteria. Be specific about the fields that you want to include in the search. For example, to view DHCP lease history for a specific member, change the search criteria to include only that member. Alerting after searches works in similar fashion.
- When you define how often the reporting server runs a search, be aware of other searches that the server is running as well. For example, when you schedule the server to run many searches at the same time, the server performance can be negatively affected. Try to stagger your searches whenever possible.
- Scheduling a search can minimize the workload on the reporting server. When you schedule a search, Grid Manager displays pre-existing report data, and this can reduce the workload on the reporting server. Though you can open a report each time you want to see up-to-date report data, on-demand searches can put more workload on the reporting server and may affect other searches and system performance.
- Create custom searches or detailed searches only when necessary. Infoblox provides predefined summary reports for most of the data and trends in your Grid. These reports are tailored for optimal performance so your reporting server can run them in a reasonable speed under normal circumstances. Summary reports collect events that have happened during a certain time period, and then summarizes the data before an update occurs. Since less raw data is involved in a summary report than in a detailed report, summary reports take less time to generate. Though you cannot see identical results between a summary report and a detailed report, creating too many detailed searches or reports may affect system performance. The reporting server categorizes each detailed and summary report by indexes; each index updates its data at different time intervals. You can use this information to determine the best way to obtain information you need without overloading the reporting server. For information about update frequencies for each report, see [Reporting Indexes and Update Time Intervals](#) on page 1134.
- Consider the daily maximum data consumption for your reporting server. Ensure that you select only the reports that you need from each Grid member so the reporting server is not overloaded with traffic. You can disable the reporting service on specific members to avoid unnecessary data transmissions. For information about daily maximum data consumption, see [Supported Platforms for Reporting](#) on page 1119.

- Review the daily data usage to avoid license violation. You can view the data usage from the *Member Status* widget in the Status Dashboard. When the data usage on the reporting server approaches or reaches the daily maximum, the appliance sends an SNMP trap and email notification, if configured. When you receive five (5) violation notifications in a rolling period of 30 days, you cannot view reports or configure report related functions. You must then contact Infoblox Technical Support to resolve the issue.

Note: The reporting server continues to process incoming data during the violation state. However, you cannot view any reports or manage any report related functions until you fix the violation issue.

Reporting Indexes and Update Time Intervals

[Table 38.3](#) lists the search indexes that the reporting server uses to generate reports. It contains information about the frequency of summary report updates for each report. Use this information to plan your reporting strategy for the Grid so you can optimize the performance of the reporting server.

Each summary report or search has its own update frequency. For example, the *DNS Top Requested Domain* report updates its data every 30 minutes, starting at the 4th minute of each half hour. It collects report data during the first 30 minutes of the previous 60 minutes. For example, if the report starts an update at 6:04 a.m., the data it collects is from 5:04 a.m. to 5:34 a.m.

The reporting server also uses this information to generate alerts. For example, once configured, Top Devices Identified alerts are executed at the 17th and 47th minutes of each hour (one minute after each update), regardless of whether DHCP fingerprint detection is enabled or disabled. For information about alerts, see [About Alerts](#) on page 1140.

Table 38.3 Reporting Indexes

Indexes	Reports	Report Data Updates
DHCP Lease History	DHCP Lease History (Detailed)	N/A
	DHCP Top Lease Clients (Summary)	Every 30 minutes, starting at the 16 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
DHCP	DHCP Message Rate Trend (Detailed)	N/A
	DHCPv4 Usage Statistics (Detailed)	N/A
	DHCPv4 Top Utilized Networks (Detailed)	N/A
	DHCPv4 Usage Trend (Summary)	Once per eight hours, starting at the 22 nd minute of each hour.
	DHCPv4 Range Utilization Trend (Summary)	Once per eight hours, starting at the 24th minute of each hour.
	Top Devices Identified (Summary)	Every 30 minutes, starting at every 16 th and 46 th minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.
	Device Trend (Summary)	Every 30 minutes, starting at every 16 th and 46 th minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.
	Device Class Trend (Summary)	Every 30 minutes, starting at every 16 th and 46 th minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.

Indexes	Reports	Report Data Updates
	Top Device Classes (Summary)	Every 30 minutes, starting at every 16 th and 46 th minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.
	Top Devices Denied an IP Address (Summary)	Every 30 minutes, starting at every 19 th and 49 th minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.
	Device Fingerprint Change Detected (Detailed)	Every 24 hours, starting at 00.15 AM.
IPAMv4	IPAM Network Usage (Detailed)	N/A
	IPAM Top Networks (Detailed)	N/A
	DNS Statistics per DNS View (Detailed)	N/A
	DNS Statistics per Zone (Detailed)	N/A
DNS	DDNS Update Rate Trend (Summary)	(Summary only) Every 30 minutes, starting at the 6 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Replies Trend (Summary)	(Summary only) Every 30 minutes, starting at the 18 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Response Latency Trend (Summary)	(Summary only) Every 30 minutes, starting at the 20 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top Clients (Summary)	(Summary only) Every 30 minutes, starting at the 2nd minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top Requested Domain Names (Summary)	(Summary only) Every 30 minutes, starting at the 4 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top RPZ Hits	(Summary only) Every 10 minutes of data summarizing when there is a constant load on RPZ.
	DNS Top RPZ Hits by Clients	(Summary only) Every 10 minutes of data summarizing when there is a constant load on RPZ.
	FireEye Alerts	(Summary only) Updated immediately when alerts are logged in the syslog.
	DNS Query Rate by Query Type (Detailed)	N/A
	DNS Query Rate by Server (Detailed)	N/A
	DNS Query Rate by Server (Detailed)	N/A
	DNS Daily Query Rate by Server (Summary)	Every 60 minutes, starting at the 32 nd minute of each hour.
	DNS Daily Peak Hour Query Rate by Server (Summary)	Every 60 minutes, starting at the 34 th minute of each hour.
	DNS Top Timed-Out Recursive Queries (Summary)	Every 30 minutes, starting at the 8 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.

Indexes	Reports	Report Data Updates
	DNS Top SERVFAIL Errors Sent (Summary)	Every 30 minutes, starting at the 6 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top SERVFAIL Errors Received (Summary)	Every 30 minutes, starting at the 7 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top NXDOMAIN / NOERROR (no data) (Summary)	Every 30 minutes, starting at the 5 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top Clients Per Domain (Summary)	Every 30 minutes, starting at the 3 rd minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
DNS Summary	DNS Top Clients (Summary)	Every 30 minutes, starting at the 2 nd minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Top Requested Domain Names (Summary)	Every 30 minutes, starting at the 4 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DDNS Update Rate Trend (Summary)	Every 30 minutes, starting at the 6 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Cache Hit Rate Trend (Summary)	Every 30 minutes, starting at the 8 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Query Rate by Query Type (Summary)	Every 30 minutes, starting at the 10 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Query Rate by Server (Summary)	Every 30 minutes, starting at the 12 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Replies Trend (Summary)	Every 30 minutes, starting at the 18 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DNS Response Latency Trend (Summary)	Every 30 minutes, starting at the 20 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
DHCP Summary	DHCP Message Rate Trend (Summary)	Every 30 minutes, starting at the 14 th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	DHCPv4 Usage Trend (Summary)	Every 8 hours, starting at the 22 th minute of each half hour. Data covers the first 8 hours of the previous 8.25 hours.
	DHCPv4 Range Utilization Trend (Summary)	Every 8 hours, starting at the 24 th minute of each half hour. Data covers the first 8 hours of the previous 8.25 hours.

Indexes	Reports	Report Data Updates
System	CPU Utilization Trend (Summary)	Once per 30 minutes, starting at the top of each half hour.
	Memory Utilization Trend (Detailed)	N/A
	Traffic Rate (Detailed)	N/A
System Summary	Memory Utilization Trend (Summary)	Every 30 minutes, starting at the 26th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	Traffic Rate (Summary)	Every 30 minutes, starting at the 28th minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
	CPU Utilization Trend (Summary)	Every 30 minutes, starting at the top of each half hour. Data covers the first 30 minutes of the previous 60 minutes.
Security	Threat Protection Event Count By Severity Trend (Summary)	Every 5 minutes.
	Threat Protection Event Count By Member Trend (Summary)	Every 5 minutes
	Threat Protection Event Count By Rule (Summary)	Every 5 minutes
	Threat Protection Event Count By Time (Summary)	Every 5 minutes
	Threat Protection Event Count By Category (Summary)	Every 5 minutes
	Threat Protection Event Count By Member (Summary)	Every 5 minutes

Note: When you filter a report by a time frame that is larger than the maximum retention period, the reporting server returns data within the maximum retention period. For example, when you try to view data of the *CPU Utilization Trend* report for the past six months, the server only returns data up to the last two months.

Cloning Searches

You can select an existing search and clone it. When you clone a search, you can give it a different name and make it globally available to other users.

To clone a search:

1. From the **Reporting** tab -> **Searches** tab, select a search you want to clone.
2. Click the Clone icon.
3. In the *Clone a Search* dialog box, complete the following :
 - **Clone Search:** Display the name of the original search.
 - **Name:** Enter a name for the new search. You can enter up to 255 ASCII characters.
 - **Set as Global Search:** Select this check box if you want to make this search globally available to all users.

The new search contains the same search criteria as the original search. You can modify the search criteria for the new search. For information, see [Modifying Searches](#).

Modifying Searches

1. From the **Reporting** tab -> **Searches** tab, select the search you want to modify.
2. Click the Edit icon.
3. The *Searches* editor provides the following tabs from which you can modify data:
 - **General** tab: This tab displays the search name, its description, category, and scope. You cannot modify these fields. You can however modify the report type and comments.

Note: When you change the report type, the appliance removes all the query items from the current search.

- **Settings** tab: You can redefine the filter settings for the search in this tab, as follows:
 - In the first drop-down list, select a field as the filter.
 - In the second drop-down list, select an operator for the filter.
 - Enter or select a value for the selected field and operator. Depending on the field and operator that you select, the field can be a text or an integer field. It can also be a drop-down list or a calendar widget.
 - Optionally, click + to add another filter. You can also click - to delete a filter.

Note: Filters of different type use the AND logic and filters of the same type use the OR logic. For example, the appliance uses the AND logic when you apply both the “Time” and “A record” filters. It uses the OR logic when you apply the “Member equals marketing.corp100.com” and “Member equals hr.corp100.com” filters. Be specific and restrictive when you define the search criteria. For information about best practices for searches, see [Guidelines for Customizing Searches](#) on page 1133.

- **Schedule** tab: Modify the search schedule, as described in [Scheduling Searches](#).
 - **Extensible Attribute** tab: Add or modify extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions** tab: Add and modify administrative permissions for this search object. For information, see [About Administrative Permissions](#) on page 160.
4. Click **Save & Close**.

Scheduling Searches

When you schedule a search, the appliance runs the search based on your configuration. The appliance also overrides the scheduled search if you open a report before the first scheduled search of the report. This also applies to the predefined *DHCP Lease History* report. Though you can only schedule searches for user-defined reports, you can run a search for a predefined report each time you open the report.

To schedule a search:

1. From the **Reporting** tab -> **Searches** tab, select a user-defined search.
2. Click the Edit icon.
3. In the *Searches* editor, select the **Schedule** tab and complete the following:
 - **Run search each time report dashboard loads:** When you select this, the appliance updates the report dashboard each time you open the report, based on the search criteria.
 - **Run search once:** When you select this, the appliance runs the search based on your settings:
 - **Day of Month:** Enter the day of the month you want to run the search.
 - **Time:** Enter a time in hh:mm AM/PM format or select a time from the wizard.
 - **Recurrence:** When you select this, the appliance updates the report dashboard based on the frequency you define here.

- **Run search every:** Enter a frequency and then from the drop-down list, select **Months, Weeks, Days,** or **Hours.** Note that the frequency you enter and select here may affect the overall performance of your reporting server. For more information about how to optimize your searches, see [Guidelines for Customizing Searches](#) on page 1133.

4. Click **Save & Close.**

Exporting Searches

You can export the data in the selected search in CSV (comma separated value) or XML format. This may take a long time depending on the amount of data you want to export.

To export data in a selected search:

1. From the **Reporting** tab -> **Searches** tab, select the search that you want to export.
2. Select the Export icon -> **Export Search Results.**
3. In the *Export Search Results* editor, complete the following:
 - **File name:** Enter the file name. You can enter up to 255 ASCII characters. Note that when you export a search, the appliance includes the file extension of .csv and .xml in the total number of characters. If you have a file name that contains 252 characters, the appliance cannot export the file because the file name exceeds the maximum number of characters.
 - **Format:** Select a format from the drop-down list.
 - **Max # of results to export:** Define the maximum number of data lines you want to export. The default is 10,000.

Scheduling Exports of Search Results

You can schedule the export of search data.

To schedule the export of search data:

1. From the **Reporting** tab -> **Searches** tab, select the search that you want to export.
2. Select the Export icon -> **Schedule Export Search Results.**
3. In the *Schedule Export* editor, complete the following:
 - **Reporting Search:** Click **Select** and select a user-defined report from the *Reporting Search Selector*. Note that you can only select user-defined searches.
 - **Disable Schedule:** Select this check box to disable the scheduling at this time. You can still save the configuration of this task.
 - **Format:** Select a format from the drop-down list.
 - **Export to:** Select the type of server to which you plan to export the data from the drop-down list.
 - **FTP:** Export the search results to an FTP server. This is the default.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Directory Path:** Enter the directory path and the file name of the export file. For example, you can enter /export/Infoblox_2009_10_20_15_30 on a Linux server, or c:\export\Infoblox_2009_10_20_15_30 on a Microsoft Windows server.
 - **Username:** Enter the username of your FTP account.
 - **Password:** Enter the password of your FTP account.
 - **SCP:** Export the search results to a SCP server.
 - **IP Address of SCP Server:** The IP address of the SCP server.
 - **Directory Path:** Enter the directory path and the file name of the export file. For example, you can enter /export/Infoblox_2009_10_20_15_30 on a Linux server, or c:\export\Infoblox_2009_10_20_15_30 on a Microsoft Windows server.
 - **Username:** Enter the username of your SCP account.

- **Password:** Enter the password of your SCP account.
- **TFTP:** Export the search results to a TFTP server.
 - **IP Address of FTP Server:** The IP address of the TFTP server.
 - **Directory Path:** Enter the directory path and the file name of the export file. For example, you can enter /export/Infoblox_2009_10_20_15_30 on a Linux server, or c:\export\Infoblox_2009_10_20_15_30 on a Microsoft Windows server.
- **Recurrence:** Select **Once** or **Repeat** from the drop-down list, and then complete the following accordingly:
- **Frequency:** For **Repeat** only. From the drop-down list, select DAILY, HOURLY, WEEKLY, or MONTHLY.
- **Day of Month:** Enter the day of month you want to run the search.
- **Time:** Enter a time in hh:mm:ss AM/PM format or select a time from the wizard.

ABOUT ALERTS

You can extend search functionality through the alerting feature. Alerts provide the following benefits:

- **Trend alerting:** For reports that show trends, alerts may be used to notify administrators when a trend reaches a specified threshold. You can enable searches to generate alerts on **summary events**, in which the alert reflects a trend of a given event.
- **Notification:** For detailed reports such as [DHCP Lease History](#), alerts may be used to notify administrators when certain conditions are met, such as a lease being handed out to a DHCP client with a specific MAC address.
- **Activity spike alerts:** You can enable searches to generate alerts for unusual spikes of activity. These events are called **raw/detail** events. They are more real-time than summary events and may not reflect the trend of a given event. However, they may reflect a network phenomenon that the administrator needs to know about.

You define alert settings against individual search types. You can edit alert settings for a single search at any time. One alert can be defined for each search. Note that you cannot define unlimited numbers of alerts without imposing possible performance penalties on both searches and alerting. Too many alerts, particularly [DHCP Lease History](#) alerts, may result in slower searching and alerting performance. In extreme cases, certain alerts may not be delivered.

One way to economize on the use of reports that may have a performance impact is to define individual alerts that use larger numbers of filtering conditions, instead of matching filtering conditions with each report. For example, consider an admin who wants to be alerted through the [DHCP Lease History](#) report for two conditions: A lease given out from “Member1” for 0C-0C-0B-14-CD-98 and a lease given out for “Member2” for 0C-0C-0B-23-FA-92. Instead of configuring two separate reports, define a single report with two filtering conditions, which use the OR logic, and then put those two sets of conditions into both groups.

Note that alerts are executed based on update frequencies for each corresponding search. For example, DHCP Lease History alerts are executed every 10 minutes, and Device Trend alerts are executed every 30 minutes at the 17th and 47th minutes of each hour (one minute after the search updates). For information about search indexes and update time intervals, see [Reporting Indexes and Update Time Intervals](#) on page 1134.

Alerting Logic

Carefully consider alerting logic when building your alerts.

You build new alerting expressions from left to right, using logical AND and OR boolean expressions to build more-complex alerting conditions.

You use three objects to define alerting expressions: **filters**, **operators**, and **values**. Working from left to right, choose the filter, choose the operator, and enter the required value where necessary in the third text field.

- Filter choices differ according to the chosen filter type. Two examples:
 - **DHCP Lease History** filters: **Members** (NIOS appliances) and **MAC/DUID**.
 - **DNS Query Rate by Server** filters: **Members** (NIOS appliances) and **QPS** (Queries per Second).
- Operators vary according to the chosen filter. Examples include **equals**, **does not equal**, **begins with**, **does not begin with** and **contains**. For example, if you choose a **Members** filter, **equals** is the only applicable operator.

- The third field is for data entry, where you enter the values to match against, or choose the member selection (for a **Members** filter, or a **QPS → Drops By** filter, where you enter a percentage value and choose the time frame for the filter).
- A plus (+) icon is directly associated with the new expression, which you click to add another filter. The logical AND is enforced using this (+) icon.
- When you add another filter to the expression, a new minus (-) icon appears, which allows deletion of its associated filter.
- A second plus (+) icon allows definition of a second (or more) expression block; this plus icon enforces the logical OR for the complete expression. This allows matching against multiple sets of filters, as in “match against this filter set OR match against this filter set OR match against this filter.”

Defining Alerts

To build or edit an alerting expression, do the following:

1. Go to **Reporting → Searches → search** check box, and then click **Edit → Alerting**.
2. Select the **Enable Alerting** check box. The expression builder and email features are activated.
3. Fill out the **Alert Subject** field, which is used by default for email, the syslog and trap alert messages.
4. Select the **Set Alert Severity** setting: **Critical**, **High**, **Medium**, **Low**, or **Info**. To enable emails for the chosen alert, select the **Send Email** check box.
5. Edit the **Email Subject**. Ensure that the default subject is removed and a new, more-specific subject is used. You are able to save an alert even if you leave this field blank.
6. Click the Add icon for the **Email List** and enter the admin email addresses.
7. The first expression with a **Choose Filter** drop-down list, a Choose Operator drop-down and a plus (+) icon directly associated with the new filter.
8. If necessary, select the **Send SNMP Trap** check box.
9. If necessary, select the **Send Syslog** check box. If you configure this option with an alert, the message goes to the syslog on the reporting member or indexer.
10. You now build your filtering expression for the alert.

Note: Email alerts are not generated if you change the host name of the member. Also, ensure that the search name or alert name and filter values do not contain any non-ASCII characters. Email notification and syslog events support non-ASCII characters.

Previewing Alerting Logic

1. If necessary, scroll down the Alerting dialog box to locate the Alert Filtering section. Select the **Enable Alerting** check box to enable this feature. You can do the following:
Click the **Preview** button at any time to preview the logical filtering expression.
Click the **Reset** button to remove the current expression settings.
2. In the first expression filter block, choose the new filter type from the **Choose Filter** drop down list.
3. Select the required operator from the **Choose Operator** drop-down menu (when applicable). Some filter types will have only a single operator, which is automatically selected.
4. In the third drop-down, enter the desired value or choose it (such as a Grid member) from the menu.
5. Click the Add icon (+) within the block to create a new filter using a logical AND with the first one, or click the Add icon (+) outside the block to create a new filter using a logical OR.
6. Note that the **Choose Filter** drop down list is the fundamental building block for each expression.
7. Save your configuration.

Click the **Advanced** tab to set alert throttling, which prevents excessive alert notifications from being sent to administrators.

Defining Advanced Alert Settings

Alert throttling prevents excessive or unnecessary instances of triggered alerts. Once an alert is triggered for the first time, subsequent instances of the same alert can be suppressed for specified frequencies between minutes up to weeks. This allows alerts to be sent at appropriate intervals based on the necessity of the alert type.

To define alert throttling and email settings for any specific report, do the following:

1. Select the check box for a report in the **Reporting → Searches** tab.
2. Go to **Reporting → Searches → search** check box → click **Edit → Alerting**.
3. Click the **Advanced** tab.
4. Select **Set Throttling** to enable prevention of excessive numbers of alerts after the first instance occurs.
5. In the **After triggering the alert, do not trigger until** section, enter the numeric value and select **Minutes, Hours, Days** or **Weeks** based on the alert type and the needs of the organization.
6. Save the configuration.

ABOUT IP BLOCKS AND IP BLOCK GROUPS

You can configure IP addresses, subnets, or a mix of multiple IP addresses and subnets into IP blocks, and then assign them to IP block groups for monitoring and tracking queries made to specific IP blocks. You can also configure as many groups as necessary and assign them to specific clients. Note that assigning more IP block groups results in monitoring more queries, which may affect the performance of the reporting server. You can generate a report to monitor queries made to these user-defined IP block groups or IP blocks. For information, see [DNS Query Trend per IP Block Group](#) on page 1162.

Guidelines while configuring IP blocks:

- You cannot configure arbitrary IP address ranges, such as 192.168.0.1 to 192.168.0.100 as an IP block.
- You cannot add or modify an IP block that overlaps with another IP block in a different group. However, you can add an IP block that overlaps with another IP block in the same IP block group.

Note: The appliance restarts the DNS service after you assign or unassign an IP block group at the Grid or member level. Also, the appliance restarts the DNS service when you modify or delete an IP address or IP block group assigned to the Grid or member or when you add, modify, or delete an IP block in such IP block groups.

You can do the following in the Groups panel:

- Add IP block groups, as described in [Adding IP Block Groups](#) on page 1143.
- Modify IP block groups, as described in [Modifying IP Block Groups](#) on page 1143.
- Add IP block, as described in [Adding IP Blocks](#) on page 1143.
- Modify IP blocks, as described in [Modifying IP Blocks](#) on page 1144.
- Delete IP block groups and IP blocks, as described in [Deleting IP Block Groups and IP Blocks](#) on page 1144.
- Print IP block groups and IP blocks, as described in [Printing IP Block Groups and IP Blocks](#) on page 1144.
- Export IP block groups and IP blocks, as described in [Exporting IP Block Groups and IP Blocks](#) on page 1144.

In addition, you can also do the following:

- Use filters or the **Go to** function to navigate to a specific group. You can also create quick filters to save frequently used filter criteria. For information about how to use quick filters, see [Using Quick Filters](#) on page 68.
- Use Global Search to search for IP block groups and IP blocks. For information, see [Global Search](#) on page 59.

- Use Smart Folders to organize IP block groups and IP blocks. For information, see [Smart Folders](#) on page 139.
- Import and export groups in CSV format. For more information about CSV import feature, see [About CSV Import](#) on page 86.

Adding IP Block Groups

To add a new group:

1. From the **Reporting** tab, select the **Groups** tab -> *Group* -> Add.
or
From the **Groups** tab, expand the Toolbar and click **Add** -> IP Block Group.
2. In the *Add IP Block Group* wizard, complete the following:
 - **Name:** Enter the name of the group.
 - **Comment:** Enter useful information about the group.
3. Do one of the following:
 - Click **Save & Close** to add the IP block group and close the wizard.
 - Click **Save & Edit** to add the IP block group and launch the editor. You can edit the details.
 - Click **Save & New** to add the IP block group and launch the wizard again to add another group.
 - Click **Save & Open** to add the IP block group and open the IP block group.

Modifying IP Block Groups

1. From the **Reporting** tab, select the **Groups** tab -> *Group*.
2. Click the Edit icon.
3. In the **General** tab of the *IP Block Group* editor, you can modify the group name and comment.
4. Click **Save & Close** to save the configuration.

Note: You can perform inline editing by double-clicking the row of data that you want to modify. The appliance displays the inline editing editor in the selected row. Click Save after modifying the data.

Adding IP Blocks

In a group, you can add as many subnets/IP addresses as necessary. Note that adding more IP addresses results in monitoring more queries, which may affect the performance of the reporting server.

1. From the **Reporting** tab, select the **Groups** tab -> *Group* -> Add.
or
From the **Groups** tab, expand the Toolbar and click **Add** -> IP Block.
2. In the *Add IP Block* wizard, complete the following:
 - **Group:** Click **Select** to select a group. When there are multiple groups, Grid Manager displays the *Group Selector* dialog box to select a group. Click a group name in the dialog box. You can use filters or the Go to function to narrow down the list.
 - **Address:** Enter the source IPv4/IPv6 addresses or the IPv4/IPv6 subnets.
 - **Comment:** Enter useful information about the IP block.
3. Do one of the following:
 - Click **Save & Edit** to add an IP address or IP block and launch the editor. You can edit the details.
 - Click **Save & New** to add an IP address or IP block and launch the wizard again to add another IP block.
 - Click **Save & Close** to add an IP address or IP block and close the wizard.

Modifying IP Blocks

1. From the **Reporting** tab, select the **Groups** tab -> *Group*.
2. Select an IP address or IP block you want to modify and click the Edit icon.
3. In the **General** tab of the *IP Block* editor, modify the IP address or comment.
4. Click **Save** to save the configuration.

Note: You can modify description by using inline editing. Double-click the row that you want to modify, the appliance displays the inline editing editor in the selected row. Click **Save** after modifying comment. You cannot modify IP address using inline editing editor.

Deleting IP Block Groups and IP Blocks

1. For IP block groups: From the **Reporting** tab, select the **Groups** tab -> *Group*.
For IP blocks: From the **Reporting** tab -> select the **Groups** tab -> *Group* -> *IP block*.
2. Click the Delete icon.
3. In the *Delete Confirmation* dialog box, click **Yes**.

Exporting IP Block Groups and IP Blocks

You can export displayed data or you can export the group list in CSV (comma separated value) format. Exporting group lists or group data may take a few moments based on the amount of exported data.

To export displayed data:

1. For IP block groups: From the **Reporting** tab, select the **Groups** tab -> *Group*.
For IP blocks: From the **Reporting** tab -> select the **Groups** tab -> *Group* -> *IP block*.
2. From the Export drop-down menu, select **Export visible data**. For more information on how to export, see [Exporting Displayed Data](#) on page 91.

To export all data to a CSV file:

1. For IP block groups: From the **Reporting** tab, select the **Groups** tab -> *Group*.
For IP blocks: From the **Reporting** tab, select the **Groups** tab -> *Group* -> *IP block*.
2. From the Export drop-down menu, select **Export data in Infoblox CSV Import format**. For more information on how to export, see [Exporting Data to Files](#) on page 89.

Printing IP Block Groups and IP Blocks

1. For IP block groups: From the **Reporting** tab, select the **Groups** tab -> *Group*.
For IP blocks: From the **Reporting** tab -> select the **Groups** tab -> *Group* -> *IP block*.
2. Click the Print icon. For more information on how to print, see [Printing from Grid Manager](#) on page 91.

PREDEFINED REPORTS

To view all available reports, from the **Reporting** tab, select the **Reports** tab. Grid Manager displays a list of predefined reports that you have selected when you set up your Grid or member reporting properties. You cannot modify or delete predefined reports or the search criteria of predefined reports. You can however modify user-defined reports. For information about how to create user-defined reports, see [Managing Reports](#) on page 1172.

Note: You cannot schedule a predefined report. GUI will throw error if you attempt to schedule a predefined report. However, you can configure a predefined report that contains charts to send in PDF format to email addresses that you define. For more information about scheduling reports, see [Scheduling Report Deliveries](#).

Predefined reports are classified in the categories listed in [Table 38.4](#). You can select the categories of report data you want the members to forward to the reporting server. For information about selecting report categories, see [Configuring General Grid Reporting Properties](#) on page 1121. No report categories are selected by default.

Predefined Report Categories

[Table 38.4](#) lists the report categories and their corresponding reports.

Note: The DNS Zones Last Queried report and DNS Resource Records Last Queried report support IDNs.

Table 38.4 Report Categories

Report Category	Corresponding Reports	Displays IDNs in Punycode (Yes/No)
DDNS Query	DDNS Update Rate Trend	Yes
DHCP Lease	DHCP Lease History	Yes
	DHCP Top Lease Clients	IDN is not supported
DHCP Fingerprints	DHCP Fingerprint Reports	
	Top Devices Identified	
	Device Trend	
	Device Class Trend	
	Top Device Classes	
	Top Devices Denied an IP Address	
DHCP Performance	Device Fingerprint Change Detected	
	DHCPv4 Usage Statistics	
	DHCPv4 Range Utilization Trend	
	DHCPv4 Usage Trend	
DNS Performance	DHCP Message Rate Trend	
	DNS Query Rate by Query Type	
	DNS Query Rate by Server	
	DNS Response Latency Trend	
	DNS Daily Query Rate by Server	

Report Category	Corresponding Reports	Displays IDNs in Punycode (Yes/No)
DNS Query	DNS Daily Peak Hour Query Rate by Server	
	DNS Top Requested Domain Names	Yes
	DNS Cache Hit Rate Trend	
	DNS Top Clients	
	DNS Replies Trend	
	DNS Top NXDOMAIN / NOERROR (no data)	Yes
	DNS Top Clients Per Domain	Yes
	DNS Top SERVFAIL Errors Received	Yes
	DNS Top SERVFAIL Errors Sent	Yes
	DNS Top Timeout Recursive Queries	Yes
	DNS Daily Query Rate by Server	
	DNS Query Trend per IP Block Group	
DNS Last Queried	DNS Resource Records Last Queried	
	DNS Zones Last Queried	
IPAMv4 Utilization Reports	IPAMv4 Network Usage Statistics	
	DNS Statistics per DNS View	
	DNS Statistics per Zone	Yes
	IPAMv4 Top Utilized Networks	
	DHCPv4 Top Utilized Networks	
System Utilization	CPU Utilization Trend	
	Memory Utilization Trend	
	Traffic Rate	
Security	DNS Top RPZ Hits	
	DNS Top RPZ Hits by Clients	
	FireEye Alerts	
	Threat Protection Event Count By Severity	
	Threat Protection Event Count By Member	
	Threat Protection Event Count By Rule	
	Threat Protection Event Count By Time	
	Threat Protection Event Count By Category	
	Threat Protection Event Count By Member	

Applying Time Filters

You can generate a report for a specific time interval by applying time filters. These filters display the date and time based on the time zone set in your user profile by default. For more information about how to configure a time zone, see [Setting the Browser Time Zone](#) on page 51.

To apply a time filter:

1. From the **Reporting** tab -> **Reports** tab, select a report.
2. Click **Show Filter** to enable the function.
3. In the filter section, complete the following:
 - **First (filter criterion):** Select Time or Start Time from the drop-down list.
 - **Operator:** Displays Equals when you select Time or Start Time.
 - **Value:** Select the time you want. You can also select time zone when you select Start Time.
 - **End Time:** Select an end time. Ensure that the end time is after the start time.
4. Do one of the following:
 - Click **Save** to save the filter criteria.

Note: By default, the Start time and End time are based on the time zone set in your user profile. When you save the time filter and keep the report open, the filter displays the Start time and End time you have selected. However, when you later retrieve the report or search, these filters are displayed based on the time zone in your user profile.

- Click **Apply** to apply the filter criteria to the report without saving the filter criteria into the database.

Changing Report Formats

All predefined reports are displayed in default formats, such as line graphs or tables. You can change the presentation of certain reports by selecting an available format.

To change the report format:

1. From the **Reporting** tab -> **Reports** tab, select a report.
2. In the Reporting Dashboard, click the Configure icon.
3. In the Configuration panel, complete the following:
 - **Panel Type:** Select the report type from the drop-down list. Depending on the search category, you can select different types of format, such as table, line chart, or stacked area. Grid Manager displays the report in the selected format.

DDNS Query Reports

DDNS Update Rate Trend

The *DDNS Update Rate Trend* report provides information about the dynamic DNS (DDNS) updates that occur on the DNS service. The default report shows a line graph that tracks the rate of DDNS updates (counts per second) by query type in the given time frame.

This report displays DDNS updates per second by the following query type: Success, Failure, Reject, and Prerequisite Reject. The time is displayed according to the time zone specified on the reporting server in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.

- **Members:** Filter by a specific member.
- **Response Type:** Filter by Success, Failure, Reject, or Prerequisite Reject.
- **Source IP Address:** Specify an IP address of the requesting source.
- **DNS Zone:** Filter by a specific DNS zone.

DHCP Lease Reports

DHCP Lease History

The *DHCP Lease History* report provides DHCP lease history in a given time frame. The search of the *DHCP Lease History* report is scheduled hourly by default.

DHCP Lease History reports can impose heavier system loads than for other alert types in the NIOS system. Avoid defining too many custom reports or alerts of this type for Grid reporting. Other types of reports do not impose significant performance restrictions. Also see [About Alerts](#) on page 1140 for methods to avoid this issue.

Note: When you join a new member to the Grid and do not start reporting service on the member, lease history for this member is not captured in the *DHCP Lease History* report. You can view lease history for this member in the **Data Management** tab -> **DHCP** tab -> **Leases** tab.

The default report displays the following information in table format:

- **Time:** Filter by last minute, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** The DHCP member that granted the lease.
- **Member IP:** The IP address of the DHCP member that granted the lease.
- **Lease IP:** The IP address of the lease.
- **Protocol:** Indicates whether the lease is for an IPv4 or IPv6 address.
- **Action:** The status of the lease. This can be one of the following: **Issued**, **Renewed**, **Freed**, or **Abandoned**.
- **Hostname:** The host name that the DHCP client sent to the appliance using DHCP option 12.
- **MAC/DUID:** For an IPv4 address, this is the MAC address of the lease. For an IPv6 address, this is the DUID (DHCP Unique Identifier) of the DHCP client that received the lease.
- **Lease Start:** The timestamp when the lease started.
- **Lease End:** The timestamp when the lease ended.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the leased client that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not match the filter criteria and those that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.

You can use any of the displayed fields, plus the start time and end time, as filters to get specific information in this report.

DHCP Top Lease Clients

The *DHCP Top Lease Clients* report provides information about the DHCP clients that have issued, renewed, and freed within a certain time frame.

This report shows the following information:

- **MAC/DUID:** The MAC address or DUID of the DHCP client.
- **Issued:** The total number of DHCP lease issued.
- **Renewed:** The number of DHCP lease renewals.
- **Freed:** The number of leases that were released.
- **MAC/DUID Total:** The total number of DHCP leases that were being requested, renewed, and released.

- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the leased client that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not match the filter criteria and those that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#) on page 1031.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by a specific member.
- **TopN:** The number of DHCP clients that have requested the most leases.
- **Report On:** Filter by leases that were issued, renewed, or freed.
- **Fingerprint:** Filter by DHCP fingerprint.

DHCP Fingerprint Reports

Top Devices Identified

The *Top Devices Identified* report lists the top DHCP fingerprints or detected operating systems for requesting clients. The appliance uses DHCP fingerprint detection to identify the operating systems or vendor IDs of remote clients. For more information about DHCP fingerprint detection, see [DHCP Fingerprint Detection](#) on page 1031. The default report displays the top 10 operating systems on which requesting clients are running within the last 24 hours.

The appliance lists the top detected operating systems or vendor IDs in table format. This report shows the total number of different MAC devices that requested a lease during each four hour interval. You can click a specific row in the table to view a list of leased clients that belong to the selected operating system or device type. Grid Manager displays another report that specifies more detailed information, such as the leased IPs and MAC addresses for each device that matches the selected DHCP fingerprint. The lease history for a fingerprint shows all the lease events that occurred during the four hour interval. It represents the number of devices that use the MAC/DUID as the unique identifier. Note that a single MAC address may have several lease events that occur within the four hour interval. Hence, the total number of each fingerprint will not be equal to the lease history of a fingerprint.

Note: You can use all available filters for the parent *Top Devices Identified* report, but you can filter the detailed report using only the **Fingerprint** column.

This report displays a table that contains the following information for each top DHCP fingerprint:

- **Fingerprint:** The name of the DHCP fingerprint or vendor ID for the requesting clients.
- **Total:** The total number of leased clients that belong to this DHCP fingerprint.
- **% of all devices:** The percentage of the leased clients belonging to this DHCP fingerprint over the total number of requesting clients.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **TopN:** Filter by the number of DHCP fingerprints that have the most requesting clients.
- **Time:** Filter by last hour, last day, last month, or last year.
- **Start Time:** Specify a start time. When you select **within the last**, the filter expression changes to “1” and displays selections of **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Enter the desired value and select the time unit, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You select **within the last** for this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member on which the filter runs.
- **Network View:** Filter by a specific network view.
- **Network:** Filter by a specific network.
- **DHCP Range:** Filter by a specific network range.

- **CIDR:** Filter by the subnet mask in CIDR format.
- **Fingerprint:** Filter by the name of the DHCP fingerprint or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Device Class:** Filter by the device category to which the leased client belongs.

Device Trend

The *Device Trend* report provides trends for the top operating systems used by remote clients in a given time frame. The default report displays line graphs for the top 10 operating systems used by remote clients over the last 24 hours. Each of the operating system is represented with a different color line graph. For more information about DHCP fingerprint detection, see [DHCP Fingerprint Detection](#) on page 1031.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last hour, last day, last month, or last year.
- **Start Time:** Specify a start time. When you select **within the last**, the filter expression changes to “1” and displays selections of **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Enter the desired value and select the time unit, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can select **within the last** for this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member on which the filter runs.
- **Network View:** Filter by a specific network view.
- **Network:** Filter by a specific network.
- **DHCP Range:** Filter by a specific network range.
- **CIDR:** Filter by the subnet mask in CIDR format.
- **Fingerprint:** Filter by the name of the DHCP fingerprint or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Device Class:** Filter by the device category to which the leased client belongs.

Device Class Trend

The *Device Class Trend* report provides trends for the top device classes used by remote clients in a given time frame. The default report displays line graphs for the top device classes used by remote clients over the last 24 hours. Each of the device class is represented with a different color line graph.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time. When you select **within the last**, the filter expression changes to “1” and displays selections of **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Enter the desired value and select the time unit, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can select **within the last** for this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member on which the filter runs.
- **Network View:** Filter by a specific network view.
- **Network:** Filter by a specific network.
- **DHCP Range:** Filter by a specific network range.
- **CIDR:** Filter by the subnet mask in CIDR format.
- **Fingerprint:** Filter by the name of the DHCP fingerprint or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Device Class:** Filter by the device category to which the leased client belongs.

Top Device Classes

The *Top Device Classes* report lists the top DHCP fingerprint device class for requesting clients. The default report displays the top 10 device classes along with the percentage of leased devices within the last 24 hours. The appliance lists the top detected device class in table format. You can click a specific row in the table to view all the devices in the class that belong to the selected device class. GUI displays the fingerprints that are detected under a selected device class. The total number of fingerprints of a specific device class is equal to the total number that is displayed against the corresponding device class.

This report displays a table that contains the following information for each top DHCP fingerprint device class:

- **Device Class:** The device category or fingerprint class for the requesting clients.
- **Total:** The total number of leased clients that belong to this DHCP fingerprint class.
- **% of all devices:** The percentage of the leased clients belonging to this DHCP fingerprint class over the total number of requesting clients.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **TopN:** Filter by the number of DHCP fingerprints that have the most requesting clients.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time. When you select **within the last**, the filter expression changes to “1” and displays selections of **min(s)**, **hour(s)**, **day(s)**, **week(s)**, or **month(s)**. Enter the desired value and select the time unit, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You select **within the last** for this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member on which the filter runs.
- **Network View:** Filter by a specific network view.
- **Network:** Filter by a specific network.
- **DHCP Range:** Filter by a specific network range.
- **CIDR:** Filter by the subnet mask in CIDR format.
- **Fingerprint:** Filter by the name of the DHCP fingerprint or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Device Class:** Filter by the device category to which the leased client belongs.

Top Devices Denied an IP Address

The *Top Devices Denied an IP Address* report lists the top DHCP fingerprint devices used by remote clients that were denied a lease or an IP address based on the fingerprint filter criteria you specified. The default report displays the top 10 devices per combination of fingerprint and network which were denied an IP address within the last 24 hours. For example, if the same device is denied from two separate networks during the past 24 hours, and/or with different fingerprints, then multiple events will be listed in the table corresponding to this device.

This report displays a table that contains the following information for each denied DHCP fingerprint device class:

- **Mac/DUID:** The Mac address or DUID of the client’s device.
- **Fingerprint:** The fingerprint description of the device used by remote clients.
- **Device Class:** The DHCP fingerprint class of the device used by remote clients.
- **Network:** The network to which the DHCP range belongs. For shared network, the network is the first network where the lease is prohibited due to fingerprint filter.
- **Attempts:** The total number of attempts by remote clients for an IP address in a given time frame.
- **Last Attempt:** The time stamp of the last attempt by remote client for an IP address in a given time frame.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **TopN:** Filter by the number of DHCP fingerprints that have the most requesting clients.
- **Time:** Filter by last day, last week, last month, or last year.

- **Start Time:** Specify a start time. When you select **within the last**, the filter expression changes to “1” and displays selections of **min(s)**, **hour(s)**, **day(s)**, **week(s)**, or **month(s)**. Enter the desired value and select the time unit, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You select **within the last** for this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member on which the filter runs.
- **Network View:** Filter by a specific network view.
- **Network:** Filter by a specific network.
- **CIDR:** Filter by the subnet mask in CIDR format.
- **Fingerprint:** Filter by the name of the DHCP fingerprint or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Device Class:** Filter by the device category to which the leased client belongs.

Device Fingerprint Change Detected

The *Device Fingerprint Change Detected* report provides information about the devices whose fingerprint data gets changed in a given time frame. In other words, this report includes all devices used by remote clients that were detected to have the same Mac address but different device class in a given time frame.

The following example illustrates how the fingerprint data can change in a given time frame:

A client device having dual boot option may request for an IP address while switching between operating systems, resulting in a change of fingerprint data and if a client’s device uses Mac Boot Camp, the mac address remains unchanged, but fingerprint data changes when it switches operating system.

Note: The *Device Fingerprint Change Detected* report includes all devices whose fingerprint data has been changed within the last seven days. It ignores devices whose fingerprint data has been changed for more than seven days.

This report displays a table that contains the following information:

- **Time:** The time the lease was obtained.
- **Mac/DUID:** The Mac address or DUID of the client’s device.
- **Current Device Type:** The current fingerprint description of the device.
- **Current Device Class:** The current fingerprint class of the device.
- **Previous Device Type:** The fingerprint description of the device before changing the fingerprint data.
- **Previous Device Class:** The fingerprint class of the device before changing the fingerprint data.
- **Lease IP:** The lease IP address of the device.
- **Action:** The current status of the lease. The lease status can be one of the following: **Issued**, **Renewed**, **Freed**, or **Abandoned**.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **Time:** Filter by last hour, last day, or last week. The default value is last day.
- **Start Time:** Specify a start time. When you select **within the last**, the filter expression changes to “1” and displays selections of **min(s)**, **hour(s)**, **day(s)**, **week(s)**, or **month(s)**. Enter the desired value and select the time unit, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You select **within the last** for this filter.

Note: Make sure that the specified duration is less than seven days while applying Start Time and End Time filters.

- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member on which the filter runs.
- **Network View:** Filter by a specific network view.

- **Network:** Filter by a specific network.
- **DHCP Range:** Filter by a specific network range.
- **CIDR:** Filter by the subnet mask in CIDR format.
- **Previous Fingerprint:** Filter by the previously defined DHCP fingerprint name or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Current Fingerprint:** Filter by the current DHCP fingerprint name or vendor ID of the leased client. Note that when you select vendor ID as the filter criterion, the appliance can return lease information that belongs to different DHCP fingerprints with the same vendor ID.
- **Previous Device Class:** Filter by the previous device category to which the leased client belongs.
- **Current Device Class:** Filter by the current device category to which the leased client belongs.
- **Device Class Changed:** Filter by the change in device category. The default value is Yes.

DHCP Performance Reports

DHCPv4 Usage Statistics

The *DHCPv4 Usage Statistics* report provides the overall DHCPv4 usage in a given time frame. The default report includes all network views, all members, all subnets, all IPv4 addresses, and all DHCP ranges, and the default time frame is the last hour. The table is sorted by DHCP utilization rate.

This report displays the following information in table format:

- **Timestamps:** The date and time of the event.
- **Network View:** Filter by a specific network view.
- **Network:** The network address.
- **CIDR:** The subnet mask in CIDR format.
- **DHCPv4 Utilization:** The percentage of DHCP address in use over the total number of DHCP addresses provisioned.
- **Ranges:** The total number of IP address ranges in the network.
- **Provisioned:** The total number of DHCP addresses configured.
- **Dynamic:** The number of dynamic DHCP leases issued.
- **Static:** The number of static DHCP addresses configured.
- **Free:** The number of free DHCP addresses.
- **Used:** The total number of DHCP addresses in use.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by a specific member.
- **Network View:** Filter by a specific network view.
- **Network:** Filter by a specific network.
- **CIDR:** Filter by the subnet mask in CIDR format.
- **Utilization Rate:** Filter by the percentage of DHCP address in use over the total number of DHCP addresses provisioned.
- **Microsoft Servers:** specify Microsoft servers assigned to networks and DHCP ranges. This filter is available even if no MS Management license installed on GM and Grid members.

DHCPv4 Range Utilization Trend

The *DHCPv4 Range Utilization Trend* report provides DHCP usage trends for the top five most utilized address ranges in a given time frame. The default report includes the top five most utilized DHCP ranges among all network views, all members, all subnets, and all IPv4 addresses.

The default report displays line graphs for the top five most utilized address ranges and shows their DHCPv4 usage trends over the last 24 hours. Each of the five address ranges is represented with a different color line graph.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **Network:** Filter by IPv4 addresses.
- **DHCP Ranges:** Filter by a specific DHCP range. The default is the top five most utilized address ranges.
- **Microsoft Servers:** specify Microsoft servers assigned to networks and DHCP ranges. This filter is available even if no MS Management license installed on GM and Grid members.

DHCPv4 Usage Trend

The *DHCPv4 Usage Trend* report provides the overall DHCP usage trend for all members in a given time frame. The default report includes information about all DHCP ranges in all network views, all members, all subnets, and all IPv4 addresses. It displays line graphs for the dynamic, static, and free DHCPv4 leases and shows their DHCPv4 usage trends over the last 24 hours. Each of the DHCPv4 leases is represented with a different color line graph.

You can also select to display this report in table format with the following information:

- **Time:** The timestamp of the event.
- **Dynamic:** The number of dynamic DHCP leases issued.
- **Static:** The number of static DHCP addresses configured.
- **Free:** The number of free DHCP addresses.

Each of the line graphs is represented with a different color.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **DHCP Ranges:** Filter by a specific DHCP range.
- **Network:** Filter by IPv4 networks.

When you select more than one member as the filter criteria, the report displays line graphs for each of the following data: Dynamic, Static, and Free, for each selected member.

- **Microsoft Servers:** specify Microsoft servers assigned to networks and DHCP ranges. This filter is available even if no MS Management license installed on GM and Grid members.

DHCP Message Rate Trend

The *DHCP Message Rate Trend* report provides the overall DHCP message rate trends for DHCP message types in a given time frame. The default report displays the actual, maximum, average, and minimum rate trends in the last 24 hours for the following message types: DHCPDISCOVER, DHCPOFFER, DHCPREQUEST, and DHCPACK.

You can also select to display this report in table format with the following information:

- **Time:** The timestamp of the event.
- **DHCPDISCOVER:** The actual rate trend of the DHCPDISCOVER messages.
- **DHCPOFFER:** The actual rate trend of the DHCPOFFER messages.

- **DHCPREQUEST:** The actual rate trend of the DHCPREQUEST messages.
- **DHCPACK:** The actual rate trend of the DHCPACK messages.

Each of the line graphs is represented with a different color.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **DHCP Message Type:** Filter by a specific DHCP message type. The default includes all message types.
- **Member:** Filter by all members or a specific member.
- **Protocol:** Filter by IPv4 or IPv6 addresses.
- **Statistics:** Filter by minimum, average, or maximum rate. The default is None, which means the report displays only the actual rate trend.

DNS Performance

DNS Query Rate by Query Type

The *DNS Query Rate by Query Type* report shows the trend of DNS queries per second by DNS record type. This report displays line graphs of DNS query trends for selected DNS record types over a given time frame.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Member:** Filter by all members or a specific member.
- **Record Type:** Filter by a specific record type.

DNS Query Rate by Server

The *DNS Query Rate by Server* report shows the trend of DNS queries for selected members. This report displays line graphs of DNS query trends for the selected members over a given time frame.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.

DNS Daily Query Rate by Server

The *DNS Daily Query Rate by Server* report shows the trend of the average or maximum daily DNS Query rate by member. This report can help you identify the average or maximum daily load that is being carried by each DNS Server. This can report help you plan better for capacity and reduce the risk of overloading DNS devices.

This report displays the following information:

- **Time:** Timestamp of events.
- **QPS:** Query per second. QPS is calculated by the max/avg of 24 hourly QPS data per day (between midnights)."

You can use the following filters to get specific information in this report:

- **Time:** Specify time span to load events.
- **Start Time:** Specify start time.
- **End Time:** Specify end time.
- **Members:** Specify members for which QPS need to be displayed.

- **Stats:** Specify what statistic value needs to be displayed [Max, Average].
 - GUI will throw error if user attempts to specify a time span smaller than 1 week (7 days).
 - There will be a single event reported per server per day.
 - Behind the scene, summary fill search generates summary events for this search once a day, and each time it collects raw events from the beginning of yesterday to the beginning of today.

DNS Daily Peak Hour Query Rate by Server

The *DNS Daily Peak Hour Query Rate by Server* report shows the average or peak DNS Query rate at the busiest hour within a day. This report will help you identify the load that is being carried by each DNS Server during busy hours. This report can help you plan better for capacity and reduce the risk of overloading DNS devices.

This report displays the following information

- **Time:** Timestamp of events.
- **QPS:** Query per second. QPS is calculated with two steps: 1) find out the busiest hour (on the top of hours such as from 8:00am to 9:00am) by average hourly QPS, and 2) use that hour's max/avg QPS as they daily max/avg QPS.

You can use the following filters to get specific information in this report:

- **Time:** Specify time span to load events.
- **Start Time:** Specify start time.
- **End Time:** Specify end time.
- **Members:** Specify members for which QPS need to be displayed.
- **Members:** Specify what statistic value needs to be displayed [Max, Average].
 - GUI will throw error if user attempts to specify a time span smaller than 1 week (7 days).
 - There will be a single event reported per server per day.
 - Behind the scene, summary fill search generates summary events for this search once a day, and each time it collects raw events from the top of an hour ago to the top of the current hour.

DNS Response Latency Trend

The *DNS Response Latency Trend* report provides DNS latency response times for all or selected cache servers. This report shows line graphs of DNS latency response times for each server.

You can use filters to get specific information in the report. This report provides the following filters:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by specific members.

DNS Query Reports

DNS Top NXDOMAIN / NOERROR (no data)

The *DNS Top NXDOMAIN / NOERROR (no data)* report shows the number of responses transmitted by the specified name server(s) indicating a client-specified non-existent domain name. This report displays horizontal bar graphs of DNS query trends for the selected members over a given time frame.

- NXDOMAIN indicates that no records of any type existed for the query name;
- NOERROR (no-data) indicates that no data existed for the requested resource record type; other records may exist for the query name.

You can use filters to get specific information in this report. Unless otherwise noted, filters use the equal-sign operator (=) for matching. This report provides the following filters:

- **TopN:** Filter by the top set of records in the report, such as “Top 10.” You choose from a set of fixed values for the TopN filter setting: 10, 20, 50, 100, 200 or 500.
- **Reply Type:** choices include **NXDOMAIN** reply or **NOERROR (No Data)**.
- **Time:** Filter by last hour, last day, last month, or last year.
- **Start Time:** Specify a start time. You can apply a **within the last** operator to this filter. Should you select **within the last**, the filter expression changes to reflect “1” and choices **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Example: **within the last 1 hour(s)**. You may edit the numeric value to the desired amount, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can apply a **Within the Last** expression to this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member against which the filter runs.
- **DNS View:** Use this filter to specify the DNS View in the Infoblox Grid against which the report should run.

DNS Top Clients Per Domain

The *DNS Top Clients Per Domain* report lists the clients that have the most DNS queries for specified domain names and their subdomains. The report shows a horizontal bar chart that lists the clients that have the most total counts of DNS requests and their percentages over a given time frame. You can display the report data in bar chart form or in table form. The domain or domains are specified using filters. The default report displays the top 10 clients within the last 24 hours.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **TopN:** Filter by the number of clients that have the top most DNS queries within the domain name(s) and related sub-domains for the report.
- **Time:** Filter by last hour, last day, last month, or last year.
- **Start Time:** Specify a start time. You can apply a **within the last** operator to this filter. Should you select **within the last**, the filter expression changes to reflect “1” and choices **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Example: **within the last 1 hour(s)**. You may edit the numeric value to the desired amount, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can apply a **Within the Last** expression to this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member against which the filter runs.
- **DNS View:** Use this filter to specify the DNS View in the Infoblox Grid against which the report should run.

You define the domains for capture in the Grid Reporting Properties page (**Reporting** tab → **Reports** tab → **Grid Reporting Properties** → **DNS Queries** side tab → **Monitor queries made to the following domains** check box).

DNS Top SERVFAIL Errors Received

The *DNS Top SERVFAIL Errors Received* report lists the top queries resulting in Infoblox name servers receiving DNS response packets containing the SERVFAIL message from upstream name servers. You can display the report data in table format. The length of the list of top queries, the time period for the report, and other parameters are specified using filters. The default report displays the top upstream query names within the last 24 hours. (The upstream query name may be a query name supplied by a client, or another name that is needed while processing a client query.) When capturing queries, the Grid member matches recursive queries to generate events for the report. This report displays no DNS client information, or the identities of impacted name servers. This report reflects the exact numeric value of the number of queries.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **TopN:** Filter by the number of upstream query names.
- **Time:** Filter by last hour, last day, last month, or last year.

- **Start Time:** Specify a start time. You can apply a **within the last** operator to this filter. Should you select **within the last**, the filter expression changes to reflect “1” and choices **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Example: **within the last 1 hour(s)**. You may edit the numeric value to the desired amount, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can apply a **Within the Last** expression to this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member against which the filter runs.
- **DNS View:** Use this filter to specify the DNS View in the Infoblox Grid against which the report should run.

DNS Top SERVFAIL Errors Sent

The *DNS Top SERVFAIL Errors Sent* report lists the top query names resulting in Infoblox name servers sending DNS response packets containing the SERVFAIL message, to downstream clients. You can display the report data in table format. The length of the list of top queries, the time period for the report, and other parameters are specified using filters. The default report displays the top 10 query names within the last 24 hours. When capturing queries, the Grid member matches authoritative and recursive queries to generate events for the report. This report displays no DNS client information, or the identities of impacted name servers, when the SERVFAIL originates from an upstream server.

You can use filters to get specific information you want in this report. This report provides the following filters:

- **TopN:** Filter by the number of top query names.
- **Time:** Filter by last hour, last day, last month, or last year.
- **Start Time:** Specify a start time. You can apply a **within the last** operator to this filter. Should you select **within the last**, the filter expression changes to reflect “1” and choices **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Example: **within the last 1 hour(s)**. You may edit the numeric value to the desired amount, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can apply a **Within the Last** expression to this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member against which the filter runs.
- **DNS View:** Use this filter to specify the DNS View in the Infoblox Grid against which the report should run.

DNS Top Timeout Recursive Queries

The *DNS Top Timeout Recursive Queries* report shows the number of queries sent to Infoblox Grid member name servers, that result in timeouts after sending recursive queries to upstream name servers. This report displays horizontal bar graphs of DNS query trends for the selected members over a given time frame.

You can use filters to get specific information in this report. Unless otherwise noted, filters use the equal-sign operator (=) for matching. This report provides the following filters:

- **TopN:** Filter by the top set of records in the report, such as “Top 10.” You choose from a set of fixed values for the TopN filter setting: 10, 20, 50, 100, 200 or 500.
- **Time:** Filter by last hour, last day, last month, or last year.
- **Start Time:** Specify a start time. You can apply a **within the last** operator to this filter. Should you select **within the last**, the filter expression changes to reflect “1” and choices **min(s)**, **hour(s)**, **day(s)**, **week(s)** or **month(s)**. Example: **within the last 1 hour(s)**. You may edit the numeric value to the desired amount, such as **within the last 24 hour(s)**.
- **End Time:** Specify an end time. You can apply a **Within the Last** expression to this filter.
- **Members:** Filter by all members or a specific member. The **Members** filter uses the **contains** operator. You also select the Grid member against which the filter runs.
- **DNS View:** Specify the DNS View in the Infoblox Grid against which the report should run.

DNS Top Requested Domain Names

The *DNS Top Requested Domain Names* report lists the top most requested domain names, their counts and the percentage of request over a given time frame. The report shows horizontal bar charts that list the total counts and request percentage for the top most requested domain names. The default report displays the top 10 domain names within the last 24 hours.

You can use the following filters to get specific information in this report:

- **TopN:** Filter by the number of the top most requested upstream query names.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **TLD:** Filter by the top-level domain, such as .com, .edu, or .org.

DNS Cache Hit Rate Trend

The *DNS Cache Hit Rate Trend* report provides information about the cache hit ratio of selected Grid members. The report shows line graphs that track cache hit rates over a given time frame. Note that if you have one member with two DNS views and requests are sent to only one DNS view, the maximum hit rate is 50% (not 100%) for the member because one DNS view has 100% hit rate and the other has 0, and the average is 50%.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.

DNS Top Clients

The *DNS Top Clients* report lists clients that have the most DNS queries. The report shows horizontal bar charts that list clients that have the most total counts of DNS requests and their percentages over a given time frame. The default report displays the top 10 clients within the last 24 hours.

You can use the following filters to get specific information in this report:

- **TopN:** Filter by the number of clients that have the top most DNS queries.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Member:** Filter by all members or a specific member.

Note: To generate data for DNS Response Latency Trend report, the Grid member enabled for DNS service queries itself for PTR record 1.0.0.127.in-addr.arpa every minute. NIOS will not exclude such DNS queries and displays default client 127.0.0.1 in the DNS Top Clients report.

DNS Replies Trend

The *DNS Replies Trend* report provides information about DNS query trends by message types. The report shows line graphs that track DNS query replies by message type over a given time frame.

This report displays line graphs of DNS query replies by the following query type: Failure, NXDomain, NXRRset, Referral, Success, Refused, and Other.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.

- **End Time:** Specify an end time.
- **Member:** Filter by all members or a specific member.
- **Response Type:** Filter by NXDomain, NXRRset, Referral, Success, and Refused.

DNS Last Queried

DNS Last Queried reports lists the zones and resource records that have not been queried during the past user defined time period (days/weeks/months).

- Network and DNS View columns are hidden by default for Last Queried Reports.
- You cannot schedule or copy the Last Queried reports, nor associate these reports with a search. These reports do not consume index volume.

When a user enables monitoring for *DNS Resource Records Last Queried* or for *DNS Zones Last Queried*, initiates a backup and later restores that data set, NIOS maintains the reporting configuration; the Zones and/or resource record originally specified to be monitored will continue to do so.

DNS Zones Last Queried

The *DNS Zones Last Queried* report lists the zones that have not been queried during the past user defined time period (days/weeks/months). Before you run this report, you must select the zones you want to monitor in the Grid Reporting Properties -> Basic tab -> DNS Queries dialog box.

This report displays the following information

- **Network View:** The network view. You cannot sort on this column. This column is hidden by default.
- **DNS View:** DNS view name. This column is hidden by default.
- **Zone:** FQDN of zone.
- **Monitored Since:** Date monitoring started.
- **Last Queried:** Displays "Not Monitored", "Not Queried Since xxxx", or date of last query.

When multiple values are specified with the same filter, the filter applies *or* logic, e.g. 'a' or 'b'. Other perspectives in NIOS UI apply *and* logic, e.g., 'a' and 'b'. You can use the following filters to get specific information in this report:

- **Not Queried:** Filter by servers that have not been queried since the user defined date.
- **DNS View:** DNS view name.
- **Zone:** FQDN of zone.

Note: If a Grid secondary server uses zone transfer to update zone data from a Grid primary server, NIOS does not monitor queries made to the Grid secondary server and it does not update the last queried timestamp for the zones.

DNS Resource Records Last Queried

The *DNS Resource Records Last Queried* report lists all DNS resource records that have not been queried during the past user defined time period (days/weeks/months). Before you run this report, you must select the zones you want to monitor in the Grid Reporting Properties editor -> DNS Queries dialog box.

Note: Exporting the *DNS Resource Record Last Queried* report may take longer than usual if the report contains a lot of records. Also, If a Grid secondary server uses zone transfer to update zone data from a Grid primary server, NIOS does not monitor queries made to the Grid secondary server and it does not update the last queried timestamp for the resource records in a zone.

This report displays the following information:

- **Network View:** Network view name. You cannot sort on this column. This column is hidden by default.
- **DNS View:** DNS view name. This column is hidden by default.
- **Zone:** FQDN of zone.

- **Name:** FQDN of resource record.
- **Record Type:** Resource record type.
- **Record Data:** Value of resource record, such as address of an A record.
- **Monitored Since:** Date monitoring started.
- **Last Queried:** Displays "Not Monitored", "Not Queried Since xxxx", or date of last query.

When multiple values are specified with the same filter, the filter applies *or* logic, e.g. 'a' or 'b'. Other perspectives in NIOS UI apply *and* logic, e.g., 'a' and 'b'. You can use the following filters to get specific information in this report:

- **Not Queried:** Specify a date when the last query was made, The only operator is "Since". If the specified date is within the past one week (including the last 7 days), GUI will display error message saying "Not Queried filter must be specified with a date more than one week ago".
- **DNS View:** DNS view name.
- **Zone:** FQDN of zone.
- **Record Type:** Only a single Record Type filter can be specified. This report shows only the following resource records:
 - A Records
 - AAAA Records
 - BulkHost
 - CNAME Records
 - DNAME Records
 - DS Records
 - Host Address
 - Host Alias
 - Host Record
 - MX Records
 - NAPTR Records
 - NS Records
 - PTR Records
 - Resource Record
 - SRV Records
 - Shared A Record
 - Shared AAAA Record
 - Shared MX Record
 - Shared SRV Record
 - TXT Records
 - Other Records

Note: NIOS does not monitor queries or update timestamp for DNSSEC records, except for DS records. As a result, the DNS Resource Records viewer displays "Not Monitored" in the Last Queried column for all DNSSEC records. In addition, the DNS Not Queried Resource Records report does not display any DNSSEC records. If this report does not display all the resource records that satisfy the filter conditions even after refreshing the page, then navigate to the Reports viewer and relaunch the DNS Resource Records Last Queried report to view all the resource records.

DNS Query Trend per IP Block Group

The *DNS Query Trend per IP Block Group* report provides trend of DNS query counts aggregated over time intervals for user-defined IP block groups.

This report displays the following information in table format:

- **Time:** Timestamp of events
- **Group:** Name of the IP block group.
- **Query Count:** Total queries made to the IP block group for a specific time interval.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last hour, last day, last week, or last month.
- **Start Time:** Specify start time.
- **End Time:** Specify end time.
- **Aggregation Time:** Specify the time span to load events.
- **DNS View:** Filter by a specific DNS view.
- **Members:** Filter by all members or a specific member.
- **Groups:** Specify a group name. The drop-down list displays all the IP block groups configured on the Grid.

IPAMv4 Utilization Reports

IPAMv4 Network Usage Statistics

The *IPAMv4 Network Usage Statistics* report provides usage statistics for each network in a given time frame.

This report displays the following information in table format:

- **Timestamp:** The timestamp when the network container was created.
- **Network View:** The network view.
- **CIDR:** The subnet mask in CIDR format.
- **DHCPv4 Utilization %:** The percentage of DHCP addresses in use over the total number of DHCP addresses provisioned.
- **CIDR:** The subnet mask in CIDR format.
- **Total:** The total number of IPAM addresses in the network.
- **Allocated:** The number of allocated IP addresses in the network.
- **Reserved:** The number of reserved IP addresses in the network.
- **Assigned:** The number of assigned IP addresses in the network.
- **Protocol:** IPv4 or IPv6.
- **Utilization %:** The percentage of IP address in use over the total number of IP addresses in the network.
- **Unmanaged:** The number of discovered IP addresses that do not have corresponding records on the appliance, such as A records, PTR records, fixed address records, host records, or leases.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Network View:** Filter by a specific network views.
- **Utilization Rate:** Filter by the utilization percentage.

DNS Statistics per DNS View

The *DNS Statistics per DNS View* report provides DNS zone statistics for each DNS view in a given time frame. The default report includes information for all network views, all members, all IPv4 and IPv6 reverse-mapping zones, all forward-mapping zones, and all DNS records by record type.

This report displays the following information in table format:

- **Timestamp:** The date and time of the event.
- **View:** The DNS view.
- **Members:** The FQDN of the member that is associated with the DNS view.
- **Forward-Mapping Zone:** The number of forward-mapping zones.
- **IPv4 Reverse-Mapping Zone:** The number of IPv4 reverse-mapping zones.
- **IPv6 Reverse-Mapping Zone:** The number of IPv6 reverse-mapping zones.
- **Signed Zone:** The number of signed zones.
- **Host:** The number of host records.
- **Total Records:** The total number of DNS resource records.

Grid Manager also displays the number of each relevant DNS resource records.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Member:** Filter by specific members.
- **DNS View:** Filter by specific DNS views.

DNS Statistics per Zone

Since every DNS view can have multiple zones and each zone can have multiple records, this report highlights the list of all zones and provides statistics based on every DNS Zone. This report allows you to identify how many and what type of DNS records each zone is serving and use these statistics for more effective planning.

The *DNS Statistics per Zone* report displays the following information:

- **Timestamp:** Timestamp of events.
- **Zone:** FQDN of zone.
- **Function:** Zone function: [Forward-Mapping, IPv4 Reverse-Mapping, IPv6 Reverse-Mapping]
- **Signed:** Boolean to indicate if the zone is signed.
- **Hosts:** Number of hosts.
- **Total Records:** Number of total resource records, Host are not counted.
 - A Records: number of A records.
 - AAAA Records: number of AAAA records.
 - CNAME Records: number of CNAME records.
 - DNAME Records: number of DNAME records.
 - DNSKEY Records: number of DNSKEY records.
 - DS Records: number of DS records.
 - MX Records: number of MX records.
 - NAPTR Records: number of NAPTR records.
 - NSEC Records: number of NSEC records.
 - NSEC3PARAM Records: number of NSEC3PARAM records.
 - NSEC3 Records: number of NSEC3records.
 - NS Records: number of NS records.

- PTR Records: number of PTR records.
- RRSIG Records: number of RRSIG records.
- SOA Records: number of SOA records.
- SRV Records: number of SRV records.
- TXT Records: number of TXT records.
- Other Records: number of other records.

You can use the following filters to get specific information in this report:

- **Time:** Specify time span to load events.
- **Start Time:** Specify start time.
- **End Time:** Specify end time.
- **Grid Primary:** Specify in-Grid primary assigned to zones.
- **Microsoft Primary:** Specify the Microsoft server assigned to DNS zones as the primary. This filter is available even if no Microsoft Management license installed on Grid Master and Grid members.
- **DNS View:** Specify the DNS view the zone belongs to.
- **Zone:** FQDN of zone.
- **Zone Function:** Specify zone function with enum [Forward-Mapping, IPv4 Reverse-Mapping, IPv6 Reverse-Mapping]
- **Signed:** Specify if the zone is signed.

IPAMv4 Top Utilized Networks

The *IPAMv4 Top Utilized Networks* report provides statistics about the top most utilized IPv4 networks. The default report includes the top 10 most utilized networks within the last hour.

This report displays the following information in table format:

- **Timestamp:** The date and time of the recorded utilization.
- **Network View:** The network view.
- **Network:** The network address.
- **CIDR:** The subnet mask in CIDR format.
- **DHCPv4 Utilization %:** The percentage of IP address in use over the total number of IP addresses in the network.
- **Total:** The total number of IP addresses in the network.
- **Assigned:** The total number of IP addresses assigned in the network.
- **Reserved:** The total number of reserved IP addresses in the network.
- **Unmanaged:** The number of discovered IP addresses that do not have corresponding records on the appliance, such as A records, PTR records, fixed address records, host records, or leases.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Member:** Filter by all members or a specific member.
- **TopN:** The number of the top most utilized networks. The default is 10.
- **Microsoft Servers:** Specify Microsoft servers assigned to networks and DHCP ranges. This filter is available even if no MS Management license installed on GM and Grid members.

DHCPv4 Top Utilized Networks

The *DHCPv4 Top Utilized Ranges* report provides statistics about the top most utilized DHCPv4 networks. The default report includes the top 10 most utilized DHCPv4 networks within the last 24 hours.

This report displays the following information in table format:

- **Timestamp:** The date and time of the recorded utilization.
- **Network View:** The network view.
- **Network:** The network address.
- **CIDR:** The subnet mask in CIDR format.
- **DHCPv4 Utilization %:** The percentage of DHCP addresses in use over the total number of DHCP addresses provisioned.
- **Ranges:** The number of DHCP address ranges in the network.
- **Provisioned:** The total number of IP addresses in the range.
- **Dynamic:** The number of dynamic IP addresses in the range.
- **Static:** The number of static IP addresses in the range.
- **Free:** The number of free DHCP addresses.
- **Used:** The total number of IP addresses in use.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Member:** Filter by all members or a specific member.
- **TopN:** The number of the top most utilized networks. The default is 10.

System Utilization Reports

CPU Utilization Trend

The *CPU Utilization Trend* report provides CPU usage trends over a given time frame. The default report displays line graphs that show CPU usage trends for up to five members in the Grid over the last 24 hours. Each of the members is represented with a different color line graph.

You can use the following filters to get specific information in this report:

- **Members:** Filter by specific members.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.

Memory Utilization Trend

The *Memory Utilization Trend* report provides memory usage trends over a given time frame. The default report displays line graphs that show memory usage trends for up to five members in the Grid over the last 24 hours. Each of the members is represented with a different color line graph.

You can use the following filters to get specific information in this report:

- **Members:** Filter by specific members.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.

Traffic Rate

The *Traffic Rate* report provides inbound and outbound traffic over a given time frame. The report displays line graphs that show traffic rate for members with reporting service enabled within the last 24 hours. Grid Manager uses different color line graphs to distinguish inbound and outbound traffic for different members.

Note: The member details are not updated for alert emails in the Traffic Rate report if the rises-by/drops-by operator is used. However, the member details are updated if you use the operators like greater than/ less than in the alert filter.

You can use the following filters to get specific information in this report:

- **Members:** Filter by all members or specific members.
- **Time:** Filter by the last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.

Security

DNS Top RPZ Hits

The *DNS Top RPZ Hits* report lists the top clients who received re-written responses through RPZ. The report displays the total client hits and total rule hits over a given time frame.

The appliance lists the top RPZ hits in table format. You can click a specific row in the table to view the lease history of a client. Grid Manager displays another report that specifies more detailed information, such as the leased IPs, host name, and MAC addresses for each client. For more information about RPZs, see [About Infoblox DNS Firewall](#) on page 1233.

You can compare the domain name and mitigation action in this report with the RPZ rules and mitigation actions in the **FireEye Alerts** report to determine the RPZ hits received due to FireEye alerts.

Note: To enable this report, you must select the **DNS Query** and **Security** check boxes in the **Grid Reporting Properties** editor. To select the check boxes, to the **Reporting** tab -> **Grid Reporting Properties** -> **General** tab -> **Basic** tab -> select the check boxes **DNS Query** and **Security** under **Report Category**.

This report displays the following information in table format:

- **Client ID:** The IP address of the client that queried the domain name that is listed in the RPZ ruleset.
- **Total Client Hits:** The total number of hits received from the respective client.
- **Domain Name:** The domain name that was queried.
- **RPZ Entry:** The RPZ rule that was triggered based on client queries.
- **Total Rule Hits:** The total number of hits received for a specific RPZ rule.
- **Mitigation Action:** The ruleset specified for the blocked domain name or IP address.
- **Substitute Addresses:** The address which was substituted for the blocked domain.
- **Time:** The date and time when the last hit was received.

You can use the following filters to get specific information in this report:

- **TopN:** Filter by the number of top clients.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Client:** The IP address of the client that attempted to query a domain name that is listed in the ruleset defined for RPZ.
- **Domain Name:** The blocked domain name specified in the RPZ ruleset.

- **DNS View:** DNS view name.
- **Members:** Filter by all members or a specific member.
- **Mitigation Action:** The ruleset specified for the blocked domain name or IP address. The value in this column is displayed as:
 - **Passthru**—if you select *Passthru* rule from the drop-down list for **Policy Override** field when you define the rule.
 - **Block (No Such Domain)**—if you select *Block (No Such Domain)* rule from the drop-down list for Policy Override field when you define the rule.
 - **Block (No Data)**—if you select *Block (No Data)* rule from the drop-down list for Policy Override field when you define the rule.
 - **Substitute**—if you select *None (Given)* from the drop-down list for **Policy Override** field when you define a Substitute (Domain Name) rule, as the rule behaves as a CNAME record for RPZ or if the A or AAAA address does not appear in the DNS servers output.
 - **Substitute (A/AAAA)**—if the A and AAAA address only appears in the DNS servers output.
 - **Substitute (A)**—if the A address only appears in the DNS servers output.
 - **Substitute (AAAA)**—if the AAAA address only appears in the DNS servers output.
 - **Substitute (Domain Name)**—if you select *Substitute (Domain Name)* from the drop-down list for Policy Override field when you define a Substitute (Domain Name) rule for the RPZ.
 - **None**—if the DNS resolution does not match any rule.
- **RPZ Zone:** Filter by RPZ zones. When you select an RPZ zone, the operator is set to suffix match by default. This operator performs a suffix match with the label before the RPZ rule name. For example, if you select the value fireeye.com for an RPZ zone, suffix match will match fireeye.com with xia.qisihuisheng.net.fireeye.com.
- **RPZ Entry:** Filter by RPZ rules.

DNS Top RPZ Hits by Clients

The *DNS Top RPZ Hits by Clients* report lists the total number of RPZ hits from a client during an interval, irrespective of the rules and mitigation actions. You can view the IP address of the client, total hits and the date and time during which the hits were received.

The appliance lists the top RPZ hits by clients in table format. You can click a specific row in the table to view the lease history of a client. Grid Manager displays another report that specifies more detailed information, such as the leased IPs, host name, and MAC addresses for each client. For more information about RPZs, see [About Infoblox DNS Firewall](#) on page 1233.

This report displays the following information in table format:

- **Client ID:** The IP address of the client that queried the domain name that is listed in the RPZ ruleset.
- **Total Client Hits:** The total number of hits received from the respective client.
- **Time:** The date and time when the last hit was received.

You can use the following filters to get specific information in this report:

- **TopN:** Filter by the number of top clients.
- **Time:** Filter by last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Client:** The IP address of the client that attempted to query a domain name that is listed in the ruleset defined for RPZ.
- **Domain Name:** The blocked domain name specified in the RPZ ruleset.
- **DNS View:** DNS view name.
- **Members:** Filter by all members or a specific member.
- **Mitigation Action:** The ruleset specified for the blocked domain name or IP address. The value in this column is displayed as:

- **Passthru**—if you select *Passthru* rule from the drop-down list for **Policy Override** field when you define the rule.
- **Block (No Such Domain)**—if you select *Block (No Such Domain)* rule from the drop-down list for Policy Override field when you define the rule.
- **Block (No Data)**—if you select *Block (No Data)* rule from the drop-down list for Policy Override field when you define the rule.
- **Substitute**—if you select *None (Given)* from the drop-down list for **Policy Override** field when you define a Substitute (Domain Name) rule, as the rule behaves as a CNAME record for RPZ or if the A or AAAA address does not appear in the DNS servers output.
- **Substitute (A/AAAA)**—if the A and AAAA address only appears in the DNS servers output.
- **Substitute (A)**—if the A address only appears in the DNS servers output.
- **Substitute (AAAA)**—if the AAAA address only appears in the DNS servers output.
- **Substitute (Domain Name)**—if you select *Substitute (Domain Name)* from the drop-down list for Policy Override field when you define a Substitute (Domain Name) rule for the RPZ.
- **None**—if the DNS resolution does not match any rule.

FireEye Alerts

The *FireEye Alerts* report lists the FireEye alerts that are received by the NIOS appliance. The report displays the date and time when the alert was generated, mitigation action for the alert, ruleset specified for the blocked domain or IP address, and the name of the FireEye appliance that generated the alert. For more information about FireEye integrated RPZs, see [Configuring FireEye RPZs](#) on page 1254.

Note: To enable this report, you must select the **Security** check box in the **Grid Reporting Properties** editor. To select the check boxes, go to the **Reporting** tab -> **Grid Reporting Properties** -> **General** tab -> **Basic** tab -> select the check box **Security** under **Report Category**. Note that you can receive this report only on the Grid Master, not on Grid members, even if you have selected **Security** as a report category on the members.

This report displays the following information in table format:

- **Time:** The date and time when the alert was generated.
- **Alert ID:** The alert type along with the alert ID.
- **Log Severity:** The severity of the alert, which can be **Critical**, **Major** or **Minor**.
- **Alert Type:** The type of alert received from the FireEye appliance.
- **FireEye Appliance:** The FireEye appliance that generated the alert.
- **RPZ Entry:** The RPZ rule specified for the FireEye alert.
- **Mitigation Action:** The ruleset specified for the blocked domain name or IP address.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Alert ID:** Filter by alert ID number.
- **Log Severity:** Filter by severity of alert, which can be **Critical**, **Major** or **Minor**.
- **Alert Type:** Filter by FireEye alert type. The values in this drop-down list are **Infection Events**, **Web Infection**, **Malware Object**, **Domain Match**, and **Callback Events**. For more information about FireEye alerts, see [Configuring FireEye RPZs](#) on page 1254.
- **FireEye Appliance:** Specify the FireEye appliance identification.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Mitigation Action:** The ruleset specified for the blocked domain name or IP address. The value in this column is displayed as:
 - **Passthru**—if you select *Passthru* rule from the drop-down list for **Policy Override** field when you define the rule.

- **Block (No Such Domain)**—if you select *Block (No Such Domain)* rule from the drop-down list for Policy Override field when you define the rule.
- **Block (No Data)**—if you select *Block (No Data)* rule from the drop-down list for Policy Override field when you define the rule.
- **Substitute**—if you select *None (Given)* from the drop-down list for **Policy Override** field when you define a Substitute (Domain Name) rule, as the rule behaves as a CNAME record for RPZ or if the A or AAAA address does not appear in the DNS servers output.
- **Substitute (A/AAAA)**—if the A and AAAA address only appears in the DNS servers output.
- **Substitute (A)**—if the A address only appears in the DNS servers output.
- **Substitute (AAAA)**—if the AAAA address only appears in the DNS servers output.
- **Substitute (Domain Name)**—if you select *Substitute (Domain Name)* from the drop-down list for Policy Override field when you define a Substitute (Domain Name) rule for the RPZ.
- **None**—if the DNS resolution does not match any rule.
- **RPZ Entry:** Filter by RPZ rules.

Threat Protection Event Count By Severity Trend

The *Threat Protection Event Count By Severity Trend* report provides event count trends by severity in a given time frame. You can view event counts distributed for the following severity levels: Critical, Major, Warning and Informational. Each of the severity level of an event is represented with a different color.

You can also define alerts in this report to notify administrators when a trend reaches a specified threshold. For information about how to define alerts, see [About Alerts](#) on page 1140.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **Log Severity:** Filter by severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Category:** Filter by rule category, which can be UDP Floods, ICMP Floods, DNS Attacks, and so on.
- **Rule ID:** Filter by rule ID based on the event category selection.

Threat Protection Event Count By Member Trend

The *Threat Protection Event Count By Member Trend* report provides event count trends on members that supports Advanced DNS Protection in a given time frame. This report tracks events on a member over a given time frame in table format. The default report displays the top 5 appliances in descending order.

You can use the following filters to get specific information in this report:

- **TopN:** Filter by the number of members that have the top most security event counts.
- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **Log Severity:** Filter by severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Category:** Filter by event category, which can be UDP Floods, ICMP Floods, DNS Attacks, and so on.
- **Rule ID:** Filter by unique rule ID.

Threat Protection Event Count By Rule

The *Threat Protection Event Count By Rule* report displays event counts based on violations of individual rules. The appliance displays event count by rule in table format and sorts the records by Total Event Count in descending order. You can click a specific Security ID in the table to view sub-report for the individual rule, showing aggregate event instances with timestamps for a specific rule on all members.

This report displays the following information in table format:

- **SID:** The unique rule ID.
- **Category:** The category to which the rule belongs.
- **Log Severity:** The severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Event Name:** The name of a rule.
- **Alert Count:** The alert count of an event.
- **Drop Count:** The drop count of an event.
- **Total Event Count:** The total number of event counts triggered by a match against the rule.

The sub-report *Threat Protection Event Count for Rule* displays the following information in table format:

Note: The sub-report *Threat Protection Event Count for Rule* displays all the detected events for a specific SID on all members, regardless of the filters you apply to the parent *Threat Protection Event Count By Rule* report.

- **Time:** The timestamp of an event.
- **Member:** The name of the member that supports threat protection.
- **Category:** The category to which the rule belongs.
- **Log Severity:** The severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Event Name:** The name of a rule.
- **Alert Count:** The alert count of an event.
- **Drop Count:** The drop count of an event.
- **Total Event Count:** The total number of event counts triggered by a match against the rule.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **Log Severity:** Filter by severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Category:** Filter by event category, which can be UDP Floods, ICMP Floods, DNS Attacks, and so on.
- **Rule ID:** Filter by unique rule ID.

Threat Protection Event Count By Time

The *Threat Protection Event Count By Time* report displays event counts with timestamp in table format. This report help you track security events behavior based on time of occurrence. For example, this report indicates whether security events peak at specific times or if it has steadily increase over time.

This report displays the following information in table format:

- **Time:** The timestamp of an event.
- **SID:** The unique rule ID.
- **Member:** The name of the member that supports threat protection.
- **Category:** The category to which the rule belongs.
- **Log Severity:** The severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Event Name:** The name of the rule.

- **Alert Count:** The alert count of an event.
- **Drop Count:** The drop count of an event.
- **Total Event Count:** The total number of event counts of a rule.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **Log Severity:** Filter by severity of an event, which can be **Critical**, **Major**, **Warning**, or **Informational**.
- **Category:** Filter by event category, which can be UDP Floods, ICMP Floods, DNS Attacks, and so on.
- **SID:** Filter by unique rule ID.

Threat Protection Event Count By Category

The *Threat Protection Event Count By Category* report provides event counts by rule category. You can track rule categories that are under the most pressure from adverse events. This report displays event counts in table format.

This report displays the following information in table format:

- **Category:** The category to which a rule belongs.
- **Critical Event Count:** The number of critical events in the selected rule category.
- **Major Event Count:** The number of major events.
- **Warning Event Count:** The number of warning events.
- **Informational Event Count:** The number of informational events.
- **Total Event Count:** The total number of event counts triggered against a rule category.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.
- **Category:** Filter by event category, which can be UDP Floods, ICMP Floods, DNS Attacks, and so on.

Threat Protection Event Count By Member

The *Threat Protection Event Count By Member* report provides event counts aggregated over time intervals for each member. This report displays event count for each member in table format and sorts the records by Total Event Count in descending order.

This report displays the following information in table format:

- **Member:** The name of the member that supports threat protection.
- **Critical Event Count:** The number of critical events on a member.
- **Major Event Count:** The number of major events detected on a member.
- **Warning Event Count:** The number of warning events detected on a member.
- **Informational Event Count:** The number of informational events detected on a member.
- **Total Event Count:** The total number of event counts detected on a member.

You can use the following filters to get specific information in this report:

- **Time:** Filter by last minute, last hour, last day, last week, last month, or last year.
- **Start Time:** Specify a start time.
- **End Time:** Specify an end time.
- **Members:** Filter by all members or a specific member.

- **Category:** Filter by event category, which can be UDP Floods, ICMP Floods, DNS Attacks, and so on.
- **Rule ID:** Filter by unique rule ID.

MANAGING REPORTS

You can do the following to manage predefined and user-defined reports in the **Reporting** tab:

- Print predefined and user-defined reports, as described in [Printing Reports](#).
- Back up the reporting database manually, as described in [Backing Up Reporting Data](#).
- Schedule the backup of the reporting database, as described in [Scheduling the Backup of the Reporting Database](#) on page 1173.
- Restore the reporting database on the reporting server, as described in [Restoring the Reporting Database](#) on page 1174

Printing Reports

1. From the **Reporting** tab -> **Reports** tab, click a report name.
2. In the Reporting Dashboard, click the Print icon.
Grid Manager displays the list of reports in another window.
3. Click **Print**.

Backing Up Reporting Data

Before you back up the reporting database, ensure that the reporting service is enabled on the reporting server. You cannot perform or schedule a backup if the reporting service is disabled on the reporting server. If you want to upgrade your reporting server, back up all the data. If you want to upgrade your reporting server, back up all the data before you power down the server. During an upgrade, the reporting server is automatically upgraded after the Grid Master. You cannot control or schedule when to upgrade the reporting server. For information about upgrades and upgrade groups, see [Managing Upgrade Groups](#) on page 408.

Note that reporting data backups are incremental backups, which means that backup files are copied to the designated file server only when there are new events generated since the last backup.

You can manually back up the reporting database or schedule a backup, but you cannot perform both at the same time.

You can perform the following reporting data backups:

- Manual backups, as described in [Backing Up the Reporting Database Manually](#).
- Scheduled backups, as described in [Scheduling the Backup of the Reporting Database](#) on page 1173.

Backing Up the Reporting Database Manually

1. From the **Grid** tab, select **Backup** -> **Reporting Backup** -> **Manual Backup** from the Toolbar.
2. In the *Manual Reporting Backup* editor, complete the following:
 - **Status:** Displays the status of the backup process, if in progress.
 - **Backup to:** Select the destination of the backup file from the drop-down list:
 - **FTP:** Back up the reporting database to an FTP server.
 - **Filepath:** Enter the directory path. For example, you can enter `/archive/backups/Infoblox/` on a Linux server, or `c:\archive\backups\Infoblox\` on a Microsoft Windows server.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Username:** Enter the username of your FTP account.
 - **Password:** Enter the password of your FTP account.

- **SCP:** Back up the reporting database to an SSH server that supports SCP.
 - **Filepath:** Enter the directory path. For example, you can enter `/archive/backups/Infoblox/` on a Linux server, or `c:\archive\backups\Infoblox\` on a Microsoft Windows server.
 - **IP Address of SCP Server:** The IP address of the SCP server.
 - **Username:** Enter the username of your SCP account.
 - **Password:** Enter the password of your SCP account.

Note: When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

Scheduling the Backup of the Reporting Database

1. From the **Grid** tab, select **Backup -> Reporting Backup -> Schedule Backup** from the Toolbar.
2. In the *Schedule Reporting Backup* editor, complete the following:

- **Status:** Displays the status of the backup process, if in progress.

Select the destination of the backup file from the **Backup to** drop-down list:

- **FTP:** Back up the reporting database files to an FTP server.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Directory Path:** Enter the directory path. For example, you can enter `/archive/backups` on a Linux system, or `c:\archive\backups` on a Microsoft Windows system. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
 - **Username:** Enter the username of your FTP account.
 - **Password:** Enter the password of your FTP account.
 - **Recurrence:** Select how often you want to back up the files. You can select **Weekly**, **Daily**, or **Hourly** from the drop-down list. When you select **Weekly**, complete the following:
 - **Every:** Choose a day of the week from the drop-down list.
 - **Time:** Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list. The Grid Master creates a backup file on the selected day and time every week.

When you select **Daily**, enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.

When you select **Hourly**, complete the following:

- **Minutes after the Hour:** Enter the minute after the hour when the Grid Master creates a backup file. For example, enter 5 if you want the Grid Master to create a backup file five minutes after the hour every hour.
- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now, but want to save the settings for future use.
- **SCP:** Back up the reporting database to an SSH server that supports SCP.
 - **IP Address of SCP Server:** The IP address of the SCP server.
 - **Directory Path:** Enter the directory path of the file. For example, you can enter `/archive/backups` on a Linux system, or `c:\archive\backups` on a Microsoft Windows system. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
 - **Username:** Enter the username of your SCP account.
 - **Password:** Enter the password of your SCP account.
 - **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see the FTP section.

- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now. You can still save the settings for future use.

Note: When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

Restoring the Reporting Database

Restoring the reporting database may take a long time to perform, and the reporting service is unavailable during a restore. Ensure that you must restore the reporting database before you perform the operation.

Note the following during a restore:

- The reporting service is unavailable.
- Existing reporting data is removed from the reporting server.
- Backup data is restored up to the amount the reporting server can accommodate.

Note: The Volume Used Today displayed in the Device Information section will not be updated after restoring the data. Also, when you restore data or execute the CLI command `reset reporting_data`, the volume violation count will be reset to zero on the second day.

1. From the **Grid** tab, select **Restore** -> **Restore Reporting** from the Toolbar.
2. In the *Restore* dialog box, complete the following:

- **Status:** Displays the status of the restore process, if in progress.

Select the destination of the backup file from the **Restore from** drop-down list:

- **FTP:** Restore the reporting backup files from an FTP server.
 - **Filepath:** Enter the directory path. For example, you can enter `/archive/backups/Infoblox/` on a Linux server, or `c:\archive\backups\Infoblox\` on a Microsoft Windows server.
 - **IP Address of FTP Server:** The IP address of the FTP server.
 - **Username:** Enter the username of your FTP server account.
 - **Password:** Enter the password of your FTP server account.
- **SCP:** Restore the reporting backup files from a SCP server.
 - **Filepath:** Enter the directory path. For example, you can enter `/archive/backups/Infoblox/` on a Linux server, or `c:\archive\backups\Infoblox\` on a Microsoft Windows server.
 - **IP Address of SCP Server:** The IP address of the SCP server.
 - **Username:** Enter the username of your SCP server account.
 - **Password:** Enter the password of your SCP server account.

3. Click **Restore**.



PART 8 GLOBAL LOAD BALANCER INTEGRATION

This section describes the Infoblox GLB (Global Load Balancer) integration solution. It describes how you can centrally manage GLBs and their associated objects through Grid Manager. You can synchronize integrated GLB data to the Grid, and then use the Infoblox IPAM tools to facilitate permission and data management. This section includes the following chapters:

- [Chapter 39, *Managing Global Load Balancers*](#), on page 1177
- [Chapter 40, *Managing Global Load Balancer Data*](#), on page 1193



Chapter 39 Managing Global Load Balancers

This chapter explains how to configure the Infoblox Grid to integrate with GLBs (Global Load Balancers) and how to manage the GLBs and load balancer synchronization groups through Grid Manager. It includes the following sections:

- [*Integrating Global Load Balancers*](#) on page 1178
 - [*About Load Balancer Synchronization Groups*](#) on page 1180
 - [*Requirements and Permissions*](#) on page 1181
 - [*Deployment Guidelines*](#) on page 1182
- [*Configuring the Management of GLBs*](#) on page 1183
 - [*Setting Usernames and Permissions on GLBs*](#) on page 1183
 - [*Setting the Management Mode*](#) on page 1183
 - [*Configuring the Synchronization Interface*](#) on page 1184
 - [*Assigning Grid Members to Load Balancers*](#) on page 1184
 - [*Validating Load Balancer Connection*](#) on page 1186
- [*Managing Global Load Balancers*](#) on page 1187
 - [*Setting Load Balancer Properties*](#) on page 1187
 - [*Modifying Load Balancer Synchronization Group Properties*](#) on page 1188
 - [*Setting Priorities for Managed Load Balancers*](#) on page 1188
 - [*Changing the Managing Member or Management Mode*](#) on page 1189
 - [*Replicating DNS Data*](#) on page 1189
 - [*Backing Up Synchronized Data*](#) on page 1189
 - [*Disabling Synchronization*](#) on page 1189
 - [*Removing a Load Balancer*](#) on page 1190
 - [*Removing a Load Balancer Synchronization Group*](#) on page 1190
- [*Monitoring Global Load Balancers*](#) on page 1190
 - [*Viewing Global Load Balancers*](#) on page 1191
 - [*Viewing Detailed Status Information*](#) on page 1192
 - [*Viewing Detailed Status Information of a Synchronization Group*](#) on page 1192

INTEGRATING GLOBAL LOAD BALANCERS

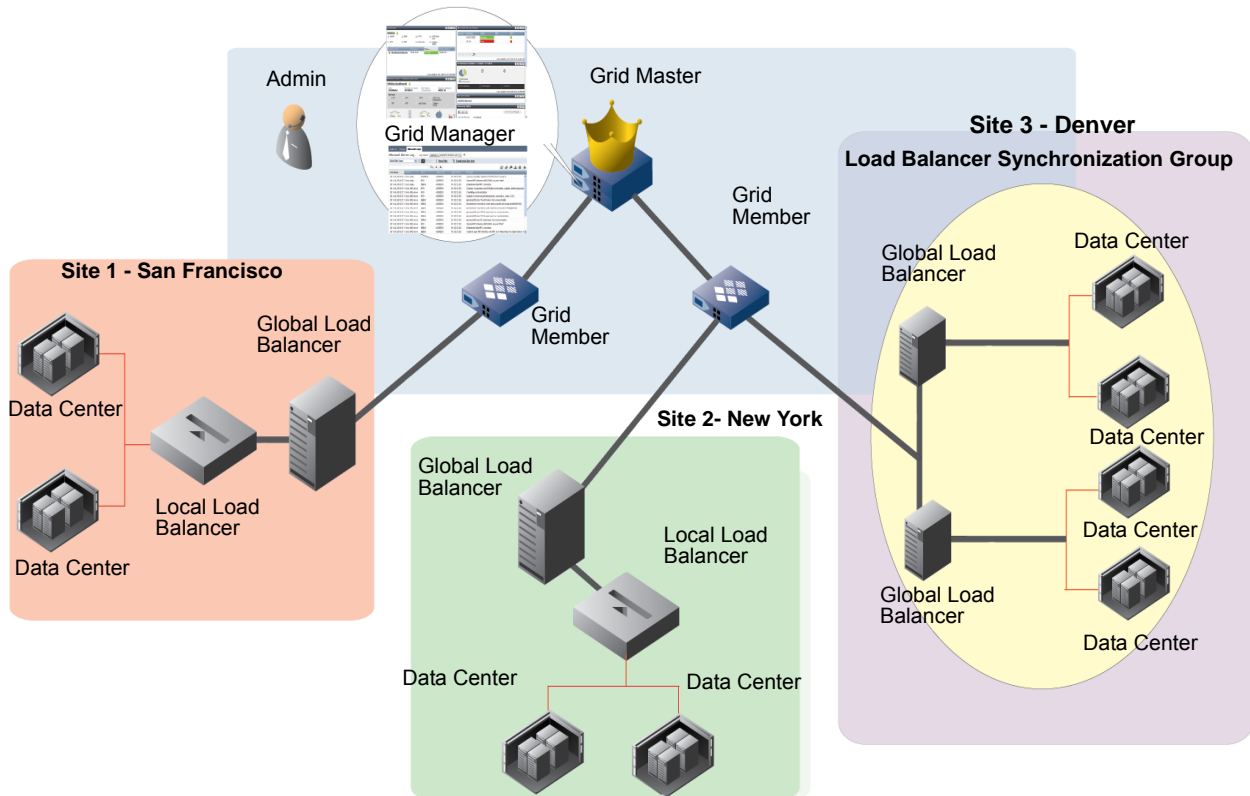
You can configure the Infoblox Grid to integrate with the F5® GTMs (Global Traffic Managers™) and load balancer synchronization groups that provide load-balancing services among multiple data centers. For more information about load balancer synchronization groups, see [About Load Balancer Synchronization Groups](#) on page 1180. Most GLB (Global Load Balancer) solutions use DNS to inspect users' IP addresses and direct them to the most efficient data centers for content delivery. By responding to client queries with the best available IP addresses, GLBs can distribute workload across multiple data centers, networks, and other resources, to achieve optimal resource utilization, maximize throughput, minimize response time, and avoid data traffic overload.

The Infoblox GLB integration solution enables the Infoblox Grid to manage GLBs, load balancer synchronization groups, and their associated objects and data. The solution provides the following:

- An easy-to-use and centralized interface (Grid Manager) for managing your GLBs, load balancer synchronization groups, and related objects. For information about Grid Manager, see [About Grid Manager](#) on page 48.
- Configuration of LBDN (Load Balanced Domain Name) and other supported objects with the Infoblox DNS service. For more information, see [Managing Global Load Balancer Data](#) on page 1193.
- Extension of the current permission models to support the newly added GLB objects and the flexibility in delegating tasks to different admins. For information about admin permissions for GLBs, see [Administrative Permissions for Load Balancers](#) on page 216.
- Associating Extensible Attribute meta data with GLB related objects to enable Smart Folders, search and filters in Grid Manager. For information about extensible attributes, see [About Extensible Attributes](#) on page 322.
- Grid Manager provides a hierarchical map view that you can use to quickly capture an overall traffic management structure of a selected GLB object. For information, see [Viewing Traffic Management Structures](#) on page 1197.

[Figure 39.1](#) illustrates an Infoblox Grid that includes a Grid Master and two Grid members that connect to two independent GLBs in two different sites and to a synchronization group in the third site. The GLBs manage LLBs (Local Load Balancers) that provide load-balancing services of data traffic among multiple data centers. The NIOS admin, with appropriate permissions, can centrally manage the GLBs, load balancer synchronization groups and their associated data and objects from a single interface, Grid Manager.

Figure 39.1 Managing GLBs and Load Balancer Synchronization Groups through the Grid Master



You do not have to configure or install any application on the GLBs for Grid members to communicate with them. Infoblox uses SOAP (Simple Object Access Protocol) to communicate and synchronize data with the GLBs and load balancer synchronization groups. You only need to add usernames and passwords on the GLBs so the Grid members can log in to them.

A Grid member synchronizes data with the GLBs, so you can view, and optionally manage the data from Grid Manager. Note that the Grid Master performs all the write operation to the GLBs. You can manage the synchronized data and manage it in two modes, read-only or read/write. In read-only mode, the Grid member synchronizes data from the GLBs to the Grid so you can use Grid member to view the synchronized data, but not update it. Read/write mode allows you to view and modify the synchronized data as well.

Configuration changes and data synchronized from the Grid Master to the GLBs apply immediately after the synchronization. Synchronization from GLBs is done periodically by the assigned synchronization Grid member. DNS objects from the Grid Master and GLBs are synchronized on the synchronization member. You do not have to restart the GLBs or the DNS service. Synchronization happens automatically every two minutes by default. You can configure this time interval through Grid Manager. For information, see [Assigning Grid Members to Load Balancers](#) on page 1184.

For information about how to set up Grid members and integrate GLBs, see [Configuring the Management of GLBs](#) on page 1183.

You can use Smart Folders to organize your data, and monitor your networks and load balancers from the Dashboard. For information about the Dashboard widget, see [Load Balancer Status](#) on page 133.

About Load Balancer Synchronization Groups

Infoblox supports dynamic grouping of GLBs that are in the same synchronization group. Synchronization groups are created by enabling synchronization on the GLBs. For information about how to create load balancer synchronization groups, refer to the appropriate GLB documentation.

GLBs in the same synchronization group replicate data across peer GLBs and share the same configuration. For more information, see [Data Replication](#) on page 1180. When the NIOS appliance synchronizes data with a load balancer synchronization group, it keeps a single configuration record for the entire group and displays a single instance of the group. This feature facilitates a more efficient data replication process and eliminates multiple copies of the same configuration records of GLBs that are in the same synchronization group.

Data Replication

GLBs in a synchronization group are fully-meshed, and they replicate data by exchanging configuration files with peer GLBs. Sharing the most updated configuration files among peers involves time synchronization, and replication among GLBs and from GLBs to NIOS is not instantaneous. As a result, data mismatch during replication can occur because grouping of GLBs is dynamic. In addition, concurrent configuration changes made on two different GLBs in the same synchronization group may overwrite each other. To avoid this issue, NIOS first synchronizes data with one of the GLBs (a lead device) in the synchronization group, which then replicates the data to other peer GLBs. Note that the lead device is not inherited from the synchronization group and it is not reflected in the configuration. Although NIOS cannot discover the lead device, you can configure one by setting priorities for the GLBs in the synchronization group. For information about how to set priorities to define the lead device, see [Setting Priorities for Managed Load Balancers](#) on page 1188.

When you change the lead device from one GLB to another, the group configuration data on NIOS may change. Note that if replication between the old and new lead device is not completed successfully, there can be permanent loss of NIOS specific data, such as extensible attributes, for objects that have been removed.

You cannot create, add, or modify load balancer synchronization groups on NIOS because configuration data is replicated directly from the GLBs. Note that load balancer synchronization groups replicate only the global load balancer configuration data. They do not replicate LLB configuration and device-specific configuration, such as Listeners, VLANs, or network configuration. DNS replication among devices in a group is optional. For more information, see [Replicating DNS Data](#) on page 1189.

Grouping Load Balancers on NIOS

To manage load balancer synchronization groups, you must configure each GLB that you wish to manage on the NIOS appliance. The appliance will then dynamically group the GLBs in the same group based on the configured information. The appliance does not automatically discover GLBs of a synchronization group. Grouping on NIOS is done by determining the configuration and connectivity of the GLBs. The appliance checks for group names, IP addresses, and network connectivity. Automatic regrouping occurs when NIOS notices connectivity or configuration changes. All GLBs in a synchronization group must actively synchronize configuration changes to remain in the group. They must stay connected to each other to receive configuration changes.

Note the following when you configure GLBs in the same synchronization group:

- GLBs in the same synchronization group must have the same group name and same configurations. Synchronization groups with the same group name but different configuration belong to different groups.
- Grouped GLBs are considered as independent groups if the grouped GLBs are not connected to each other.
- If synchronization is not enabled on a GLB, it does not belong to any synchronization group. NIOS manages these GLBs independently.
- When synchronization is disabled on a GLB that was previously grouped, then the GLB is removed from the group and becomes independent.
- Regrouping indicates if there are additional GLBs beyond the ones that are known to NIOS. However, the appliance cannot auto-discover additional GLBs.
- Grouped GLBs that are disabled or unreachable are regrouped based on their previously known configuration and connectivity. Regrouping may occur before the next synchronization.

- A synchronization group is identified by a name and this name does not have to be unique. By default, NIOS displays the group name that is inherited from the synchronization group. You can reconfigure the group name on NIOS. Note that this group name is specific to NIOS and will not be updated on the load balancer synchronization group.

Requirements and Permissions

A Grid member must have a Load Balancer license installed to manage load balancers and load balancer synchronization groups. The license allows the member to synchronize data with GLBs. It also activates the tabs, dialog boxes and other elements in Grid Manager that you need to manage load balancers and synchronization groups. For information about licenses, see [Managing Licenses](#) on page 377.

The Load Balancer license has capacity limits. Every enabled load balancer in a synchronization group consumes license capacity. If automatic regrouping adds a load balancer to a synchronization group at license capacity limit, that load balancer will be disabled.

Note that if you do not see the **Traffic Management** tab after you add a member that has a Load Balancer license, restart the Grid Master to view the tab and to manage load balancers in the Grid.

Version Requirements for GLBs and Synchronization Groups

All GLBs in a synchronization group must be running the same TMOS version. If the TMOS version on a GLB changes, it is removed from the synchronization group or the entire group is being upgraded. A group upgrade disables all GLBs in the group and stops on-going synchronization. Validating the connection updates the TMOS version and enables the load balancer synchronization group. If a GLB runs a TMOS version that is different from the group version, that load balancer is considered for regrouping. If all the GLBs in a group have the same TMOS version, but the version is different from the group version, then regrouping occurs and the TMOS version on the GLBs will be updated to match the group version. Otherwise, the GLBs that have a different TMOS version form a new synchronization group.

Infoblox has qualified the following F5® GTM (Global Traffic Manager™) TMOS versions:

Platform	Version
TMOS	11.4.x
TMOS	11.3.x
TMOS	11.1.x
TMOS	10.2.x

Note: Only GLBs with TMOS version 11.4.x, 11.3.x, 11.1.x and 10.2.x are supported for load balancer synchronization groups. NIOS does not group GLBs running TMOS version 10 in a synchronization group (this is a limitation on the load balancer side). These GLBs are managed independently. If GLBs running TMOS 10 belong to a synchronization group, then the configuration data is duplicated for these GLBs. When you delete a GLB object, the duplicated GLB object will also be deleted.

Grid members check the TMOS version before each synchronization. If a GLB reports an unsupported version before a synchronization, the member logs an error and the synchronization fails.

Administrative Permissions

By default, only superusers can configure Grid members to manage load balancers and load balancer groups. Superusers can give limited-access users read-only, read/write, or deny permission to GLBs and their associated objects. Read-only permission allows admins to view the properties and data of a GLB or a load balancer synchronization group from Grid Manager. Write permission is required to configure Grid members to manage GLBs, edit GLB object properties, and view the topology of the network with the integration of GLBs.

Note that to view and manage the data synchronized from load balancers and load balancer synchronization groups, admins must have permissions to the applicable resources. For example, to view LBDNs (Load Balanced Domain Names) synchronized from load balancers, admins must have read-only permission to the LBDN objects; and to edit them, admins need read/write permission to them. For information, see [Administrative Permissions for Load Balancers](#) on page 216.

The administrative permissions on the Grid are different from those on the synchronization groups and load balancers. These permissions are independent of each other and are not synchronized. For information about load balancer permissions, refer to the appropriate GLB documentation.

Deployment Guidelines

Following are some recommendations and considerations when configuring Grid members to manage load balancers and load balancer synchronization groups in the network:

- To maintain configuration and data integrity on GLBs, Infoblox recommends that you do not configure multiple NIOS appliances to synchronize with a single GLB.
- Both the Grid Master and Grid member configured to manage the GLBs and load balancer synchronization groups must have HTTPS access to the management port and IP address on the GLBs.
- Grid members connect to load balancers and load balancer synchronization groups using SOAP and HTTPS over TCP/IP. You must adjust your firewall policies to allow traffic between the managing Grid member and its assigned GLBs and synchronization groups.
- By default, all GLB traffic and data synchronization can originate from any interface, such as LAN1 or MGMT, that has the lowest metric for normal routing. You can change the default source interface, as described in [Configuring the Synchronization Interface](#) on page 1184.
- The appliance does not support CSV import or export for synchronized load balancer objects.
- The appliance retrieves the status of load balancer objects from the load balancer. It does not detect the status.
- Infoblox recommends that you schedule the initial synchronization at a time when your network is less busy, especially if you are synchronizing a large amount of data. In addition, if a load balancer reconnects after being disconnected for a long period of time, it could synchronize a significant amount of data and this could impact the Grid member performance.
- Not all Infoblox platforms support being configured as a GLB synchronization member. Contact Infoblox Technical Support for more information.
- The managing member must be close, in terms of network hops, latency and bandwidth, to the load balancers and the lead device of a load balancer synchronization group to which it connects. This can help reduce the synchronization time and potential retries due to network delays.
- Although a Grid member that manages GLBs and load balancer synchronization groups can run other protocols and services, to optimize performance, Infoblox recommends that you configure one or more members solely for managing load balancers.
- For a load balancer synchronization group, configure each managed load balancer on NIOS and assign them to the same group. Validate certificate information and connection to ensure reachability.
- Wait for the synchronization to complete while NIOS groups GLBs in a synchronization group.

CONFIGURING THE MANAGEMENT OF GLBs

You can configure any Grid member that has a Load Balancer license installed to synchronize data with GLBs. Any data you manage through the Grid is handled by the Grid Master.

When an HA pair manages the GLBs, the active node handles the synchronization. If an HA failover occurs during a synchronization, the failing node immediately aborts the synchronization. The new active node resumes the next synchronization.

Complete the following tasks to integrate a GLB with the Grid:

1. On the GLB, add the username and password that you use to configure the Grid member to synchronize data with the GLB. For information, see [Setting Usernames and Permissions on GLBs](#).
2. On the Grid member, configure the mode for managing the GLB. For information, see [Setting the Management Mode](#) on page 1183.
3. Configure the interface for synchronizing data with the GLB. For information, see [Configuring the Synchronization Interface](#) on page 1184.
4. Assign and configure the Grid member to manage the GLB. For information, see [Assigning Grid Members to Load Balancers](#) on page 1184.

Setting Usernames and Permissions on GLBs

Permissions configured on GLBs and on the NIOS appliance are independent of each other. Infoblox recommends that users who are responsible for synchronization be superusers on the GLB and load balancer synchronization groups. You can use NIOS permissions to limit access for GLB administration.

To enable a Grid member to synchronize data with a GLB, you must do the following on the GLB:

- Create the username and password for the Grid member to connect with the GLB.
- Grant user accounts on the GLB the necessary permissions so admins on the NIOS Grid Master can manage the GLB and GLB related objects accordingly. Infoblox recommends these admins to be superusers.

For information about how to perform these tasks on the GLB, refer to the corresponding GLB documentation. For information about how to set the same credentials on the NIOS Grid member, see [Assigning Grid Members to Load Balancers](#) on page 1184.

Setting the Management Mode

A Grid member can manage a GLB or a load balancer synchronization group in read-only mode, which is the default, or in read/write mode. In read-only mode, the Grid member copies GLB data from the GLBs to the Grid so NIOS admins can view the synchronized data. They cannot update the data or configure any properties, but they can set extensible attributes for GLB objects on NIOS. In read/write mode, NIOS admins are allowed to view and modify synchronized GLB data.

NIOS admins can modify and delete group-specific data when the lead device and synchronization group are in read/write mode. Note that the management mode of a GLB is separate from the admin permissions that the NIOS appliance requires to access the GLB and GLB resources. An admin must still have the applicable permissions to the GLBs and GLB resources they want to access. For information about admin permissions, see [About Administrative Permissions](#) on page 160.

For information about how to set the management mode on the Grid member, see [Changing the Managing Member or Management Mode](#) on page 1189.

Configuring the Synchronization Interface

For the Grid member to synchronize data with the GLB, you can configure the synchronization interface you want to use to connect and synchronize data with the GLB. Note that this configuration applies to both the Grid Master and Grid member.

1. From the **Grid** tab -> **Grid Manager** tab, select Grid Properties -> Edit from the Toolbar.
2. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab, and then complete the following:
 - **Load Balancing Source:** From the drop-down list, select the preferred load balancing interface (Any, Address, LAN1, LAN2, MGMT, or VIP) you want the appliance to use to connect and synchronize data with the load balancer. When you select **Address**, you must enter the IP address of the interface. The default is **Any**, which allows the Grid member to communicate with the load balancer through any available interfaces.

Note: If the selected interface fails, the appliance falls back to the default interface, which is **Any**.

- **Load Balancing Source Address:** If you select **Address** as the load balancing source, specify the IP address of the interface you want the appliance to connect and synchronize data with the load balancer.

Assigning Grid Members to Load Balancers

Before a Grid member can manage and synchronize data with a GLB, you must add the GLB and assign a managing member to the GLB. When you add GLBs that are part of a load balancer synchronization group, NIOS automatically groups them together. You can add GLBs of multiple synchronization groups at a time. A GLB is not grouped if synchronization is not enabled. For more information about load balancer synchronization groups, see [About Load Balancer Synchronization Groups](#) on page 1180.

To configure a Grid member to manage GLBs:

1. From the **Grid** tab -> **Load Balancers** tab, click the Add icon.
2. In the *Add load balancer(s)* wizard, complete the following and then click **Next**:
 - **Managing Member:** Specify whether you want a Grid member to manage the GLB. Click **None** to deny assigning a member. Click **Select Member** to select a Grid member.
 - **Synchronization Interval (min):** The default synchronization interval is two minutes. This is the time between the completion of one synchronization and the start of a new one. Synchronizing large data sets could take longer than the synchronization interval, causing a delay in the start of the next synchronization. For example, if the synchronization interval is two minutes but a synchronization takes five minutes, the time between the start of the first synchronization and the start of the next one is approximately seven minutes.
 - **Credentials to Connect to the Load Balancer(s):** Enter the Login name and Password that the appliance uses to connect to the load balancers. These must be the same as those you specified when you created the user account for the Grid member on the GLB.
 - **Manage Load Balancer(s) in:** Select the management mode, which is either **Read-only** or **Read/Write**. In read-only mode, the Grid member copies GLB data to the Grid so NIOS admins can view the synchronized data, but not update the data or configure any properties. In read/write mode, NIOS admins are allowed to update the synchronized GLB data. The default is **Read-only**.
 - **Sync Zones from Load Balancer(s):** Select this check box to synchronize all discovered DNS zones from the GLB to NIOS. To synchronize DNS zones, ensure that you map the GLB DNS view to the corresponding DNS view on NIOS. Otherwise, synchronization of DNS zones does not start. You must also manage the GLB in read/write mode and enable the GLB in order for the DNS zone synchronization to happen. You can deselect this check box to disable the synchronization of zones from load balancers. Note that enabling and disabling zone synchronization from GLBs does not affect the synchronization from NIOS to GLBs. When you add delegated zones on GLBs through Grid Manager, the zones are synchronized to the GLBs even when you disable synchronization here. For information about mapping DNS views, see [Viewing Global Load Balancer Objects](#) on page 1196.

- **Logging Level:** From the drop-down list, select the logging level. The appliance logs events in the syslog. The default is **Debug**.
 - **Debug:** Provides information about all events associated with the synchronization.
 - **Normal:** The Grid member is synchronizing with the GLB and these messages provide normal status information.
 - **Low:** The Grid member synchronized the data, but there was an issue, which is detailed in the Message section. If the Grid member encounters an error during the synchronization, it skips the object with the error, logs the error in the syslog, and continues to synchronize the rest of the data. The Grid member logs the error at each synchronization until you resolve the issue and it can synchronize the object successfully.
 - **High:** The Grid member failed to synchronize an object, such as a listener or an LBDN, due to the error described in the Message section.
- **Comment:** Enter additional information about the GLB.
- **Disable Synchronization:** Select this to disable synchronization between the Grid member and the load balancer. This allows you to preprovision the load balancer and then enable them at a later time.

3. Click **Next** and complete the following:

- **Validate SSL Certificate:** Select this check box to have the Grid member validate the GLB SSL certificate each time it synchronizes with the GLB.

Click the Add icon and Grid Manager adds a row to the Managed Load Balancers table. Complete the following to add GLBs:

- **Device Name:** Enter the name of the GLB.
- **FQDN or IP Address:** Enter either the FQDN or IP address of the GLB. In order for the member to resolve the FQDN of a load balancer, you must define a DNS resolver for the Grid member in the **DNS Resolver** tab of the *Member Properties* editor.

Note: If you have configured multiple load balancers in a single synchronization group, add only one load balancer from a group for synchronization with the NIOS appliance. In addition, ensure that when NIOS is synchronizing with a GLB, other iControl API client does not call `Management.Partition.set_active_partition()`. Because `Management.Partition.set_active_partition()` sets global state for all other clients connected to this GLB.

- **Port:** Enter the port for communication between the member and the GLB. The default is 443.
- **Certificate Status:** Displays the current status of the SSL certificate of the GLB. The status can be a combination of the following:
 - **Trusted:** The certificate is trusted.
 - **Accepted:** The admin has accepted the certificate.
 - **Untrusted:** The certificated is not trusted.
 - **Unaccepted:** The admin has not accepted the certificate.

You can also do the following:

- Test GLB connection with the Grid member and view certificate information by selecting a load balancer check box and click the Validate LB Device Connection icon in the horizontal bar to verify whether the appliance can successfully connect to the GLB. For more information, see [Validating Load Balancer Connection](#) on page 1186.
 - Select a GLB and click the Delete icon to delete it.
4. Click **Next** to define extensible attributes for the GLBs. For information, see [Using Extensible Attributes](#) on page 332.
5. Save the configuration.

After you configure a Grid member to manage a load balancer, the member automatically connects to the load balancer and starts synchronizing data. NIOS automatically groups together the GLBs that are in the same synchronization group. A GLB will not join any group if synchronization is not enabled. NIOS manages these GLBs independently. You can then do the following:

- View the status of the synchronization groups and independent GLBs in the **Grid** tab -> **Load Balancers** tab, as described in [Monitoring Global Load Balancers](#) on page 1190. Newly configured groups and independent GLBs first display the status of **Unknown** as the Grid member contacts the synchronization groups and GLBs. The status changes to **OK** after the Grid member successfully connects to their respective groups and load balancers.
- View the data synchronized from the load balancer. To view and manage load balancer data, navigate to the **Data Management** tab -> **Traffic Management** tab. For more information, see [Chapter 40, Managing Global Load Balancer Data](#), on page 1193.
Network conditions and the amount of data can affect the synchronization time. Therefore, you might not be able to view all of the synchronized data immediately.
- Use Smart Folders to organize the load balancers and their data. For example, you can create a folder for a specific site and another folder for synchronized LBDN objects. For information about Smart Folders, see [Chapter 3, Smart Folders](#), on page 139.
- Use Global Search to search for synchronized data, such as LBDN objects and load balancer pool members. For information, see [Global Search](#) on page 59.
- Perform inline editing for certain fields such as Name, Comment and Site by double-clicking the row that you want to modify. The appliance displays the inline editing editor in the selected row. Click **Save** after modifying the data.

Validating Load Balancer Connection

You must validate the connection of each load balancer in a synchronization group to ensure that the appliance can successfully connect to the GLB. Do the following to validate the connection and view certificate information of a GLB:

1. From the **Grid** tab -> **Load Balancers** tab, select a GLB check box, and then click the Validate LB Device Connection icon.
A dialog that contains the load balancer certificate and GLB information appears. Review the certificate information.
2. Click **Accept** to accept the certificate. The certification information is saved to the NIOS database and the certificate status will appear as **Accepted**. If you do not accept the certificate, the certificate status is **Unaccepted**.

Note: You cannot validate the connection of a load balancer synchronization group. Also, load balancers cannot synchronize with NIOS if their certificates are not self-signed. You must import the signed CA certificates on NIOS so you can synchronize the load balancers successfully.

MANAGING GLOBAL LOAD BALANCERS

After you configure Grid members to manage GLBs, you can set certain properties and manage them as follows:

- Set GLB properties, as described in [Setting Load Balancer Properties](#).
- Change the managing member or the management mode, as described in [Changing the Managing Member or Management Mode](#) on page 1189.
- Back up the synchronized data, as described in [Backing Up Synchronized Data](#) on page 1189.
- Disable synchronization with a load balancer, as described in [Disabling Synchronization](#) on page 1189.
- Remove a GLB, as described in [Removing a Load Balancer](#) on page 1190.

In addition, you can do the following to manage a load balancer group:

- Change certain group properties, as described in [Modifying Load Balancer Synchronization Group Properties](#) on page 1188.
- Configure DNS replication, as described in [Replicating DNS Data](#) on page 1189.
- Remove a load balancer group, as described in [Removing a Load Balancer Synchronization Group](#) on page 1190.

Setting Load Balancer Properties

You can modify certain load balancer properties you previously configured, such as the logging level, extensible attributes, and administrative permissions. Extensible attributes and permissions apply to the data only when they are managed from Grid Manager. Note that extensible attributes and permissions are not synchronized to the GLB.

To set the GLB properties:

1. From the **Grid** tab, select the **Load Balancers** tab -> *load_balancer* check box, and then click the Edit icon.
2. In the *Load Balancer Configuration* editor, you can set properties in the following tabs. For information about the fields in the **General** and **Certificate** tabs, see [Assigning Grid Members to Load Balancers](#) on page 1184.
 - **General -> Basic:** Modify the general settings such as the device name and synchronization interval. When you change the IP address of a GLB, you must click the Validate LB Device Connection icon to test the connection before the synchronization resumes. For information about how to test connection, see [Validating Load Balancer Connection](#) on page 1186.
 - **General -> Advanced:** Click **Override** and modify the load balancing interface.
 - **Certificate:** This tab displays the SSL certificate information of the GLB. You can modify the SSL validation option.
 - **Extensible Attributes:** Define extensible attributes for the GLBs. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** Define administrative permissions that apply to the server. For information see [About Administrative Permissions](#) on page 160.
3. Save the configuration.

Modifying Load Balancer Synchronization Group Properties

You can modify certain properties of a load balancer synchronization group by selecting a group or selecting a GLB, which is part of a synchronization group. To edit group properties, do the following:

1. From the **Grid** tab, select the **Load Balancers** tab -> *Load Balancer Group* check box, and then click the Edit icon.
2. In the *Load Balancers Group* editor, select the **General** -> **Basic** tab and do any of the following:
 - **LB Group Name:** Displays the name of a load balancer synchronization group. You can modify the group name on NIOS. This name is not updated on the GLBs.
 - **Synchronization Interval (min):** The default synchronization interval is two minutes. You can modify the synchronization interval.
 - **Managing Member:** Click **Select Member** and select another Grid member.
 - **Manage Load Balancer(s):** Select either **Read-only** or **Read/Write**. Depending on the management mode, NIOS admins can perform different tasks. For information, see [Setting the Management Mode](#) on page 1183.
 - **Sync Zones from Load Balancer(s):** Select this check box to synchronize all discovered DNS zones from the GLB to NIOS or clear this check box to stop the DNS zones synchronization.
 - **Comment:** Enter additional information.
 - **Disabled:** Select this check box to disable the synchronization group. Clear the check box to enable the load balancer synchronization group. The status of the load balancer synchronization group is updated only after it synchronizes with Grid.
3. Save the configuration.

Setting Priorities for Managed Load Balancers

You can configure priorities for managed GLBs that belong to a load balancer synchronization group. Priorities determine the order in which NIOS attempts to contact the managed GLBs. In a group, the GLB with the lowest priority value is selected as the lead device. For example, a load balancer with the priority set to 0 has the highest priority and is considered as the lead device. NIOS synchronizes group data with the lead device and other grouped GLBs synchronize their own GLB data. The lead device must be online, not disabled, and reachable from NIOS. In addition, the lead device must have a replication path to all other GLBs in the group so it can replicate changes that NIOS makes on the lead device. If the lead device is unavailable, NIOS determines the next lead device based on the priority value and synchronizes data with the new lead device. Note that automatic regrouping occurs when NIOS selects a new lead device.

To set the priority:

1. From the **Grid** tab, select the **Load Balancers** tab -> *load_balancer* check box, and then click the Edit icon.
2. In the *Load Balancer Configuration* editor, select the **Managed Load Balancer(s)** tab.
3. Select a *load_balancer* check box and enter an integer from 0 to 255 in the **Priority** column.

Note: Infoblox recommends that you assign the highest priority (0) to the designated lead device if it exists for a synchronization group. You can optionally set other GLBs in the group to read-only mode to ensure that NIOS updates only the lead device. Doing this makes replication more predictable and reduces redundancy.

4. Save the configuration.

Changing the Managing Member or Management Mode

You can change the managing member and the management mode of a GLB. If you change the managing member, the previous member aborts any ongoing synchronization, and the newly assigned member resumes the synchronization process.

To change the member or management mode:

1. From the **Grid** tab, select the **Load Balancers** tab -> *load_balancer* check box, and then click the Edit icon.
2. In the *Load Balancer Configuration* editor, select the **General** -> **Basic** tab and do any of the following:
 - **Managing Member:** Click **Select Member** and select another Grid member.
 - **Manage Load Balancer(s):** Select either **Read-only** or **Read/Write**.
 - **Disabled:** Select this check box to disable the load balancer. The Grid member does not try to connect to disabled GLB. Clear this check box to enable the load balancer. Status of the load balancer is updated only after it synchronizes with Grid.
3. Save the configuration.

Replicating DNS Data

You can enable a load balancer synchronization group to replicate DNS data. When DNS data replication is enabled on the group, NIOS synchronizes DNS data with the lead device. If DNS replication is disabled for the group, NIOS synchronizes DNS data with each individual GLB in the group. You must enable this feature on the load balancer synchronization group. Though you cannot enable this on NIOS, you can view whether the feature is enabled on the group or not.

To view DNS replication:

1. From the **Grid** tab, select the **Load Balancers** tab -> *load_balancer* check box, and then click the Edit icon.
2. In the *Load Balancers Group* editor, select the **Replication** tab.
3. The appliance displays the DNS replication information:
 - **Replicates DNS objects:** Indicates whether DNS replication is enabled for the synchronization group.

Backing Up Synchronized Data

When you back up the Grid, it includes all synchronized load balancer data. If you restore a backup, the data is restored on the Grid only. It is not synchronized to the GLBs. When the Grid member synchronizes the data after the restore operation, it overrides the data on the Grid with the data from the GLBs. For information about backing up and restoring data, see [Chapter 9, Managing NIOS Software and Configuration Files](#), on page 405.

Disabling Synchronization

When you disable synchronization, the Grid member completes any on-going synchronization and does not start a new one. Setting this option only affects data synchronization and does not affect the operations of the load balancer. Synchronization resumes when the load balancer is re-enabled.

To disable a load balancer:

1. From the **Grid** tab, select the **Load Balancers** tab -> *load_balancer* check box, and click the Edit icon.
2. In the **General** tab -> **Basic** tab, select the **Disable Synchronization** option.
3. Save the configuration.

Removing a Load Balancer

When you remove a load balancer from the Grid, the managing member stops any on-going synchronization and does not start a new one. If the load balancer served DNS, the synchronized DNS data remains unchanged in the Grid.

Removing a managed load balancer from the Grid does not affect the operations of the load balancer.

To remove a managed GLB:

1. From the **Grid** tab, select the **Load Balancers** tab -> *load_balancer* check box, and click the Delete icon.
2. When the *Delete Confirmation* dialog box appears, click **Yes**.

Removing a Load Balancer Synchronization Group

When you remove a load balancer group from the Grid, all GLBs associated with that group are removed from NIOS. Data on the managed GLBs remains unchanged, but NIOS specific data such as extensible attributes, comments, permissions on the load balancer synchronization group, and associated GLBs are removed. When you remove the last GLB in a group, the entire group will be removed.

To remove a load balancer group:

1. From the **Grid** tab, select the **Load Balancers** tab -> *Load Balancer Group* check box, and click the Delete icon.
2. When the *Delete Confirmation* dialog box appears, click **Yes**.

Note: Scheduled tasks for GLBs of a synchronization group are not removed from the **Task Manager** tab. When you try to add the deleted GLB again, Grid Manager displays an error message indicating that there is a pending task for the selected GLB. You cannot add the deleted GLB until the task is executed or deleted.

MONITORING GLOBAL LOAD BALANCERS

You can monitor the status of managed GLBs and synchronized groups from the Dashboard and from various panels in the **Grid** tab. The appliance also logs events related to the synchronization process in the Infoblox syslog, depending on the logging level that you configured in the **General** -> **Basic** tab of the *Load Balancer Properties* editor described in [Setting Load Balancer Properties](#) on page 1187.

You can monitor load balancers and their objects as follows:

- View the *Load Balancers Status* widget on the Dashboard. For information, see [Load Balancer Status](#) on page 133.
- View the status of load balancers and load balancer synchronization groups. For information, see [Viewing Global Load Balancers](#).
- View the load balancer synchronization events in the syslog. For information, see [Viewing the Syslog](#) on page 1016.

Viewing Global Load Balancers

You can view details about the managed GLBs and load balancer groups by navigating to the **Grid** tab -> **Load Balancers** tab.

For each load balancer and load balancer synchronization group, the panel displays the following:

- **Name:** The name of the load balancer or the name of the load balancer group. To view the members of a specific load balancer group, click the arrow next to the group name to expand the group. All groups are collapsed by default.
- **Status:** The connection status, which can be one of the following:
 - **OK:** The Grid member is connected to the GLB. For a load balancer group, the Grid member is connected to the lead device.
 - **Unknown:** The Grid member is unable to contact the GLB or the lead device and cannot retrieve any status details. This can be caused by incorrect IP address, FQDN, username, or password.
 - **Error:** The GLB has a connection error. Click the Detailed Status icon to view detailed information or check the syslog for any error messages.
 - **Warning:** Certain issues, such as Grid member failures or licensing issues, have occurred. Click the Detailed Status icon to view detailed information or check the syslog for messages to determine the reason for the warning.

Note: Synchronization may fail and the connection status continues to display “Warning” if the F5 server name contains invalid characters. Infoblox recommends that the name of the server contains only ASCII characters.

- **Disabled:** The load balancer is disabled. The Grid member does not try to connect to disabled GLB.
- **Comment:** Information about the load balancer or the load balancer group.
- **IPv4 Address:** The IPv4 address of the load balancer.
- **IPv6 Address:** The IPv6 address of the load balancer.
- **Address or FQDN:** The FQDN or the address of the load balancer.
- **Version:** The supported version of the load balancer.
- **Managing Member:** The FQDN of the Grid member that manages the load balancer.
- **Site:** The location to which the load balancer belongs. This is one of the predefined extensible attributes.

You can also do the following:

- Click the Add icon to add a load balancer. For information, see [Assigning Grid Members to Load Balancers](#) on page 1184.
- Select a load balancer and click the Edit icon to edit the load balancer configuration. For information, see [Setting Load Balancer Properties](#) on page 1187.
- Select a load balancer and click the Delete icon to delete it. For information, see [Removing a Load Balancer](#) on page 1190.
- Select a load balancer group and click the Delete icon to delete it. For information, see [Removing a Load Balancer Synchronization Group](#) on page 1190.
- Select a load balancer and click the Detailed Status icon to view detailed information about the load balancer. For information, see [Viewing Detailed Status Information](#) on page 1192.
- Select a load balancer group and click the Detailed Status icon to view detailed information about the lead device. For information, see [Viewing Detailed Status Information of a Synchronization Group](#) on page 1192.
- Click the Validate LB Device Connection icon to test the connection and to view certificate information of the load balancer. For information, see [Validating Load Balancer Connection](#) on page 1186.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

- Click the Export icon to export the list of load balancers to a .csv file.
- Click the Print icon to print the list of load balancers.

Viewing Detailed Status Information

You can view more status information by selecting a load balancer from the **Load Balancers** tab and clicking the Detailed Status icon. The Detailed Status panel displays the following Information:

- **Device Status:** The status icon indicates the load balancer as follows:
 - Green: The load balancer is synchronizing data with the Grid member.
 - Red: There is a connection error caused by incorrect IP address, FQDN, username, or password. You can check the messages in the load balancer log to determine the reason.
 - Black: The load balancer is disabled. View **Status Detail** for more information.
 - Yellow: There are some issues with the synchronization between the member and server. You can check the messages in the load balancer log to determine the reason for the warning.
- **Status Detail:** Information about the current device status.
- **Product Name:** The product name of the load balancer.
- **Last Updated:** The date and time of the last synchronization.
- **Version:** The version number of the load balancer and the build number.
- **Uptime:** The total time that the load balancer has been up and running.

Viewing Detailed Status Information of a Synchronization Group

You can view detailed status information by selecting a load balancer synchronization group from the **Load Balancers** tab and clicking the Detailed Status icon. The Detailed Status panel displays the following Information:

- The name of the synchronization group.
- **Lead Device:** IP address of the lead device.
- **Version:** The version number of the lead device and the build number.
- **Group Status:** The status icon can be one of the following:
 - Green: The load balancer synchronization group is synchronizing data with the Grid member.
 - Red: There is a connection error and this can be caused by incorrect IP address, FQDN, username, or password.
 - Yellow: There are some issues with the synchronization between NIOS and GLBs or GLBs are regrouping. Note that the status changes to OK (green) after the GLBs of a synchronization group regroup successfully. You can check the messages in the load balancer log to determine the reason for the warning.
 - Black: A load balancer is disabled or synchronization group is disabled.
- **Status Detail:** Information about the current group status.
- **Last Updated:** The timestamp of the last synchronization.



Chapter 40 Managing Global Load Balancer Data

This chapter provides guidelines for using Grid Manager to manage GLBs (Global Load Balancers) related data. It also describes how to monitor GLB data. It includes the following sections:

- [*About Global Load Balancer Data*](#) on page 1195
 - [*Supported Global Load Balancer Objects*](#) on page 1195
 - [*Viewing Global Load Balancer Objects*](#) on page 1196
 - [*Viewing Traffic Management Structures*](#) on page 1197
- [*Managing Synchronized DNS Data*](#) on page 1198
 - [*Mapping DNS Views*](#) on page 1198
 - [*Resolving DESYNC Conflicts*](#) on page 1199
 - [*Configuring Delegations to Zones on Load Balancers*](#) on page 1199
 - [*Modifying Delegations to Zones on Load Balancers*](#) on page 1200
- [*Managing LBDN on GLBs*](#) on page 1201
 - [*Associating GLBs with Authoritative Zones*](#) on page 1201
 - [*Adding LBDNs*](#) on page 1201
 - [*Modifying LBDNs*](#) on page 1202
 - [*Viewing LBDNs and LBDN Records*](#) on page 1203
- [*Managing Listeners*](#) on page 1203
 - [*Adding Listeners*](#) on page 1203
 - [*Modifying Listeners*](#) on page 1204
- [*Managing Data Centers*](#) on page 1204
 - [*Adding Data Centers*](#) on page 1204
 - [*Modifying Data Centers*](#) on page 1205
- [*Managing Global Load Balancing Pools*](#) on page 1205
 - [*Creating Pools and Adding Pool Members*](#) on page 1205
 - [*Modifying Load Balancing Pools*](#) on page 1206
- [*Managing Global Load Balancer Servers*](#) on page 1207
 - [*Adding GLB Servers*](#) on page 1207
 - [*Modifying GLB Servers*](#) on page 1208

- [*Managing Global Load Balancer Virtual Servers*](#) on page 1209
 - [*Adding GLB Virtual Servers*](#) on page 1209
 - [*Modifying GLB Virtual Servers*](#) on page 1210

ABOUT GLOBAL LOAD BALANCER DATA

After you configure a Grid member to manage a GLB or a load balancer synchronization group, the Grid member connects to the GLB and starts synchronizing data to the NIOS database. The synchronization time varies, depending on factors such as the number of managed load balancers and the amount of data being synchronized. The synchronized data is then replicated to the Grid Master through the Grid replication process.

The Grid member can synchronize DNS zones on the GLBs, and it supports certain types of GLB objects and data. For more information about synchronizing DNS zones, see [Managing Synchronized DNS Data](#) on page 1198. For information about the supported GLB objects, see [Supported Global Load Balancer Objects](#). To configure the Grid member to synchronize DNS zones, you must first map the GLB DNS view to the corresponding DNS view on NIOS. For information about how to map DNS views for GLBs, see [Viewing Global Load Balancer Objects](#) on page 1196. Note that the member cannot synchronize discovered DNS zones if you have not mapped the DNS views even if you have enabled the member to synchronize DNS zones from GLBs.

Grid Manager displays synchronized GLB objects in the **Traffic Management** tab. For information about viewing GLB objects, see [Viewing Global Load Balancer Objects](#) on page 1196. You can also view the hierarchy of a GLB and its associated objects that Infoblox supports. For information about how to view the hierarchical map of GLB objects, see [Viewing Traffic Management Structures](#) on page 1197.

Supported Global Load Balancer Objects

If the load balancer and GLBs in a synchronization group is managed in read/write mode, NIOS admins can update the synchronized GLB data. For information about admin permissions, see [Administrative Permissions for Load Balancers](#) on page 216.

Grid members synchronize the following supported GLB objects:

- **Load Balancer Delegation:** You can add delegations to zones on load balancers. For information about how to configure them, see [Configuring Delegations to Zones on Load Balancers](#) on page 1199.
- **LBDN:** Load Balanced Domain Name. This is most likely the domain name that corresponds to the pool of a local load balancer. You can add an LBDN record to an authoritative zone or to delegated load balancer zone. For more information about how to add and modify LBDN objects, see [Managing LBDN on GLBs](#) on page 1201.
- **Listener:** A listener is an IP address that allows a GLB to listen on the network for traffic it is responsible for. For information about how to add and modify listeners, see [Managing Listeners](#) on page 1203.
- **Data Center:** A data center is a centralized repository for the storage, management, and dissemination of data. You must configure a data center for each physical or virtual location in your network. For information about how to add and modify data center objects associated with GLBs, see [Managing Data Centers](#) on page 1204.
- **Pool:** A load balancing pool is a logical set of devices that you group together to receive and process traffic. A GLB can have associated load balancing pools that contain various members for performing load balancing tasks. For information about how to add and modify load balancing pool members, see [Managing Global Load Balancing Pools](#) on page 1205.
- **Server:** A server is a physical device on which you can configure one or more virtual servers. For information about how to add and modify servers, see [Managing Global Load Balancer Servers](#) on page 1207.
- **Virtual Server:** A virtual server is a traffic management object that is represented by an IP address and a service. A virtual server is often used to balance traffic among a pool of servers on the network. When you assign a load balancing pool to a virtual server, the GLB directs traffic coming into the virtual server to one of the pool members. For information about how to add and modify virtual servers, see [Managing Global Load Balancer Virtual Servers](#) on page 1209.

Note that you cannot use Grid Manager to create unsupported objects and assign them to a load balancer. For more information about these supported objects, refer to the corresponding GLB documentation.

Viewing Global Load Balancer Objects

After the Grid member synchronizes data with the GLBs, Grid Manager lists all GLB objects in the **Data Management** tab -> **Traffic Management** tab.

Note that when you modify or delete a GLB object, such as a GLB virtual server, it may take a little while for the data synchronization to complete before you can view the updated information in the **Traffic Management** tab. For example, Grid Manager may display the “Selected object could not be found” message when you try to modify a deleted virtual server. Grid Manager removes the virtual server from the viewer when data synchronization is complete.

To view GLB objects:

1. From the **Data Management** tab, select the **Traffic Management** tab.
2. In the Traffic Management panel, select the GLB object you want Grid Manager to display. Infoblox supports the following object types: **LBDN**, **Listener**, **Data Center**, **GLB Pool**, **GLB Server**, and **GLB Virtual Server**. For more information, see [Supported Global Load Balancer Objects](#) on page 1195.
3. Based on the selected objects, Grid Manager displays the following for each GLB object:
 - **Name:** The name of the GLB object.
 - **Type:** The object type. This can be one of the supported objects: LBDN, Listener, Data Center, Pool, Server or Virtual Server.
 - **Status:** Displays the connection status, which can be one of the following:
 - **Checking:** The appliance is checking the status of the GLB object.
 - **Unlicensed:** The GLB object is not licensed.
 - **Running:** The GLB object is running properly.
 - **None:** The appliance cannot retrieve the status of the GLB object.
 - **Error:** The GLB object has an error. You can check the syslog for any messages.
 - **Warning:** Certain issues, such as object failures, have occurred. You can check the syslog for any messages.
 - **IPv4 Address:** The IPv4 address of the GLB object, if applicable.
 - **Device/Group:** Displays the IP address of the load balancer. For GLB groups, this field displays the name of the group.
 - **Lead Device Name:** Displays the name of the lead device of a group.
 - **Comment:** Displays any comments that were entered for the GLB object.
 - **Site:** Displays any values that were entered for this pre-defined attribute.

You can add the following columns for display:

- **IPv6 Address:** The IPv6 address of the GLB object, if applicable.
- **Partition:** The name of the GLB partition, if applicable.

You can do the following in the **Traffic Management** tab:

- Click **Toggle multi-line view** or **Toggle single-line view** to view objects.
- Click the Add icon to add a GLB object.
- Select a GLB object and click the Edit icon to edit the configuration.
- Select a GLB object and click the Delete icon to delete it.
- Click the Configure icon next to the GLB check box to display the configuration menu for different tasks.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- Click the Export icon to export the list of GLB objects to a .csv file.
- Click the Print icon to print the list of GLB objects.

Viewing Traffic Management Structures

Grid Manager provides a hierarchical map view that you can use to quickly capture an overall traffic management structure of a selected GLB object. The map displays objects in hierarchical order of LBDN -> pool -> pool members. It starts with a selected GLB object as the central point and displays its relationship with other associated objects. The objects are represented as nodes in the map. You can hover your mouse over the node to display the tooltip that contains the following information:

- **Name:** The device name of the GLB object.
- **Status:** The current status of the GLB object.
- **Type:** The object type.
- **Last Discovered Time:** The timestamp when the object was last discovered.
- **Health Status:** The current status of the GLB object.

When you select an object, the map displays the selected object at the center and uses lines to represent its connections with other associated objects. For example, if you start with an LBDN object that has associated pool members and GLB servers, the map displays connectors from the LBDN object to all the related objects in a hierarchical tree format.

Note: If your browser has a pop-up blocker enabled, you must turn off the pop-up blocker or configure your browser to allow pop-ups in order to view the traffic management visualizer.

To view the hierarchical map of a GLB object:

1. From the **Data Management** tab -> **Traffic Management** tab, select a GLB object that has associated objects and data, and then click the Configure icon next to the check box.
2. At the top of the configuration menu, select **LBDN -> *glb_object_name***.

Grid Manager opens a separate browser window and displays the hierarchical map in the **Traffic Management Structure** panel. The map displays the selected GLB object as the starting point and its connections and relationships with other associated objects in the network. The map also contains legends that indicates the status of the displayed objects and the nature of their paths to other objects.

You can also do the following in the **Traffic Management Structure** panel:

- Configure the number of nodes on each level by clicking the Tree Configuration icon at the upper top right corner of the map. Enter the maximum number of nodes you want displayed for each level and then click **Submit**.
- Change the map orientation by clicking the Change Tree Orientation icon. The default orientation is from top to bottom.
- Click the Refresh icon to refresh the map. You can also click **Turn Auto-Refresh On** to turn on auto-refresh and **Turn Auto-Refresh Off** to turn it off.
- Click anywhere in the tree map and hold your mouse to drag the map to a desired location in the panel.

MANAGING SYNCHRONIZED DNS DATA

When Grid members are configured to manage GLBs in read/write mode, you can use Grid Manager to view, edit and delete the DNS zones discovered on those GLBs. You can add, modify, and delete the following supported DNS objects through Grid Manager:

- Delegations to zones on GLBs. For information, see [Configuring Delegations to Zones on Load Balancers](#) on page 1199.
- LBDNs on authoritative zones or delegated zones. For information, see [Managing LBDN on GLBs](#) on page 1201.

Before you can manage DNS data on GLBs through Grid Manager, you must complete the following:

- Manage the GLB in read/write mode, as described in [Setting the Management Mode](#) on page 1183.
- Enable the GLB. Refer to the appropriate GLB documentation for details.
- Map the GLB DNS view to the corresponding DNS view on NIOS, as described in [Mapping DNS Views](#) on page 1198. Otherwise, synchronization of DNS zones does not start.
- Enable the synchronization of all discovered DNS zones from the GLB, as described in [Assigning Grid Members to Load Balancers](#) on page 1184.

Note that when you add delegated zones on GLBs through Grid Manager, the zones are synchronized to the GLBs even when you disable synchronization from the GLBs. All updates are synchronized to the GLB at regular intervals.

Note: A delegated zone configured for the GLB can be synchronized to the GLB even when the FQDN of the A record is not within the FQDN of the zone. However, the A record may not be synchronized to the GLB.

Mapping DNS Views

In order for the Grid member to synchronize DNS zones it discovers on the GLB, you must map the GLB DNS view to the corresponding DNS view on the Grid member. Otherwise, even if you have configured NIOS to synchronize DNS zones discovered on the GLB, the synchronization does not happen. In addition, before you configure NIOS to synchronize DNS zones from a GLB, ensure that all GLB DNS views in which you create DNS zones are viewable from the GLB user interface. Otherwise, NIOS logs an “NOTAUTH” error in the syslog. You can check the GLB user interface to see if the same error occurs on the GLB. If so, correct the DNS view configuration. For information about how to properly configure GLB DNS views, refer to the GLB documentation.

Note that GLB DNS views, even those on the same GLB, are independently mapped to the NIOS view. NIOS synchronizes each DNS view mapping independently.

You can map one GLB DNS view or multiple zones from different GLB views to a single NIOS view. When you map multiple GLB DNS views to one NIOS view, consider the following:

- The SOA serial numbers on the GLB views are updated independently and have no relations to the SOA serial number of the NIOS GLB delegation.
- When you change zone data on the single NIOS view, the data is synchronized to all GLB DNS views that are mapped to the NIOS view.
- When you change zone data on one of the GLB DNS views, the data is first synchronized to the NIOS view, and then to other GLB DNS views that are mapped.
- When you make different zone data changes on multiple GLB DNS views at the same time, the changes create a conflict and cause the DNS zone synchronization to stop. The NIOS DNS zone goes into the DESYNC state. When a zone is in the DESYNC state, no data synchronization happens on the zone. You can view the syslog for DESYNC issues. To resolve a DESYNC issue, see [Resolving DESYNC Conflicts](#) on page 1199.

Note: The DESYNC state is for a particular GLB DNS view mapping. Other valid mappings for the same zone may be unaffected. The zone may still be updating while in a DESYNC state because another valid mapping exists and is synchronizing

To map DNS views:

1. From the **Grid** tab, select the **Load Balancers** tab, and click **Map DNS View** from the Toolbar.
2. In the *Load Balancer View to DNS View Mapping* editor, do the following:
 - Click a **Device/Group Name** row, and then double click the corresponding **Grid View**. The *Load Balancer View to DNS View Mapping* editor displays an individual load balancer, a GLB group with DNS synchronization enabled, and a GLB, which is part of a group with DNS synchronization enabled.
 - Select the corresponding DNS view from the **Grid View** drop-down list.
3. Save the configuration.

Note: DNS synchronization starts only after the mapping is complete.

Resolving DESYNC Conflicts

When zone data on different GLBs that are mapped to a NIOS view are updated concurrently, the changes may create synchronization conflicts and cause the NIOS DNS zone to go into a DESYNC state. You can view the DESYNC issue in the syslog.

To resolve the conflict and get the DNS zone out of the DESYNC state:

- Change zone data for all GLB DNS views to match that of the corresponding NIOS zone. The next DNS synchronization will identify the updates and put the DNS zone back in the SYNC state.
- Change zone data for the DESYNC zone on NIOS. This will manually push the data to all the mapped GLB DNS views and force the GLBs to accept the updated data and put the zone back in to the SYNC state.

Configuring Delegations to Zones on Load Balancers

When you configure a delegation to a GLB, the delegation appears as an authoritative zone on the GLB.

To add a delegation to the GLB:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* -> **Subzones** tab, select **Add -> Load Balancer Delegation** from the toolbar.
2. In the *Add Load balancer Delegation* wizard, complete the following:
 - **Name:** Click **Select Zone** to select a zone. The field displays a dot followed by the domain name of the selected or current zone. Enter one or more labels before the dot to specify the domain name of the delegation.
 - **DNS View:** This field appears only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
 - **Comment:** Enter information about the delegation.
 - **Disable:** Select this option to temporarily disable this delegated zone. Disabling the load balancer delegated zone disables the delegation NS and glue A records in the parent authoritative zone, but the delegation is still synchronized to the GLB.
 - **Lock:** Select this option to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
3. Click **Next** and complete the following to add name servers:

Click the Add icon. The appliance adds a row to the Name Servers table. Click the row and specify the following:

 - **Name:** Enter the name of a remote name server to which you want the local server to redirect queries for data for the zone.
 - **Address:** Enter the IP address of the delegated server.
4. Click **Next** and complete the following to add the managed load balancers or load balancer groups:

Click the Add icon, and select a load balance or a group from the *Load Balancer Device/Group Selector*. The appliance displays the following in the All Devices/Groups table:

- **Name:** The name of the load balancer or the name of a GLB group.

You can select the check box and click the Delete icon to delete the load balancer or a group.

5. Click **Next** and complete the following to configure SOA record settings:

- **Primary name server (for SOA MNAME field):** Enter the name of the primary name server that serves this zone. This appears in the MNAME field of the SOA record.
- **Email address (for SOA RNAME field):** Enter a valid email address of the admin who manages the domain zone files. This appears in the RNAME field of the SOA record.
- **Refresh:** Define the time and the time unit the secondary DNS server waits before querying the primary DNS server's SOA record to check for changes. When the refresh time expires, the secondary DNS server requests a copy of the current SOA record from the primary. The secondary DNS server compares the serial number of the primary DNS server's current SOA record and the serial number in its own SOA record. If they are different, the secondary DNS server will request a zone transfer from the primary DNS server. The default value is three seconds.
- **Retry:** Define the time the secondary server waits before retrying a failed zone transfer. Normally, the retry time is less than the refresh time. The default value is one second.
- **Expire:** Define the time and the time unit that the secondary server keeps trying to complete a zone transfer. If this time expires prior to a successful zone transfer, the secondary server expires the zone file. The secondary server stops answering queries, as it considers its data too old to be reliable. The default value is four seconds.
- **Default TTL:** Define the minimum time-to-live value that applies to all resource records in the zone file. This value is supplied in query responses to inform other servers how long they should keep the data in cache. The default value is eight seconds.
- **Negative caching TTL:** Define the negative caching time-to-live value. The default value is 15 seconds.

6. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

7. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the Schedule Change panel, click **Later**, and then specify a date, time, and time zone.

Modifying Delegations to Zones on Load Balancers

To modify the settings of a load balancer delegation:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* -> **Subzones** tab, select a delegation and click the Edit icon.
 2. The *LB Delegation* editor contains the following basic tabs from which you can modify data. For information about how to modify data, see [Configuring Delegations to Zones on Load Balancers](#) on page 1199.
 - **General:** This tab displays the **Name** and **Type** of the load balancer delegation. You can modify the **Comment**, **Disable**, and **Lock** fields.
 - **SOA Configuration:** You can override the current settings and enter unique ones for the delegation.
 - **Delegated Name Servers:** Displays the delegated name servers. You can add or delete name servers in this tab.
 - **Load Balancer Devices/Groups:** Displays the managed load balancers and load balancer groups. You can add or delete load balancers or load balancer groups in this tab.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the GLB delegation. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
 3. Save the configuration.
- or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

MANAGING LBDN ON GLBs

When you add an LBDN (Load Balanced Domain Name) to a GLB, one or more LBDN records may be created as proxies for the LBDN if there are zones associated with the GLB that match the name or alias of the GLB. You can add LBDNs to GLBs, and modify their LBDN records. When you modify LBDN records, the changes will affect the corresponding LBDNs. For example, when you delete an LBDN record, the associated LBDNs will be deleted. Once you create an LBDN, you cannot change or assign it to another GLB. Note that DNS views and DNS view mapping have no effects on how you manage LBDNs and their corresponding records.

Grid Manager displays LBDN and LBND records based on your configuration. All LBDN appear in the **Traffic Management** tab, but only the LBDN records that match a zone name appear in the **DNS** tab -> *authoritative_zone* or *glb_delegated_zone* -> **Records** tab. For more information about how to view LBDN records, see [Viewing LBDNs and LBDN Records](#) on page 1203.

Associating GLBs with Authoritative Zones

In order to manage LBDN in an authoritative or delegated zone through Grid Manager, you must enable the authoritative zone and associate it with the GLB to which the LBDN belongs. You can create LBDN on different GLBs using the same name. When you associate a zone with multiple GLBs, all LBDN proxies for the LBDNs (including those with the same names) on the associated GLBs appear in the **Records** tab of the zone. For more information, see [Viewing LBDNs and LBDN Records](#) on page 1203.

To associate a zone with its GLB:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *authoritative_zone* check box, and then click the Edit icon.
2. In the *Zone* editor, select the **Load Balancer Devices/Groups** tab, and complete the following:
Click the Add icon, and the appliance displays the *Load Balancer Device/Group Selector*. In the dialog box, select the load balancer or group and the appliance displays the following:
 - **Name:** The name of the load balancer or the name of the GLB group.
 You can select the check box and click the Delete icon to delete the load balancer or a group.
3. Save the configuration.

Adding LBDNs

To add an LBDN to a GLB:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *authoritative_zone* or *glb_delegated_zone* -> **Subzones** tab, select **Add** -> **Record** -> **LBDN** from the toolbar.
or
From the **Data Management** tab, select the **Traffic Management** tab, and then click **Add** -> **LBDN** from the Toolbar.
2. In the *LBDN* wizard, complete the following:
 - **Load Balancer Device/Group:** Click **Select** and select a GLB or a load balancer group from the *Load Balancer Device/Group Selector*.
 - **Partition:** Select a partition from the drop-down list.
3. Click **Next** and complete the following information:
 - **Name:** Enter the domain name of the load balancer.
 - **Alias:** Click the Add icon. The appliance adds a row to the Alias table. Click the row and enter the name of the alias.

- **Load balancing method:** From the drop-down list, select the method you want to use for load balancing. The default is **Round Robin**, which causes the GLB to send each incoming request to the next available load balancing member.
 - **Persistence:** Select this check box to enable persistence on load balancing members.
 - **Comment:** Enter additional information about the LBDN object.
 - **Disabled:** Select this to disable the LBDN record.
4. Click **Next** and complete the following to add iRules:

Click the Add icon. In the *Load Balancer iRules Selector*, select the iRules you want to add to the LBDN. An iRule is a user-written script designed to inspect and direct individual connections in specific ways. You define iRules on the GLB. The appliance adds the selected iRules to the table and displays the following:

 - **Name:** The name of the iRule.
 5. Click **Next** and complete the following to add load balancing pool members:

Click the Add icon. In the *Load Balancer Pools Selector*, select the load balancing pools you want to associate with this LBDN record. The appliance adds the selected pools to the table and displays the following:

 - **Name:** The name of the pool.
 - **Ratio:** The ratio weight you want to assign to the pool member. The ratio weight determines the amount of traffic the pool member receives. You can click this field and modify the ratio. The GLB distributes connections among members according to the ratio weight. Enter a number between 0 to 65535. The default is 1.

You can add another pool or select a pool check box and click the Delete icon to delete the selected pool.

 - **Last Resort Pool:** From the drop-down list, select the last resort pool.
 6. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 7. Save the configuration.
- To schedule this task, click the Schedule icon at the top of the wizard. In the Schedule Change panel, click **Later**, and then specify a date, time, and time zone.

Modifying LBDNs

To modify an LBDN:

1. From the **Data Management** tab -> **Traffic Management** tab, select an LBDN and click the Edit icon.
 2. The *LBDN* editor contains the following basic tabs from which you can modify data. For information about how to modify data, see [Adding LBDNs](#) on page 1201.
 - **General:** This tab displays the **Load Balancer**, **Partition**, **Name** and **Description** of the LBDN object. You can modify the rest of the fields.
 - **iRules:** Displays the selected iRules. You can add or delete iRules in this tab.
 - **Pools:** Displays the associated load balancing pools. You can add or delete pools in this tab.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the LBDN record. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
 3. Save the configuration.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

Viewing LBDNs and LBDN Records

Depending on your configuration, Grid Manager displays LBDNs in the following tabs:

- **Data Management** tab -> **Traffic Management** tab
Grid Manager displays all LBDNs in this tab. For more information, see [Viewing Global Load Balancer Objects](#) on page 1196.
- **Data Management** tab -> **DNS** tab -> **Zones** tab -> *authoritative_zone* or *delegated_zone* -> **Records** tab
Only LBDN records that match the domain name of a specific zone appear in this tab. For example, in an authoritative zone named “com”, only LBDN records with names that contain “com”, such as “V11.com” or “corp100.com”, appear in the **Records** tab. For LBDN record proxies for LBDNs to show up in this tab, the zone must be associated with the GLB. To associate a zone with a GLB, see [Associating GLBs with Authoritative Zones](#) on page 1201. For more information, see [Viewing Resource Records](#) on page 678.

MANAGING LISTENERS

To communicate with the rest of your network, you must configure the GLB so it can correctly identify the resolution requests for which it is responsible. A listener is an object that monitors the network for DNS queries. The listener instructs the system to monitor the network traffic destined for a specific IP address. Listeners are device specific and they are not replicated by synchronization groups.

Adding Listeners

To add a listener:

1. From the **Data Management** tab, select the **Traffic Management** tab, and then click **Add** -> **Listener** from the Toolbar.
2. In the *Listener* wizard, complete the following:
 - **Load Balancer:** Click **Select** and select a GLB from the *Load Balancer Device Selector*.
 - **Partition:** Select a partition from the drop-down list.
3. Click **Next** and complete the following:
 - **Destination (Name):** Enter a valid IPv4 or IPv6 address for destination address.
 - **VLAN Traffic:** Select one of the following from the drop-down list:
 - **All VLANs:** Select this to include all VLAN for the traffic.
 - **Enabled on...:** Use the arrows to move VLANs from the Available table to the Active table and vice versa. The appliance enables traffic for the listener on all VLANs and tunnels in the Active table.
 - **Disabled on...:** Use the arrows to move VLANs from the Available table to the Active table and vice versa. The appliance disables traffic for the listener on all VLANs and tunnels in the Active table.
 - **Protocol:** Select the protocol you want to use from the drop-down list.
 - **DNS Profile:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the DNS profile for the listener.
 - **Route Advertisement:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the route advertisement for the data center.
 - **Comment:** Enter additional information about the load balancer object.
4. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
5. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

Modifying Listeners

To modify a listener:

1. From the **Data Management** tab -> **Traffic Management** tab, select a listener and click the Edit icon.
2. The *Listener* editor contains the following basic tabs from which you can modify data. For information about how to modify data, see [Adding Listeners](#) on page 1203.
 - **General:** This tab displays the **Load Balancer, Partition, Destinations (Name)** and **Description** of the listener. You can modify the rest of the fields.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the listener. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Save the configuration.
or
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

MANAGING DATA CENTERS

A data center is a centralized repository that facilitates the storage, management, and dissemination of data for a specific organization or business. It can contain any type of servers. GLBs distribute workloads and route users to the most efficient data centers to perform specific tasks. You must configure a data center for each physical location in your network.

Adding Data Centers

To add a data center:

1. From the **Data Management** tab, select the **Traffic Management** tab, and then click **Add -> Data Center** from the Toolbar.
2. In the *Data Center* wizard, complete the following:
 - **Load Balancer Device/Group:** Click **Select** and select a GLB or a group from the *Load Balancer Device/Group Selector*.
 - **Partition:** Select a partition from the drop-down list.
3. Click **Next** and complete the following:
 - **Name:** Enter a name for the data center.
 - **Location:** Enter the location of the data center.
 - **Contact:** Enter the name of the admin who manages the data center.
 - **Prober Pool:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the prober pool for the data center.
 - **DNS Profile:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the DNS profile for the data center.
 - **Route Advertisement:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the route advertisement for the data center.
 - **Comment:** Enter additional information about the data center.
 - **Disabled:** Select the check box to disable this data center.
4. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

5. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

Modifying Data Centers

1. From the **Data Management** tab -> **Traffic Management** tab, select a data center and click the Edit icon.
2. The *Data Center* editor contains the following basic tabs from which you can modify data. For information about how to modify data, see [Adding Data Centers](#) on page 1204.
 - **General:** This tab displays the **Load Balancer**, **Partition**, **Name**, and **Description** of the data center. You can modify the rest of the fields.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the data center. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Save the configuration.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

MANAGING GLOBAL LOAD BALANCING POOLS

For sites with a large amount of incoming traffic, you can configure the GLB to distribute client requests to multiple servers instead of to the specified destination IP address in the client requests by creating a load balancing pool. A pool is a collection of virtual servers that can reside on multiple network servers. An individual pool member can belong to one or multiple load balancing pools, depending on how you want to manage your network traffic. When you add virtual servers to a pool, you can configure the GLB to direct traffic to a specific virtual server within a pool, using a specific load balancing method.

Creating Pools and Adding Pool Members

To add a global load balancing pool:

1. From the **Data Management** tab, select the **Traffic Management** tab, and then click **Add** -> **GLB Pool** from the Toolbar.
2. In the *GLB Pool* wizard, complete the following:
 - **Load Balancer Device/Group:** Click **Select** and select a GLB or a load balancer group from the *Load Balancer Device/Group Selector*.
 - **Partition:** Select a partition from the drop-down list.
3. Click **Next** and complete the following:
 - **Name:** Enter a name for the global load balancing pool.
 - **TTL:** Define the TTL (Time to Live) value by selecting the time and time unit from the drop-down lists.
 - **CNAME:** Enter the CNAME for the global load balancing pool.
 - **Prober Pool:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the prober pool.
 - **Comment:** Enter additional information about the pool.
 - **Disabled:** Select the check box to disable this pool.

If the GLB is configured as a Big-IP system redundant, enter peer addresses in the list.

4. Click **Next** to complete the health monitor information:

You can use health monitors to ensure that pool members are able to receive traffic. There are pre-defined monitors in the Available table you can use to associate with a pool, depending on the type of traffic you want to monitor. Use the arrows to move health monitors from the Available table to the Active table and vice versa. The appliance monitors the status of all monitors in the Active table.

- **Availability Requirements:** Define a minimum number of health monitors that report whether a pool member is available. From the drop-down list, select **All** to include all monitors in the Active table or select **At least** and then enter a number to define the minimum number of health monitors. Before the GLB can declare a pool member being up and running, the minimum number of health monitors must report that the pool member is available to receive traffic.

5. Click **Next** to complete load balancing methods:

- **Load Balancing Method:** Select the **Preferred**, **Alternate**, and **Fallback** load balancing methods from the drop-down lists. **Round Robin** is the default for the preferred and alternate method.
- **Fallback IPv4:** Enter the IPv4 addresses for the fallback method.
- **Fallback IPv6:** Enter the IPv6 addresses for the fallback method.

6. Click **Next** to add load balancing pool members:

Click the Add icon. In the *Load Balancer Virtual Server Selector*, select the virtual server you want to add to the pool. The appliance adds the selected pool to the table and displays the following:

- **Name:** The name of the virtual server.
- **Sever Name:** The server name of the pool member.
- **Address:** The IP address of the pool member.
- **Port:** The port used for communication on the pool member. The default is 443.
- **Ratio:** The ratio weight you want to assign to the pool member. The ratio weight determines the amount of traffic the pool member receives. You can click this field and modify the ratio. The GLB distributes connections among members according to the ratio weight. Enter a number between 0 to 65535. The default is 1.
- **Disabled:** Select the check box to disable this virtual server.

You can add multiple virtual servers to a pool. You can also select a pool member check box and click the Delete icon to delete the selected pool.

7. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

8. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

Modifying Load Balancing Pools

To modify a load balancing pool:

1. From the **Data Management** tab -> **Traffic Management** tab, select a pool and click the Edit icon.
2. The *GLB Pool* editor contains the following basic tabs from which you can modify data. For information about the fields, see [Creating Pools and Adding Pool Members](#) on page 1205.
 - **General:** This tab displays the **Load Balancer**, **Partition**, **Name** and **Description** of the pool. You can modify the rest of the fields.
 - **Health Monitors:** Displays the health monitors you have selected and the availability requirement. You can add or remove health monitors in this tab.
 - **Load Balancing Method:** Displays the preferred, alternate, and fallback load balancing methods you have selected. You can modify the methods in this tab.
 - **Pool Members List:** Displays the global load balancer pool members you added. You can add or delete pool members in this tab.

- **Extensible Attributes:** Add and delete extensible attributes that are associated with the pool. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Save the configuration.
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

MANAGING GLOBAL LOAD BALANCER SERVERS

Servers are physical devices on which you can configure one or more virtual servers. The servers that you define for the GLB to manage can include other GLBs and LLBs (Local Load Balancers). A server contains at least one virtual server. When you add virtual servers to a pool as a pool member, Grid Manager automatically picks up information from the associated servers. You can also assign health monitors with the server to report the status of that server.

Adding GLB Servers

To add a server:

1. From the **Data Management** tab, select the **Traffic Management** tab, and then click **Add -> GLB Server** from the Toolbar.
2. In the *GLB Server* wizard, complete the following:
 - **Load Balancer Device/Group:** Click **Select** and select a GLB or a load balancer group from the *Load Balancer Device/Group Selector*.
 - **Partition:** Select a partition from the drop-down list.
3. Click **Next** and complete the following:
 - **Name:** Enter a name for the global load balancer server.
 - **Product:** From the drop-down list, select the server type.
 - **Address List:** Click the Add icon. The appliance adds a row to the table. Click the row and enter the IP address of the server. You can also select a server check box and click the Delete icon to delete the address.
 - **Peer Address List:** This field is displayed only when you select BIG-IP System (Redundant) from the **Product** drop-down list. Click the Add icon. The appliance adds a row to the table. Click the row and enter the IP address of the server. You can also select a server check box and click the Delete icon to delete the address.
 - **Data Center:** From the drop-down list, select the data center to which the server belongs.
 - **Prober Pool:** This field appears only when the GLB runs TMOS v11.x. This does not appear for TMOS v10.x. From the drop-down list, select the prober pool for the data center.
 - **Comment:** Enter additional information about the server.
 - **Disabled:** Select the check box to disable this server.
4. Click **Next** to complete health monitor information:

You can use health monitors to ensure that servers are able to receive traffic. There are pre-defined monitors in the Available table you can use to associate with a server, depending on the type of traffic you want to monitor. Use the arrows to move health monitors from the Available table to the Active table and vice versa. The appliance monitors the status of all monitors in the Active table.

 - **Availability Requirements:** Define a minimum number of health monitors that report whether a server is available. From the drop-down list, select **All** to include all monitors in the Active table or select **At least** and then enter a number to define the minimum number of health monitors. Before the GLB can declare a server being up and running, the minimum number of health monitors must report that the server is available to receive traffic.

5. Click **Next** to complete virtual server discovery information:

- **Virtual Server Discovery:** Select one of the following from the drop-down list:
 - **Disabled:** Select this to disable the automatic discovery of virtual servers associated with this server. You can then add specific virtual servers to the Virtual Servers table later.
 - **Enabled:** Select this to enable the automatic discovery of virtual servers associated with this server. When you select this, you cannot add specific virtual servers. The appliance automatically detects them.
 - **Enabled (No Delete):** Select this to enable the discovery of virtual servers associated with this server. Once the virtual servers are detected, you cannot delete them.

When you select **Disabled** for virtual server discovery, you can add specific virtual servers. Click the Add icon. The appliance adds a row to the Virtual Servers table. Click the row and complete the following:

- **Name:** Enter the virtual server name.
- **Address:** Enter the IP address of the virtual server.
- **Service Port:** Enter the service port number of the virtual server.
- **Translation:** Enter the IP address of for NAT translation.
- **Translation Port:** Enter the NAT port number.
- **Link Discovery:** A link is a logical representation of a router that connects your network to the Internet. Select one of the following from the drop-down list:
 - **Disabled:** Select this to disable the automatic discovery of links.
 - **Enabled:** Select this to enable the automatic discovery of links
 - **Enable (No Delete):** Select this to enable the discovery of links. Once the links are detected, you cannot delete them.

6. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

7. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

Modifying GLB Servers

To modify a server:

1. From the **Data Management** tab -> **Traffic Management** tab, select a GLB server and click the Edit icon.
2. The *GLB Server* editor contains the following basic tabs from which you can modify data. For information about the fields, see [Adding GLB Servers](#) on page 1207.
 - **General:** This tab displays the **Load Balancer**, **Partition**, **Name**, **Description**, and **Product** of the GLB server object. You can modify the rest of the fields.
 - **Health Monitors:** Displays the health monitors you selected and the availability requirements. You can add or remove health monitors in this tab.
 - **Virtual Servers:** Displays information about the virtual server discovery and virtual servers you have added. You can modify the information in this tab.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the GLB server. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Save the configuration.

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later**, and enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

MANAGING GLOBAL LOAD BALANCER VIRTUAL SERVERS

A virtual server is a traffic management object that is represented by an IP address and a port number that points to a service. A virtual server is often used to balance traffic among a pool of servers on the network. A virtual server must be attached to a server. When you assign a virtual server to a load balancing pool, the GLB directs traffic coming into the virtual server to one of the pool members. Virtual servers increase the availability of resources for processing client requests.

Virtual servers also treat varying types of traffic differently, depending on your traffic-management needs. For example, a virtual server can enable compression on HTTP request data as it passes through the GLB, or decrypt and re-encrypt SSL connections and verify SSL certificates. A virtual server can also apply iRules you have selected for an LBDN object.

Adding GLB Virtual Servers

To add a GLB virtual server:

1. From the **Data Management** tab, select the **Traffic Management** tab, and then click **Add** -> **GLB Virtual Server** from the Toolbar.
2. In the *GLB Virtual Server* wizard, complete the following:
 - **Load Balancer Device/Group:** Click **Select** and select a GLB or a load balancer group from the *Load Balancer Device/Group Selector*.
 - **Partition:** Select a partition from the drop-down list.
 - **GLB Server:** Click **Select GLB Server** to select a GLB server from the *Load Balancer Server Selector*. A virtual server must be associated with a GLB server.
3. Click **Next** and complete the following:
 - **Name:** Enter a name for the GLB virtual server.
 - **Service:** Enter the following:
 - **Address:** Enter the IP address of the service provided by the virtual server.
 - **Protocol:** Select the protocol used for the service.
 - **Port:** Enter the port number for the service.
 - **Translation:** Enter the following:
 - **Address:** Enter the IP address for NAT, if any.
 - **Protocol:** Select the protocol used for NAT.
 - **Port:** Enter the port number for NAT.
 - **Comment:** Enter additional information about the load balancer object.
 - **Disabled:** Select the check box to disable this virtual server.
4. Click **Next** to complete the health monitor information:

You can use health monitors to ensure that servers are able to receive traffic. There are pre-defined monitors in the Available table you can use to associate with a server, depending on the type of traffic you want to monitor. Use the arrows to move health monitors from the Available table to the Active table and vice versa. The appliance monitors the status of all monitors in the Active table.

 - **Availability Requirements:** Define a minimum number of health monitors that report whether a virtual server is available. From the drop-down list, select **All** to include all monitors in the Active table or select **At least** and then enter a number to define the minimum number of health monitors. Before the GLB can declare a server being up and running, the minimum number of health monitors must report that the server is available to receive traffic.
5. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
6. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

Modifying GLB Virtual Servers

To modify a GLB virtual server:

1. From the **Data Management** tab -> **Traffic Management** tab, select a GLB virtual server and click the Edit icon.
2. The *GLB Virtual Server* editor contains the following basic tabs from which you can modify data. For information about the fields, see [Adding GLB Virtual Servers](#) on page 1209.
 - **General:** This tab displays the **Load Balancer, Partition, GLB Server, Name, Description,** and **Service** of the GLB virtual server object. You can modify the rest of the fields.
 - **Health Monitors:** Displays the health monitors you selected and the availability requirements. You can add or remove health monitors in this tab.
 - **Extensible Attributes:** Add and delete extensible attributes that are associated with the GLB virtual server. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#) on page 20.
3. Save the configuration.
or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later**, enter a date, time, and time zone. For information, see [Scheduling Tasks](#) on page 75.

PART 9 INFOBLOX INFRASTRUCTURE SECURITY

This section contains information about Infoblox infrastructure security features. Infoblox offers a few security features that you can use to control user access to certain core service resources and to protect your network against certain DNS threats. Based on your network requirements, you can enable Infoblox Advanced DNS Protection or Infoblox DNS Firewall to implement policy controls. You can also set up DNS blacklists or configure a security banner. Following is a list of security features:

- **Advanced DNS Protection**

The Infoblox Advanced DNS solution employs hardware-accelerated security rules to detect, report upon, and stop DoS (Denial of Service), DDoS (Distributed Denial of Service) and other network attacks targeting DNS caching and authoritative applications. This feature helps minimize “false positives” and ensures that your mission-critical DNS services continue to function even when under attack. For more information, see [Infoblox Advanced DNS Protection](#) on page 1213.

- **Infoblox DNS Firewall**

Infoblox DNS Firewall uses DNS RPZ (Response Policy Zones), a technology developed by the ISC (Internet System Consortium) for allowing reputable sources to dynamically communicate reputation domain names so you can implement policy controls for DNS lookups. For more information, see [Infoblox DNS Firewall](#) on page 1231.

- **Access Control (Named ACLs)**

To effectively manage your core network services, you can grant legitimate hosts access to specific operations on the appliance using an ACL (access control list) or anonymous ACEs (access control entries). You can also configure a named ACL and apply it to multiple operations, such as file distribution and DNS zone transfers. For more information, see [Configuring Access Control](#) on page 306.

- **DNS blacklists**

Your organization can prevent customers or employees from accessing certain Internet resources, particularly web sites, by prohibiting a recursive DNS member from resolving queries for domain names that you specify. You can configure a recursive DNS member to redirect the DNS client to predefined IP addresses or return a REFUSED response code (indicating that resolution is not performed because of local policy), depending on the domain name. For more information, see [About Blacklists](#) on page 579.

- **Security Banner**

You can configure and publish a notice and consent banner as the first login screen that includes specific terms and conditions you want end users to accept before they log in to the Infoblox Grid. When you enable the notice and consent banner, users must accept the terms and conditions displayed on the consent screen before accessing the login screen of Grid Manager. For more information, see [Configuring Notice and Consent Banner](#) on page 268.



Chapter 41 Infoblox Advanced DNS Protection

This chapter describes the Infoblox Advanced DNS Protection solution and its features. It explains how to enable and disable the threat protection service, define threat protection rule settings, and manage threat protection rules. It contains the following sections:

- [About Infoblox Advanced DNS Protection](#) on page 1214
 - [Configuring Advanced DNS Protection](#) on page 1215
 - [License Requirements](#) on page 1216
 - [Administrative Permissions](#) on page 1216
 - [Starting and Stopping Threat Protection Service](#) on page 1216
- [Understanding Threat Protection Rules](#) on page 1217
 - [System and Auto Rules](#) on page 1217
 - [Custom Rules](#) on page 1218
 - [About Rule Versions](#) on page 1220
 - [Using the Events Per Second Rule Setting](#) on page 1220
- [Configuring Grid Security Properties](#) on page 1221
 - [Enabling Multiple DNS Requests through a Single TCP Session](#) on page 1222
- [Creating Custom Rules](#) on page 1222
- [Managing Threat Protection Rules](#) on page 1223
 - [Viewing Threat Protection Rules](#) on page 1223
 - [Enabling and Disabling Rules](#) on page 1224
 - [Manually Uploading Rule Updates](#) on page 1224
 - [Publishing Rule Updates](#) on page 1225
 - [Modifying System and Auto Rules](#) on page 1226
 - [Modifying Custom Rules](#) on page 1226
- [Monitoring Threat Protection Events](#) on page 1227
 - [Monitoring through Syslog](#) on page 1227
 - [Threat Protection Statistics Widget](#) on page 1228
 - [Threat Protection Reports](#) on page 1229
- [DNS and Network-Flood Threats](#) on page 1229

ABOUT INFOBLOX ADVANCED DNS PROTECTION

The Infoblox Advanced DNS Protection solution employs hardware-accelerated security rules to detect, report upon, and stop DoS (Denial of Service), DDoS (Distributed Denial of Service) and other network attacks targeting DNS caching and authoritative applications. This feature helps minimize “false positives” and ensures that your mission-critical DNS services continue to function even when under attack. For information about possible DNS threats, see [DNS and Network-Flood Threats](#) on page 1229.

Advanced DNS Protection is designed to provide visibility and protection against network floods and DNS threats. It detects DNS attacks through predefined and custom threat protection rules, and mitigates DNS threats by dropping problematic packets while responding only to legitimate traffic. With valid licenses installed, you can subscribe to automatic rule updates that deliver near real-time protection against new and emerging attacks. You may also manually perform the rule update process based on your configuration.

Advanced DNS Protection runs on Infoblox Advanced Appliances that support subscriber-facing DNS caching and external DNS authoritative applications. With valid licenses installed, Advanced DNS Protection supports both IPv4 and IPv6 and can be enabled on the Infoblox-4030 Rev-2 and the following Infoblox Advanced Appliances: PT-1400, PT-2200, and PT-4000. For more information about these appliances, refer to the respective installation guides on the Infoblox Support site at <https://support.infoblox.com>, under the **Technical Documentation** tab.

Infoblox Advanced Appliances support all existing DNS features that are applicable to DNS caching and authoritative applications, except the following:

- Configuration of multiple interfaces on the same subnet
- HA configuration
- VLAN interfaces and VLAN tagging
- DSCP support on the Infoblox-4030 Rev-2 appliance
- 10/100-Mbps gigabit Ethernet mode and fixed speed/duplex settings
- Dynamic DNS updates (DDNS traffic will be blocked)

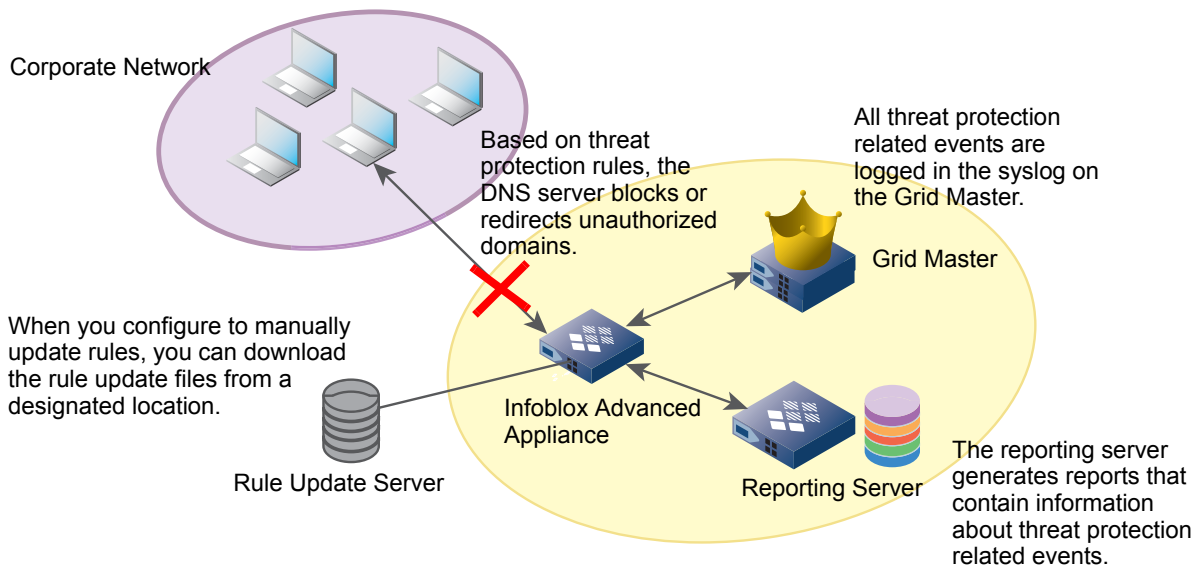
Note: Even though you can configure static routes on the Infoblox-4030 Rev-2 appliance when DNS cache acceleration is enabled, cached DNS responses are always sent through the interface on which the queries arrive, not the interface that is configured for the static route.

Consider the following when threat protection service is enabled on the Advanced Appliance:

- Protected interfaces (LAN1 and LAN2) are limited to DNS traffic, protocols in support of DNS anycast (BGP and OSPF) and the standard IP protocols such as ICMP, as well as connections to NTP servers.
- The MGMT interface is used for other traffic, such as Grid, SSH, SNMP, NTP, and it will not be protected against DDoS attacks.
- You cannot run other services, such as FTP, TFTP, HTTP, and DHCP, on the advanced appliance.
- The appliance terminates TCP connections for incoming DNS requests after handling the initial request through each TCP connection. The exception for this default Grid setting is for an SOA query sent by a client that is accepted in the allow-transfer ACL. In the case of an SOA query, the TCP connection remains open for subsequent DNS requests. This exception also covers the case in which an AXFR query follows the SOA query through the same TCP connection. For more information about how to override this default Grid setting, see [Enabling Multiple DNS Requests through a Single TCP Session](#) on page 1222.

Advanced DNS Protection supports a set of predefined threat protection rules that detect and mitigate possible DNS threats. You can modify some of the parameters and assign actions such as logging events and applying mitigation to these rules. You can also create custom rules to suit your security needs. As illustrated in [Figure 41.1](#) on page 1215, the Infoblox Advanced Appliance acting as an authoritative DNS server is added to the Grid. After installing valid threat protection licenses and configuring the appliance to serve as an advanced appliance, it can now detect DNS threats and mitigate DNS threats based on threat protection rules. All threat protection related events, conformed to CEF (Common Event Format), are logged in the syslog on the Grid Master. To perform further investigation about possible threats, the reporting server generates specific threat protection related reports. For information about how to monitor threat protection related events and reports, see [Monitoring Threat Protection Events](#) on page 1227.

Figure 41.1 Infoblox Advanced DNS Protection Solution



Configuring Advanced DNS Protection

To enable and configure Advanced DNS Protection on supported Infoblox appliances, complete the following:

1. Obtain valid Threat Protection and Threat Protection Update licenses from Infoblox and install them on the Infoblox Advanced Appliance. For information about license requirements, see [License Requirements](#) on page 1216.
2. Enable threat protection service, as described in [Starting and Stopping Threat Protection Service](#) on page 1216.
3. Configure threat protection rule settings for the Grid, including automatic rule updates, as described in [Configuring Grid Security Properties](#) on page 1221.
4. Optionally, you can do the following:
 - Override the default Grid setting that disables multiple DNS requests through one TCP session, as described in [Enabling Multiple DNS Requests through a Single TCP Session](#) on page 1222.
 - Modify system rules, as described in [Modifying System and Auto Rules](#) on page 1226.
 - Create custom rules from rule templates, as described in [System and Auto Rule Categories](#) on page 1217.

After you have successfully set up Advanced DNS Protection, you can do the following:

- View the current threat protection rules, as described in [Viewing Threat Protection Rules](#) on page 1223.
- Modify system and custom threat protection rules, as described in [Managing Threat Protection Rules](#) on page 1223.
- Manually upload rule updates, as described in [Manually Uploading Rule Updates](#) on page 1224.
- Publish uploaded rule updates, as described in [Publishing Rule Updates](#) on page 1225.
- Monitor threat protection related events and reports, as described in [Monitoring Threat Protection Events](#) on page 1227.

License Requirements

You must install the **Threat Protection** and **Threat Protection Update** licenses on the Infoblox Advanced Appliance before you can use the Advanced DNS Protection feature. With valid licenses installed, the Infoblox Advanced Appliance can be used only as a DNS caching or DNS authoritative server. You can join the appliance to the Grid and treat it as a Grid member. Note that if you install a Threat Protection license on a member, you can enable threat protection only on this member.

To receive initial and subsequent threat protection rules and rule updates, you must have the **Threat Protection Update** license installed. You can then configure NIOS to automatically download and publish threat protection rules or you can manually complete the process. For information, see [Manually Uploading Rule Updates](#) on page 1224 and [Publishing Rule Updates](#) on page 1225.

Contact your Infoblox representative to obtain the Threat Protection and Threat Protection Update licenses. Note that you cannot install temporary licenses for this feature. For information about licenses, see [Managing Licenses](#) on page 377.

Administrative Permissions

Superusers can configure all threat protection related tasks. You can assign **Security Permissions** to specific admin groups and roles. You can also add a global permission for managing Grid security properties or add an object permission for managing member security properties. For more information about security permissions, see [Administrative Permissions for DNS Threat Protection](#) on page 217.

Starting and Stopping Threat Protection Service

After you install the Threat Protection licenses on the appliance, you can start the threat protection service so you can monitor and mitigate DNS threats on that appliance.

Note: Before you enable this service, you must properly configure the MGMT port to support the threat protection function. For information about how to configure the MGMT port, see [Using the MGMT Port](#) on page 359.

To start or stop threat protection service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Services** tab -> *member* check box.
2. From the Toolbar, click **Start** to start the service or **Stop** to stop the service.

Note that when you stop threat protection service, the appliance does not provide visibility or protection against network floods or DNS threats.

After you enable threat protection service, you can configure rule settings, add custom rules, and evaluate system rules to ensure that mitigation to DNS threats is handled properly. You can also temporarily disable the threat protection service when necessary. For information about how to configure Grid security settings, see [Configuring Grid Security Properties](#) on page 1221.

Note: Starting and stopping threat protection service on the Infoblox-4030 Rev-2 appliance may trigger a product restart.

UNDERSTANDING THREAT PROTECTION RULES

Advanced DNS Protection supports the following threat protection rules:

- Predefined system and auto-generated rules, as described in [System and Auto Rules](#) on page 1217.
- Custom rules, as described in [Custom Rules](#) on page 1218.

Each threat protection rule belongs to a rule category. When you first import rules, the appliance publishes the system and auto rules in their respective categories. It also provides rule templates for creating custom rules. These rule categories and templates are displayed in the **Data Management** tab -> **Security** tab -> **Grid Rules** panel. To view threat protection rules, see [Viewing Threat Protection Rules](#) on page 1223.

NIOS automatically manages rule categories and you cannot add, delete, or modify them. NIOS may add new categories or remove old categories during a rule update. These actions are performed without intervention after the updates are authorized or automatically executed. You cannot add or delete system and auto rules, but you can create custom rules through predefined rule templates.

To obtain initial rules and subsequent rule updates, you can configure the appliance to automatically download and publish these rules or you can download them from the Infoblox Support web site and apply them manually. For information about how to configure rule settings, see [Configuring Grid Security Properties](#) on page 1221. Note that only the Grid Master in the Grid receives rules and rule updates. Grid member receives rules and updates through standard Grid replication from the Grid Master.

When rule updates occur, new rules are added while old ones can be removed. When an old rule is used by the Grid or a member, NIOS stores both the current and new versions of the same rule so you can switch between them. For more information about rule versions, see [About Rule Versions](#) on page 1220.

Note: Rules, rule templates and categories that are removed through updates are permanently deleted and cannot be restored from the NIOS Recycle Bin. Users can recover custom rules from the Recycle Bin, if enabled.

System and Auto Rules

System rules are predefined threat protection rules that are built into the advanced appliances. New system rules are added through rule updates. You can enable an entire category of system rules as well as individual rules. Note that you cannot add or delete system rules, though you can change some parameters. For most system rules, you can modify the **Active Version**, **Action** and **Log Severity**. For more information, see [Modifying System and Auto Rules](#) on page 1226.

Auto-generated rules are firewall rules that are automatically defined by NIOS for blocking traffic for disabled services and ports. They do not support functionality such as rate limiting. These rules can be grouped into different rule categories and are enabled or disabled by default. You cannot enable or disable auto rules in this release of Advanced DNS Protection, though you can set log severity and control logging for these rules.

System and Auto Rule Categories

The appliance supports the following system and auto rule categories. For detailed descriptions about each system and auto rule, see [Appendix H, "Threat Protection Rules"](#), on page 1371.

- **BGP:** Contains auto rules that mitigate attacks that target BGP (Border Gateway Protocol) routing parameters, such as invalid attribute lengths or invalid message types.
- **DNS Amplification and Reflection:** Contains system and auto rules that can be used to mitigate the commonly used methods of DDoS attacks. For information about DNS amplification and reflection, see [DNS Reflection and Amplification Attacks](#) on page 1230.
- **DNS Cache Poisoning:** Contains rate limiting rules that assign bandwidth restrictions rules used to mitigate DNS cache poisoning (on UDP and TCP) that is performed by sending a large volume of fake replies to a recursive server, which can result in hundreds or thousands of redirects. For more information about DNS cache poisoning, see [DNS Cache Poisoning](#) on page 1230.

- **DNS Malware:** Contains rules that protect against DNS malware that posts serious threats to the DNS infrastructure. For information about DNS malware threats, see [DNS Malware](#) on page 1230.
- **DNS Message Types:** Contains DNS system rules that can be used to filter requests that query specific DNS flags in the DNS message header.
- **DNS Protocol Anomalies:** Contains auto rules that address general DNS protocol attacks such as invalid DNS queries.
- **DNS Tunneling:** Contains auto rules that mitigate against DNS tunneling attacks. For more information, see [Inside-Out Attacks](#) on page 1229.
- **Default Drop:** Contains system rules that automatically drop IP packets when unusual UDP, TCP, and ICMP traffic is detected.
- **General DDoS:** Contains auto rules that address general DDoS (Distributed Denial of Service) attacks such as loopback address spoofing, and UDP or TCP packets that contain the same source and destination addresses.
- **ICMP:** Contains auto rules that mitigate ICMP and ICMPv6 ping attacks. ICMP ping size (for IPv4 and IPv6) for these rules is limited to 792 bytes. For information about ICMP, see [Internet Control Message Protocol \(ICMP\) Flood](#) on page 1229.
- **NTP:** Contains auto rules that mitigate attacks that target the NTP (Network Time Protocol).
- **OSPF:** Contains auto rules that mitigate attacks that target OSPF (Open Shortest Path First) routing parameters, such as invalid attribute lengths or invalid message types.
- **Reconnaissance:** Contains auto rules that mitigate network reconnaissance attacks, in which unauthorized remote attackers attempt to access networks by exploiting network standards and communications.
- **TCP/UDP Floods:** Contains DNS system rules that are used to mitigate DNS TCP and UDP floods. For information about TCP/UDP floods, see [UDP DNS Flood](#) on page 1229.

Custom Rules

Based on your security needs, you can define custom rules using predefined rule templates. Custom rules are typically whitelisting and blacklisting rules. You can create up to 1,000 custom rules from each rule template offered by Advanced DNS Protection. You can add or delete custom rules for the Grid only. You cannot add or delete them for members, but you can enable, disable, and modify certain rule parameters at the member level.

When you create custom rules, NIOS automatically generates the rule ID from the template used. Note that custom rules do not support IDNs (Internationalized Domain Names). You must first convert IDNs into puny codes before entering the data. For information about how to create custom rules, see [Creating Custom Rules](#) on page 1222.

When you create custom rules, you are essentially creating whitelisting and blacklisting entries that utilize rate limiting to detect suspicious UDP and TCP traffic. Advanced DNS Protection supports a series of rule templates for defining new custom rules. For information about rule templates, see [Custom Rule Templates](#) on page 1219.

Whitelisting rules define a list of allowed resources before they are blocked by the configured rate limit settings. They provide for only a selected set of entities to access the protected environment. Examples include company offices and their associated internal network services, which presumably use access control systems to enforce them. In effect, addresses or networks that do not match the whitelisting entries are automatically blocked.

Blacklisting rules define a list of disallowed resources through FQDN lookups as well as rate limiting. Blacklists typically allow a far broader base of access to many more entities, and cite a list of specific entities or people that do not have access. Otherwise, any devices or users theoretically have access to the protected environment.

For whitelist entries, the matching values are mandatory, in which the IP address or network of the rule is expressly permitted access. Blacklist entries are forbidden, in which the IP address or network of the rule is expressly denied access. In essence, blacklisting is convenience while whitelisting is more secure.

Note: You can create DNS-specific blacklists under the **Data Management** tab → **DNS** tab. However, you cannot use this blacklist feature as part of DNS threat protection.

Custom Rule Templates

Advanced DNS Protection supports a few custom rule templates from which you can create new custom rules. Note that when you use a specific rule template to create custom rules, the new rules reside in their respective rule categories. For information about creating custom rules, see [Creating Custom Rules](#) on page 1222.

For each rule you create, you can define the **Events per second** value to determine the number of events per second that will be logged for the rule. You can also define specific rule parameters for custom rules, as follows:

Note: Custom rules do not support IDNs (Internationalized Domain Names). To use IDNs for custom rules, you must first convert the IDNs into puny codes. You can use the **IDN Converter** from the **Toolbar** for the conversion.

- **BLACKLIST FQDN lookup TCP:** Use this rule template to create custom rules for blacklisting DNS queries by FQDN lookups on TCP. In the Rule Parameters table, complete the following:
 - **Blacklisted FQDN:** Enter the FQDN that you want the appliance to block over TCP traffic. You can specify one FQDN for each custom rule.
- **BLACKLIST FQDN lookup UDP:** Use this rule template to create custom rules for blacklisting DNS queries by FQDN lookups on UDP. In the Rule Parameters table, complete the following:
 - **Blacklisted FQDN:** Enter the FQDN that you want the appliance to block over UDP traffic. You can also enter a list of FQDNs using semicolon as the separator.
- **BLACKLIST IP TCP Drop prior to rate limiting:** Use this rule template to create rules for blocking IPv4 or IPv6 addresses on TCP before the appliance drops the packets based on rate limiting rules you have defined using the **RATELIMITED IP TCP** template. In the Rule Parameters table, complete the following:
 - **Blacklisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are dropped before any relevant rate limiting rules take effect. Note that all TCP traffic from the specified IPv4 and IPv6 addresses and networks will be blocked. Enter network addresses in address/CIDR format.
- **BLACKLIST IP UDP Drop prior to rate limiting:** Use this rule template to create rules for blocking IPv4 or IPv6 addresses on UDP before the appliance drops the packets based on rate limiting rules you have defined using the **RATELIMITED IP UDP** template. In the Rule Parameters table, complete the following:
 - **Blacklisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are dropped before any relevant rate limiting rules take effect. Note that all UDP traffic from the specified IPv4 and IPv6 addresses and networks will be blocked. Enter network addresses in address/CIDR format.
- **RATELIMITED FQDN lookup UDP:** Use this rule template to create custom rules that contains rate limiting restrictions for blocking DNS queries by FQDN lookups on UDP traffic. In the Rule Parameters table, complete the following:
 - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of UDP traffic that consists of DNS lookups for the FQDN defined in this rule. The default is 5.
 - **Drop interval:** Enter the number of seconds for which the appliance drops packets.
 - **Blacklist rate limited FQDN:** Enter the FQDN that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this FQDN when the UDP traffic of DNS lookups for this FQDN exceeds the configured rate limit value.
- **RATELIMITED IP TCP:** Use this rule template to create custom rules that contains rate limiting restrictions for blacklisting IP addresses on TCP. If there are certain IP addresses that you want to block before its traffic reaches the rate limit restrictions, you can create a rule using the **BLACKLIST IP TCP Drop prior to rate limiting** template. In the Rule Parameters table, complete the following:
 - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of TCP traffic that consists of DNS lookups for the IP address or network defined in this rule. The default is 5.
 - **Drop interval:** Enter the time interval in seconds the appliance drops IP packets sent by the rate limited IP address or network defined for this rule. The default is 30 seconds.

- **Rate limited IP address/network:** Enter the IP address or network that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this IP address based on the drop interval when the TCP traffic of DNS lookups for this IP address exceeds the configured rate limit value.
- **RATELIMITED IP UDP:** Use this rule template to create custom rules that contains rate limiting restrictions for blacklisting IP addresses on UDP. If there are certain IP addresses that you want to block before its traffic reaches the rate limit restrictions, you can create a rule using the **BLACKLIST IP UDP Drop prior to rate limiting** template. In the Rule Parameters table, complete the following:
 - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of UDP traffic that consists of DNS lookups for the IP address or network defined in this rule. The default is 5.
 - **Drop interval:** Enter the time interval in seconds the appliance drops IP packets sent by the rate limited IP address or network defined for this rule. The default is 30 seconds.
 - **Rate limited IP address/network:** Enter the IP address or network that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this IP address based on the drop interval when the TCP traffic of DNS lookups for this IP address exceeds the configured rate limit value.
- **WHITELIST IP TCP Pass prior to rate limiting:** Use this rule template to create custom rules for allowing certain IP addresses on TCP before the appliance drops the packets based on rate limiting rules you have defined using the **RATELIMITED IP TCP** template. In the Rule Parameters table, complete the following:
 - **Whitelisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are allowed before any relevant rate limiting rules take effect.
- **WHITELIST IP UDP Pass prior to rate limiting:** Use this rule template to create custom rules for allowing certain IP addresses on UDP before the appliance drops the packets based on rate limiting rules you have defined using the **RATELIMITED IP UDP** template. In the Rule Parameters table, complete the following:
 - **Whitelisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are allowed before any relevant rate limiting rules take effect.

About Rule Versions

A rule update file contains system rules, auto rules, and rule templates that you use to create custom rules. You can configure the appliance to automatically download the ruleset or you can manually apply it. For information about how to configure this, see [Configuring Grid Security Properties](#) on page 1221.

Each rule consists of a rule ID, rule name, and description. NIOS also publishes the rule version indicating whether a rule has the latest version. NIOS stores at least two versions of the same rule: the current version and the latest version. You can switch between these versions for the same rule. However, if a rule is not enabled for the Grid or any members, NIOS stores only the latest version and removes the older version. If a rule is deprecated and is not enabled for the Grid or any members, NIOS removes it from the system. Otherwise, Grid Manager displays “**Deprecated**” as the rule version. For information about how to view threat protection rules, see [Viewing Threat Protection Rules](#) on page 1223.

Using the Events Per Second Rule Setting

The **Events per second** setting allows for disabling or throttling of event logs for specific threat protection rules.

Setting the **Events per second** parameter to zero disables logging for that rule. Setting the parameter to any other number enables threat protection logging for that specific rule. For information about how to configure this, see [Configuring Grid Security Properties](#) on page 1221.

CONFIGURING GRID SECURITY PROPERTIES

After you have installed valid threat protection licenses, you can configure rule update settings for the Grid. These settings apply to all members in the Grid. You can override only the global **Event per second** filter (in the **Basic** tab) and the **Disable multiple DNS requests via single TCP session** option (in the **Advanced** tab) in the *Member Security Properties* editor by selecting a member and clicking Edit.

To configure rule settings:

1. From the **Data Management** tab, select the **Security** tab, and then click **Grid Security Properties** from the Toolbar.
2. In the *Grid Security Properties* editor, select the **Threat Protection** tab -> **Basic** tab, and complete the following:
In the Threat Protection Ruleset Updates section, define the rule update policy. The appliance automatically performs rule updates by default. You can choose to manually publish rule updates. For information about how to manually update rules, see [Manually Uploading Rule Updates](#) on page 1224 and [Publishing Rule Updates](#) on page 1225.
 - **Last Published Rule Update:** Displays the version string of the last published rule update. This field changes each time when the ruleset is updated.
 - **Last Applied On:** Displays the timestamp and time zone when the last rule update was applied on the Grid. This field changes each time when a ruleset is applied.
 - **Rule Update Policy:** Select the rule update policy from the drop-down list to determine whether updates are being applied automatically or manually. When you select **Automatic**, the appliance automatically switches to the newly added rule version and publishes the changes when a rule update is applied. Select **Manual** to manually download rule updates and publish them. Note that you must have a valid Threat Protection Update subscription license installed in order to perform rule updates. For information about how to perform a manual update, see [Manually Uploading Rule Updates](#) on page 1224 and [Publishing Rule Updates](#) on page 1225.
 - **Enable Automatic Ruleset Downloads:** Select this to enable automatic downloads for rule updates.

Note: When you select this, ensure that you configure and enable a valid DNS resolver for the Grid in the *Grid Properties* editor so the appliance can successfully access the rule update file.

- **Test Connection:** Click this to test the connectivity between the advanced appliance and the server from which you receive the rule update files. Grid Manager displays a message indicating whether the connection is successful.
- **Download Rules Now:** Click this to immediately download the latest rule update file from the Infoblox rule update server, provided that the connection between the appliance and the server is successful.

In the Schedule section, define the schedule for automatic ruleset downloads. The following options are enabled only when you have selected **Enable Automatic Ruleset Downloads**:

- **Default:** Select this to set the default schedule settings for automatic ruleset downloads.
- **Custom:** Select this to schedule downloads at a later date and time. Click the Calendar icon to select the date and time.

Note: Automatic ruleset downloads are performed within 15 minutes before or after your configured time.

In the Threat Protection Logging section, define the events per second per rule value to allow the appliance to log events in the syslog:

- **Events per Second per Rule:** Specify the number of events logged per second per rule. The default value is one and the maximum value is 700. Setting the value to 0 (zero) disables the appliance from logging events for the rules. The appliance displays an error message when you enter a value greater than the maximum value. You can override this event filter at the member level.
3. Save the configuration. To publish changes, click **Publish** if it appears at the top of the screen. Note that NIOS does not require restarting of the threat protection service after for rule updates.

Enabling Multiple DNS Requests through a Single TCP Session

The advanced appliance inspects only one DNS request sent over a single TCP connection. To avoid accepting possible malicious data following a valid DNS request, the appliance terminates the TCP connection after handling the initial DNS request over TCP. You can modify this default Grid setting at the Grid or member level.

To modify this setting, do the following:

1. From the **Data Management** tab, select the **Security** tab, and then click **Grid Security Properties** from the Toolbar.
or
From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* check box, and then click the Edit icon.
2. In the *Grid Security Properties* or *Member Security Properties* editor, select the **Threat Protection** tab -> **Advanced** tab, and complete the following:
 - **Disable multiple DNS requests via single TCP session:** This is selected by default to avoid accepting possible malicious data following a valid DNS request. When this is selected, the appliance handles the initial DNS request through TCP and then terminates the TCP session to block subsequent DNS traffic, except for an SOA query sent by a client that is accepted in the allow-transfer ACL. This exception covers the case in which an AXFR query follows the SOA query through the same TCP connection.
3. Save the configuration.

CREATING CUSTOM RULES

Advanced DNS Protection provides a few rule templates from which you can create custom rules. For information about the list of rule templates that you can use, see [Custom Rule Templates](#) on page 1219.

To create a custom rule:

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab, and then click **Add Custom Rule** from the Toolbar.
2. In the *Add Custom Rule* editor, complete the following:
 - **Template:** From the drop-down list, select the blacklisting or whitelisting rule template from which you want to create the new rule. For more information about the rule templates, see [Custom Rule Templates](#) on page 1219.
 - **Description:** Displays the description of the rule that you are about to create. You cannot modify this.
 - **Comment:** Enter comments to describe the new rule.
 - **Disable:** Select this if you want to keep the new rule disabled for later use.
3. Click **Next** and complete the following to configure rule parameters:
 - **Description:** Displays the description of the rule that you are about to create. You cannot modify this.
 - **Active Version:** Select the rule version from the drop-down list if you are making changes to an existing rule. The default is 1.
 - **Action:** Displays the operation the appliance performs when an event related to this rule occurs. Some rules are restricted to specific actions. For example, the action for all blacklisting rules is set to **Drop**, where the appliance drops IP packets when such an event occurs. The action for all whitelisting rules is set to **Pass**, where the appliance passes IP packets when such an event occurs.
 - **Log Severity:** Select **Critical**, **Major**, **Warning** or **Informational**. The log severity you select here determine the severity of the message triggered by a match against the rule.

In the Rule Parameters section, do the following:

- Click the **Value** field for **Events per second** to enter the maximum number of events per second that will be logged for this rule. Use a nonzero value if you want matches against the current rule to log events. Setting a value to 0 (zero) disables the appliance from logging events associated with this rule.

- Depending on the template you have selected, click the **Value** field and enter the appropriate parameters to configure the rule. For descriptions about the parameters for each rule template, see [System and Auto Rule Categories](#) on page 1217.

4. Click **Save & Close**.

The new rule, with an automatically assigned rule ID, is created and placed in its corresponding rule category.

MANAGING THREAT PROTECTION RULES

You can modify any previously defined custom rule, or some of the parameters for system and auto rules. For most system and auto rules, you may change the **Active Version**, the **Action** and the **Log Severity**. You can also enable or disable individual rules or an entire category of rules.

If you have selected to manually update threat protection rules, you must download updated rules from the Infoblox Support web site and then publish them to the system.

You can do the following after the initial setup, including uploading the initial ruleset:

- Look at rules that are currently installed on your system, as described in [Viewing Threat Protection Rules](#) on page 1223.
- Enable and disable certain rules, as described in [Enabling and Disabling Rules](#) on page 1224.
- Upload rule updates to the system when you have selected to manually apply rule updates, as described in [Manually Uploading Rule Updates](#) on page 1224.
- Publish rule updates that you have uploaded to the system, as described in [Publishing Rule Updates](#) on page 1225.
- Modify existing system rules, as described in [Modifying System and Auto Rules](#) on page 1226.
- Modify custom rules, as described in [Modifying Custom Rules](#) on page 1226.

Viewing Threat Protection Rules

To view the current threat protection rules:

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
2. In the Grid Rules table, Grid Manager displays threat protection rules by categories. The **Category** column lists all the category to which rules belong.
3. To view rules listed in each category, expand the list by clicking the arrow beside the check box. You can view the following information for each rule:
 - **Rule ID:** The ID of the rule.
 - **Rule Name:** The name of the rule. This can contain up to 255 characters.
 - **Type:** The rule type. This can be **Custom**, **System**, or **Auto**. For more information about each rule type, see [Understanding Threat Protection Rules](#) on page 1217.
 - **Disabled:** Displays whether the rule is disabled. A disabled rule does not perform any mitigation functions.
 - **Comment:** Comments that were entered for the rule. This can contain up to 255 characters.
 - **Action:** The operation that the appliance performs when the event occurs. This can be one of the following:
 - **Alert:** The appliance passes the packets and logs the event.
 - **Drop:** The appliance drops the packets and logs the event.
 - **Pass:** The appliance passes the packets but does not log the event.
 - **Description:** Description about the rule. This can contain up to 255 characters.
 - **Log Severity:** Log severity level. This can be **Critical**, **Major**, **Warning**, or **Informational**.
 - **Active Version:** The current rule version. When new versions of the rule are added during rule updates, the appliance stores multiple versions of the rule so you can switch between versions.

- **Latest Version:** It displays one of the following:
 - **Yes:** The rule is the latest version.
 - **No:** The rule is not the latest version. For information about switching to the latest version, see [Managing Threat Protection Rules](#) on page 1223.
 - **Deprecated:** This rule is obsolete, but is being used for the Grid or one of the members. NIOS will remove this rule when it is no longer used for the Grid or members.

You can also do the following in this panel:

- Modify some of the data in the table. Double click a row, and modify the data. Click **Save** to save the changes. Note that some fields are read-only.
- Select the check box of a rule and click the Edit icon to modify the properties of the rule.
- Select the check box of a custom rule and click the Delete icon to delete a custom rule.
- Print or export the data.
- Use filters and the Go to function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the Go to field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For more information, see [Using Quick Filters](#) on page 68.

Using Filters

You can use the following system filters to filter threat protection rules in the Grid Rules panel. When you select a filter, Grid Manager displays only the specified rules. Using filters makes it easier to locate specific rules for editing, enabling, and disabling.

- **All Auto Generated Rules:** This option shows all auto-generated rules defined in the Advanced DNS Protection solution.
- **All Custom Rules:** This option shows all whitelisting and blacklisting custom rules defined by NIOS users.
- **All System Rules:** All protocol-specific rules associated with particular attack phenomena.

For more information about filters, see [Using Filters](#) on page 63.

Enabling and Disabling Rules

By default, all activated threat protection rules apply across the entire Grid. Enabling or disabling a rule category will enable or disable all rules contained in that category. You can also enable or disable individual rules.

To enable or disable all rules in a category, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab → **Security** tab → **Threat Protection Rules** tab.
2. Click the Action icon and choose **Enable All Rules in Category** or **Disable All Rules in Category**. Either option can be disabled depending on the current state of the rules in the category.

To enable or disable individual rules, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab → **Security** tab → **Threat Protection Rules** tab.
2. Click the arrow beside a rule category to open a Rule category.
3. Click the Action icon for any rule and choose **Enable** or **Disable** from the menu. Either option can be disabled depending on the current state of the rule.

Manually Uploading Rule Updates

You can update threat protection rules anytime when you select to manually perform rule updates. You can choose to download rule updates but not immediately deploy them. NIOS archives and tracks the existing and new versions for the same rule, allowing for switching between these versions when necessary. After uploading the rules, you can apply them by publishing them to the system. For more information, see [Publishing Rule Updates](#) on page 1225.

Rule updates do not require restart of the DNS or threat protection service in the Grid, and they do not affect ongoing services. However, the appliance deploys updated rules only when you publish the changes. Note that all threat protection rule update events are logged in the syslog on the Grid Master only.

Note: By default, threat protection rule updates are automatic. Infoblox recommends that you retain this setting. For information about how to configure this setting, see [Configuring Grid Security Properties](#) on page 1221.

To manually upload a rule update file:

1. From the **Grid** tab, select the **Grid Manager** tab → **Security** tab, and then click **Ruleset Upload** from the Toolbar.
2. In the *Ruleset Upload* dialog box, do the following:
 - **File:** Click **Select** to navigate to the file location, and then click **Upload**. Grid Manager displays the file name in this field.
 - Click **Test** to check the changes that will occur during the rule update, without actually applying the update. You can view the update details in the Syslog Viewer. The appliance preserves the uploaded file if you do not click **Update** to update the rules. When you manually upload rule updates the next time, this file will be displayed in the dialog. You can then choose to apply the update from this file or upload a new file before performing the update.
 - Click **Update** to update the rules.
 - Click **View Update Results** to view the updated rules in the Syslog Viewer. All threat protection rule updates are logged in the syslog on the Grid Master.

Publishing Rule Updates

You can publish rule updates at any time after you have uploaded the rules. For information about uploading ruleset files, see [Manually Uploading Rule Updates](#) on page 1224.

To publish rule updates:

1. From the **Grid** tab, select the **Grid Manager** tab → **Security** tab, and then click **Publish Changes** from the Toolbar.
2. In the *Publish Changes* dialog box, complete the following:
 - **Publish Changes on all Members:** Select one of the following:
 - **Simultaneously:** Publish changes on all of the members in the Grid at the same time.
 - **Sequentially:** Publish changes on each Grid member according to the number of seconds you enter in the **Sequential every (seconds)** field. For example, if you enter every 10 seconds, the system update changes on the first member, and 10 seconds later on the second member. This is the default option.
 - **Impacted Members and Services:** Click the Poll Members icon to display the affected members in that Grid. Grid Manager displays the member names and whether each member is configured for the threat protection service:
 - **Yes:** The service is active and the system will publish rule updates on this member upon execution of this task.
 - **No:** The service is not active and the system will not publish rule updates on this member.
 - **Disabled:** The service is currently disabled on this member.

To schedule this task, click the Schedule icon at the top of the dialog box. In the Schedule Change panel, complete the following:

- **Now:** Publish rule updates upon clicking **Publish**.
- **Later:** Enter the following information to schedule publishing updates at a certain date and time:
 - **Start Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
 - **Start Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. When you enter the time in a 24-hour format such as 23:00, Grid Manager displays 11:00:00 PM. You can also select a time from the drop-down list by clicking the time icon.

- **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
3. Click **Publish** to publish changes immediately or click **Schedule Publish** to schedule the publish.

Modifying System and Auto Rules

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
2. In the Grid Rules table, expand the category to which the rule belongs, select the check box, and then click the Edit icon.
3. In the *System Rule* or *Auto Generated Rule* editor, select the **General** tab -> **Basic** tab to modify the following:
 - **Comment:** Enter information about the system rule.
 - **Disable:** Select this check box to disable the system rule. You cannot disable auto rules.

You can also view the following information (but you cannot modify it):

- **Rule ID:** Displays the rule ID of the system rule.
 - **Name:** Displays the name of the rule.
 - **Category:** Displays the category to which the system rule belongs.
 - **Description:** Displays the description of the system rule.
4. In the *System Rule* or *Auto Generated Rule* editor, select the **Settings** tab -> **Basic** tab. Depending on the rule, you may or may not be able to modify the following:
 - **Active Version:** Select the revision of the system rule from the drop-down list. Every rule provides a list of its currently stored versions. Though Infoblox recommends keeping the most recent version active for any rule, older versions can be restored to activity if an issue is found in the latest version or for testing purposes.
 - **Action:** Select one of the following: **Drop** or **Pass**. Some rules are restricted to specific actions. For example, the action for all blacklisting rules is set as **Drop** where the appliance drops the packets and logs the activity when such an event occurs. The action for all whitelisting rules is set as **Pass**, where the appliance silently passes the packets without logging when such an event occurs.
 - **Alert:** Logs the activity, and passes the packet.
 - **Drop:** Logs the activity and drops the packet.
 - **Pass:** Silently passes the packet without logging.
 - **Log Severity:** Select the log severity level from the drop-down list. You can select **Critical**, **Major**, **Warning**, or **Informational**. The selection here corresponds to the severity levels you can configure for logging in the syslog.
 - **Rule Parameters:** In the Rule Parameters table, the **Description** column displays the rule parameters. Click the row and enter the corresponding values for the rule parameters in the **Value** column. Depending on the rule, this table displays only the parameters that are relevant to the system or auto rule.

You can also view the following information (but you cannot modify it):

- **Latest Version:** It displays one of the following:
 - **Yes:** The rule is the latest version.
 - **No:** The rule is not the latest version. You may want to retrieve the latest version.
 - **Deprecated:** This rule is now obsolete, but is used for the Grid or member. NIOS will remove this rule if it is no longer used for the Grid or member.
5. Save the configuration.

Modifying Custom Rules

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
2. In the Grid Rules table, expand the category to which the rule belongs, select the check box, and then click the Edit icon.

3. In the Custom Rule editor, select the **General** tab -> **Basic** tab to modify the following:
 - **Comment:** Enter information about the custom rule.
 - **Disable:** Select this check box to disable the custom rule.

You can also view the following information (but you cannot modify it):

 - **Template:** Displays the name of the template the custom rule uses.
 - **Rule ID:** Displays the rule ID of the custom rule.
 - **Name:** Displays the name of the rule.
 - **Category:** Displays the category to which the custom rule belongs.
 - **Description:** Displays the description of the custom rule.
4. In the Custom Rule editor, select the **Settings** tab -> **Basic** tab to modify the following:
 - **Active Version:** Select the revision of the custom rule from the drop-down list. Every rule provides a list of its currently stored versions. Though Infoblox recommends keeping the most recent version active for any rule, older versions can be restored to activity if an issue is found in the latest version or for testing purposes.
 - **Log Severity:** Select the log severity level from the drop-down list. You can select **Critical**, **Major**, **Warning**, or **Informational**. Log severity may have an effect on how other Grid services respond to particular events. The selection here corresponds to the severity levels you can configure for logging in the syslog.
 - **Rule Parameters:** In the Rule Parameters table, the **Description** column displays the rule parameters. Click the row and enter the corresponding values for the rule parameters in the **Value** column.

You can also view the following information (but you cannot modify it):

 - **Latest Version:** It displays one of the following:
 - **Yes:** The rule is the latest version.
 - **No:** The rule is not the latest version. You may want to retrieve the latest version.
 - **Deprecated:** This rule is now obsolete, but is used for the Grid or member. NIOS will remove this rule if it is no longer used for the Grid or member.
 - **Action:** Displays the operation which the appliance performs when this event occurs. Some rules are restricted to specific actions. For example, the action for all blacklisting rules is set as **Drop** where the appliance drops the packets and logs the activity when such an event occurs. The action for all whitelisting rules is set as **Pass**, where the appliance silently passes the packets without logging when such an event occurs.
5. Save the configuration.

MONITORING THREAT PROTECTION EVENTS

You can monitor threat protection events through the following:

- Syslog, as described in [Monitoring through Syslog](#) on page 1227.
- Status Dashboard, as described in [Threat Protection Statistics Widget](#) on page 1228.
- Threat protection report, as described in [Threat Protection Reports](#) on page 1229.

Monitoring through Syslog

To receive threat protection events in the syslog, you must enable the **Security** option in the **DNS logging category** of the *Grid DNS Properties* editor. For information about configuring the logging category, see [Setting DNS Logging Categories](#) on page 1015. Once the **Security** option is enabled, Infoblox advanced appliances log each threat protection related event in the syslog in CEF (Common Even Format). You can get detailed information about the events by reviewing the syslog periodically. For information about how to configure the syslog server, see [Using a Syslog Server](#) on page 1012.

When a DNS attack is detected against an enabled rule, the appliance generates a log message. Note that only threat protection messages in CEF are displayed in the syslog.

Example:

When the appliance detects ICMP ping attacks that exceed the pint size against an existing auto rule that has the following configuration:

Log Severity = Critical
 Rule ID = 5090004
 Rule Name = ICMP Large Packet
 Rule Action = Drop
 Rule Category = ICMP

it generates the following threat detection event log message:

```
2013-06-24T02:36:48+00:00 daemon (none) threat-protect-log[11069]: crit
CEF:0|Infoblox|NIOS THREAT|6.10.0-223885|5090004|ICMP Large Packet|8|src=10.32.1.35
spt=0 dst=10.34.82.74 dpt=53 act="DROP" cat="ICMP"
```

The number of log messages generated is based upon your **Event per Second per Rule** setting. For example, if the setting is **5**, the appliance generates five log messages of the same event per second when the attack continues within the time duration. Each log message contains the following information:

- The timestamp when the event happened in `yyyy-mm-ddThh:mm:ss+00:00` format.
- **Infoblox|NIOS Threat|x.x.x**: Indicates the Infoblox product, and x.x.x represents the NIOS version.
- The number following the NIOS version is the rule ID. In this example, it is 5090004.
- Following the rule ID is the rule name specified in the rule.
- The number following the rule ID is the log severity. The following numbers indicate the severity levels:
 - **8 = Critical**
 - **7 = Major**
 - **6 = Warning**
 - **4 = Informational**
- **src**: Source IP address.
- **spt**: Source port.
- **dst**: Destination IP address.
- **dpt**: Destination port.
- **act**: The rule action, which can be **ALERT**, **DROP**, or **PASS**, depending on the rule configuration.
- **cat**: The rule category to which the rule belongs. In this example, the rule category is **"ICMP."** For information about rule categories, see [System and Auto Rule Categories](#) on page 1217.

To view DNS threat protection related log messages:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right corner, select the Grid member on which you want to view the syslog.
3. From the Quick Filter drop-down list, select **Threat Rule Update Events** or **Threat Detection Event Logs** to view rule update events or threat detection events respectively. To narrow down the system messages you want to view, click Show Filter and then select the filters you want to use. For information about how to use filters, see [Using Filters](#) on page 67.

Threat Protection Statistics Widget

You can also get a high-level view of the threat protection events through the Status Dashboard. The advanced appliances provide the *Threat Protection Statistics* widget so you can monitor the trend and counts of the various events. For more information about the Dashboard and the *Threat Protection Statistics* widget, see [Status Dashboards](#) on page 116.

Threat Protection Reports

NIOS provides a series of reports to monitor and analyze DNS threat protection events. When you integrate a reporting in your Grid, you can get the threat protection related reports so you can monitor event counts by severity, member, rule, and rule category. For detailed information about these reports, see [Security](#) on page 1166.

DNS AND NETWORK-FLOOD THREATS

DNS is a tempting target for attacks given that it is a core Internet service. Attackers can send malformed DNS queries or DNS responses to the targeted server, hoping to exploit bugs in its DNS implementation. Other variants include code insertion, buffer overflows, memory corruption, NULL pointer dereferencing, and specific vulnerability exploits. DNS attacks tend to follow specific patterns but can be difficult to deal with using only rate-limiting techniques, because of the sheer scale of many recent attacks. DNS threat protection is designed to grow and expand over time, through threat protection rule updates, to deal with both outside-in and inside-out attacks on network infrastructure and Internet services.

Following are some of the network-flood attacks that can target your DNS caching and authoritative servers:

Internet Control Message Protocol (ICMP) Flood

An ICMP flood attack is also known as a ping attack in which attackers send a large number of ICMP ping packets to a DNS server repeatedly in order to hinder the server's ability to respond to other requests. It can also be an attempt to send a large number of ping packets to the broadcast IP of a subnetwork, otherwise known as a Smurf attack, as a basic means of amplifying an attack across more hosts than a normal ping would typically permit. These types of attacks can be dealt with by setting a policy to disallow pings to the broadcast IP on the network.

Note: When threat protection is enabled, ICMP ping size (for IPv4 and IPv6) is limited to 16,000 bytes.

SYN Flood

A host sends a long stream of TCP SYN packets, frequently using a forged sender address. Because TCP regards a SYN packet as part of a legitimate connection request, the requested server starts a half-open connection by responding with a SYN ACK packet. Since the sender address is faked, the final ACK response from the sender never comes, and the half-open TCP socket closes only after a time out interval. A massive wave of SYN requests with fake senders can wipe out the connection resources of a network device, effectively locking it away from legitimate users.

UDP DNS Flood

UDP Flood is a denial-of-service attack that uses the connectionless UDP transport protocol and attempts to send large numbers of packets to random UDP protocol ports on a remote system, or to a specific protocol. UDP flood is a reflection attack that is often used for attacking DNS servers operating on UDP port 53. UDP flooding typically uses IP spoofing, in which the sender address is faked. The purpose is to occupy so many resources on the target that it can no longer provide its services on the network.

Inside-Out Attacks

A sophisticated form of “phishing” in which an attacker is able to inject a worm or other piece of attack software onto a host machine, which thereupon captures sensitive information such as logins, and adds that data to DNS queries that can be sent from the trusted machine to an untrusted entity for collection. DDoS Security detects data leaks of this type, logs the incident, and funnels the suspect packets to a quarantine location. In a similar vein, **DNS Tunneling** uses DNS as a covert channel to avoid firewall and IPS security mechanisms. Tunneling encapsulates Inbound and outbound packets inside DNS requests and DNS responses.

DNS Fluxing

the use of a system called a Domain Generation Algorithm botnet to perform one of the following attacks:

- **Fast Fluxing:** Forcing rapid swapping in and out of IP addresses, with extremely high frequency through changing DNS records with brief TTLs
- **Domain Fluxing:** Forcing constant changing of and allocation of multiple fully-qualified domain names (FQDNs) to a single IP address on the recursive or authoritative DNS server.

DNS Cache Poisoning

With cache poisoning, attackers attempt to insert a spoofed DNS response to a DNS resolver, which then stores the response in its cache, where it lives until the TTL expires. The cache is poisoned and subsequent requests for the domain address to recursive name servers are answered with the address of a different server, presumably controlled by the attacker. So long as the fake entry resides in the DNS server cache (persistence of a cache entry is usually governed by time to live) it can result in hundreds of thousands of dangerous redirects. In such cases the URL is legitimate, but the destination servers are not. This process is often called a “pharming” attack. Web servers and mail servers are frequent targets. Other redirection attacks include DNS Changer and DNS Replay. Man-in-the-middle is another descriptive term for many redirection attacks.

DNS Reflection and Amplification Attacks

As with UDP flood, DNS reflection attacks use a form of IP spoofing, changing the source address in their DNS queries to show the address of their intended target, such as a DNS root server or a top-level domain (TLD) name server operator. DNS reflection and amplification recognizes UDP as an asymmetrical protocol (small requests, large responses) and the existence of open DNS resolvers to the Internet cloud. The result is that small DNS queries reflect large UDP datagram responses to the target address in the original source datagrams. Some recent attacks have used this DDoS technique at a huge scale.

Because DNS runs over UDP and does not require a handshake, it’s possible to use the protocol as a means to lock down a host or a network. Designed a specific way, sending a small query to any open DNS resolver can result in a single response containing several kilobytes or more, that are sent to the unwitting spoofed victim. (This type of response typically is sent via TCP, as UDP does not allow for more than 512 bytes in a response datagram. The resulting packet usually exceeds the MTU of the recipient’s interfaces, resulting in further packet fragmentation and processing.) Open DNS resolvers may allow for launching DDoS attacks containing hundreds of gigabytes of data. Attackers may also use the EDNS0 DNS protocol extension as a means to enable larger DNS responses. Many network operators, particularly overseas, allow open DNS resolvers to run on their networks, unwittingly allowing attackers to abuse them. Many network operators do provide intelligent rate-limiting to prevent abuse, even while supporting open recursive DNS servers. Hence, issues of this type usually result from mistakes in configuration.

DNS Malware

Sophisticated malware also has emerged as a serious threat to DNS infrastructure. They are classified as Advanced Persistent Threats, and use DNS to embed themselves in the target network and stealthily communicate with external command servers to obtain malware updates, instructions, and to conduct attacks for data theft, industrial espionage and other goals.



Chapter 42 Infoblox DNS Firewall

This chapter provides information about the Infoblox DNS Firewall feature that you can configure and manage on the Infoblox appliance. It includes the following sections:

- [*About Infoblox DNS Firewall*](#) on page 1233
 - [*Setting Up Infoblox DNS Firewall*](#) on page 1234
- [*License Requirements and Admin Permissions*](#) on page 1235
 - [*For Local RPZs and RPZ Feeds*](#) on page 1235
 - [*For FireEye Integrated RPZs*](#) on page 1235
- [*Best Practices for Configuring RPZs*](#) on page 1236
 - [*General RPZ Best Practices*](#) on page 1236
 - [*Best Practices For FireEye Integrated RPZs*](#) on page 1236
- [*Enabling Recursion for RPZs*](#) on page 1237
 - [*Configuring RPZs for All Recursive Servers*](#) on page 1237
- [*Configuring Local RPZs*](#) on page 1238
- [*Configuring Rules for RPZs*](#) on page 1239
 - [*Managing Passthru Rules*](#) on page 1239
 - [*Managing Block \(No Such Domain\) Rules*](#) on page 1240
 - [*Managing Block \(No Data\) Rules*](#) on page 1241
 - [*Managing Substitute \(Domain Name\) Rules*](#) on page 1242
 - [*Managing Substitute \(Record\) Rules*](#) on page 1243
- [*Configuring RPZ Feeds*](#) on page 1248
- [*Configuring the Infoblox RPZ Feed*](#) on page 1249
 - [*Infoblox RPZ feeds*](#) on page 1251
- [*Downloading Rules of an RPZ Feed*](#) on page 1252
- [*Testing RPZ Feed Rules*](#) on page 1253
- [*About FireEye Integrated RPZs*](#) on page 1254
 - [*Configuring FireEye RPZs*](#) on page 1254
 - [*Configuring Rules for FireEye RPZs*](#) on page 1257
 - [*Configuring the FireEye appliance*](#) on page 1257
 - [*Handling Alerts from the FireEye appliance*](#) on page 1259
 - [*Logging FireEye Integrated RPZ messages*](#) on page 1259
 - [*Configuration Examples*](#) on page 1259

- [*Managing RPZs*](#) on page 1261
 - [*Viewing RPZs*](#) on page 1261
 - [*Modifying RPZs*](#) on page 1262
 - [*Reordering RPZs*](#) on page 1263
 - [*Locking and Unlocking RPZs*](#) on page 1263
- [*Managing RPZ Rules*](#) on page 1264
 - [*Viewing RPZ Rules*](#) on page 1264
 - [*Modifying RPZ Rules*](#) on page 1265
 - [*Deleting RPZ Rules*](#) on page 1265
 - [*Copying RPZ Rules*](#) on page 1265
 - [*Importing RPZ Rules*](#) on page 1266
- [*Verifying RPZ Configuration*](#) on page 1266
 - [*Viewing RPZ in the Syslog*](#) on page 1266
 - [*Viewing the Last Updated RPZs*](#) on page 1267

ABOUT INFOBLOX DNS FIREWALL

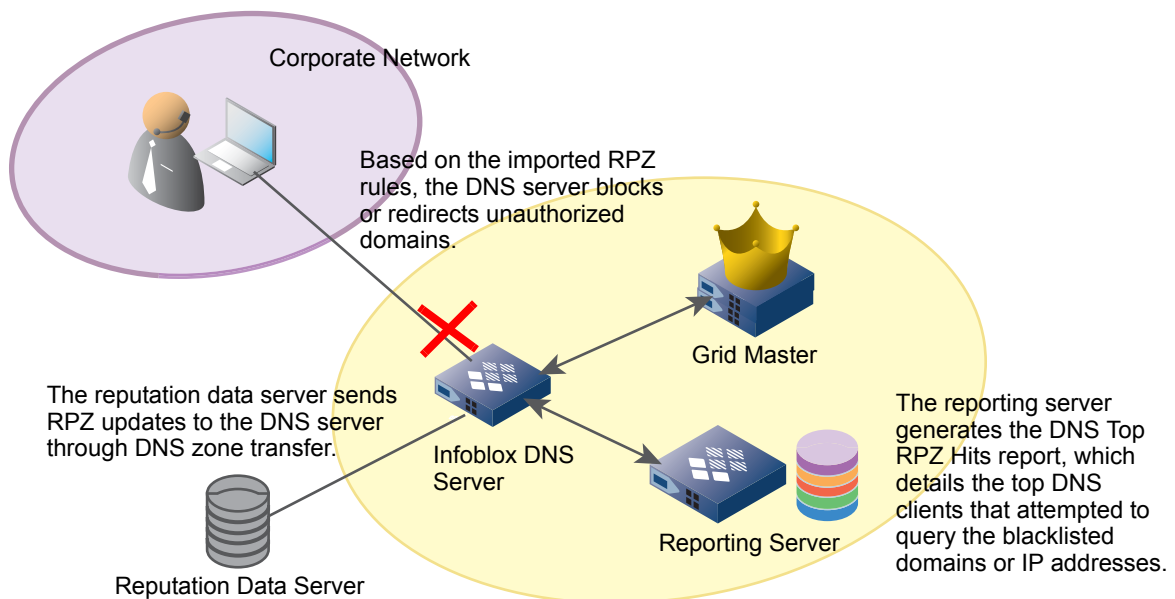
Infoblox DNS Firewall employs DNS RPZs (Response Policy Zones), a technology developed by ISC (Internet System Consortium) for allowing reputable sources to dynamically communicate domain name reputation so you can implement policy controls for DNS lookups.

On an Infoblox appliance, you can configure RPZs and define RPZ rules to block DNS resolution for malicious or unauthorized domain names, or redirect clients to a walled garden by substituting responses. You can assign actions to RPZ rules. For example, `abc.com` can have an action of pass thru or substitute (domain) with the domain `xyz.com`. You can also configure a Grid member to act as a lead secondary that receives RPZ updates from external reputation sources and redistributes the updates to other Grid members. Infoblox DNS firewall also facilitates the detection of malware and APTs (Advanced Persistent Threats) by integrating the NIOS appliance with a FireEye appliance. You can now employ APT mitigation strategy using FireEye as an external threat detection source.

An Infoblox Grid performs RPZ actions for queries that originate from external sources. The name server recursive cache on an RPZ enabled Grid member uses the address of the client from which the query originates to identify if the query is generated from an external source or an internal Grid. If the query originates from a Grid Master or a Grid member, RPZ actions are automatically bypassed for those queries. Infoblox uses an ACL, `infoblox-deny-rpz`, for RPZ, which contains a list of addresses to bypass RPZ actions if a query is sent from an internal Grid Master or member. Note that RPZ action is performed only once for a single recursion.

As illustrated in [Figure 42.1](#), the Infoblox DNS server receives RPZ updates, which include blacklisted domains and responses, from a reputation data server through a DNS zone transfer. The appliance then blocks or redirects queries and responses based on the imported RPZ rules. The reporting server can then generate the *DNS Top RPZ Hits* report that details the top DNS clients that have received redirected responses through RPZs.

Figure 42.1 Infoblox DNS Firewall



There are three types of RPZs:

- **Local RPZ**—A local RPZ is a zone that allows administrators to define multiple response policies locally. Responses sent are based on the defined rules. For information about how to configure local RPZs, see [Configuring Local RPZs](#) on page 1238.
- **RPZ Feed**—An RPZ feed receives response policies from external sources. DNS clients receive responses based on the imported rules from a reputable source, such as a commercial RPZ provider. For information about RPZ feed, see [Configuring RPZ Feeds](#) on page 1248 and [Configuring the Infoblox RPZ Feed](#) on page 1249.

- FireEye integrated RPZ—By integrating the NIOS appliance with the FireEye appliance, you can detect malware and APTs and take necessary actions to mitigate those threats. For information about FireEye integrated RPZ, see [About FireEye Integrated RPZs](#) on page 1254.

Setting Up Infoblox DNS Firewall

For a successful DNS firewall deployment to protect your endpoint devices and servers from stealthy malware and malicious domains, consider the guidelines described in [Best Practices for Configuring RPZs](#) on page 1236. To configure Infoblox DNS Firewall, complete the following tasks:

1. Install a valid RPZ license on the appliance, as described in [License Requirements and Admin Permissions](#) on page 1235. For more information about RPZ licenses, see [License Requirements and Admin Permissions](#) on page 1235.

Note: Ensure that you have installed a valid DNS license on the same appliance.

2. Enable recursive queries for a DNS view, member, or Grid, as described in [Enabling Recursion for RPZs](#) on page 1237.

Note: Ensure that you enable recursive queries for RPZ rules to take effect.

3. Configure RPZ logging to ensure that all matching and disabled rules for all queries are logged in the syslog. You can view the syslog to ensure that the rules are set up correctly before they take effect. Ensure that you enable **rpz** in the **Logging Category** of *Grid DNS Properties* editor to log these events. For information about how to set logging categories, see [Setting DNS Logging Categories](#) on page 1015.
4. You can configure a local RPZ, an RPZ feed, or a FireEye RPZ on the NIOS appliance. Complete one of the following depending on your selection:
 - On a DNS member, complete the following to create local RPZ rules:
 - a. Create an RPZ, as described in [Configuring Local RPZs](#) on page 1238.
 - b. Configure rules for the local RPZ you have created, as described in [Configuring Rules for RPZs](#) on page 1239.
 - Optionally, complete the following to receive RPZ updates from an RPZ feed:
 - a. Configure an RPZ feed, as described in [Configuring RPZ Feeds](#) on page 1248. You can also configure the Infoblox DNS feed, as described in [Configuring the Infoblox RPZ Feed](#) on page 1249. The Infoblox DNS feed is a reputable data server validated by Infoblox to provide reputation RPZ updates.
 - b. Download rules from the RPZ feed, as described in [Downloading Rules of an RPZ Feed](#) on page 1252.
 - Optionally, complete the following to receive alerts from a FireEye appliance:
 - a. Create a FireEye integrated RPZ, as described in [Configuring FireEye RPZs](#) on page 1254.
 - b. Define rules for FireEye RPZs, as described in [Configuring Rules for FireEye RPZs](#) on page 1257.
 - c. Create FireEye admin users, as described in [For FireEye Integrated RPZs](#) on page 1235.
 - d. Add URLs and user credentials on the FireEye appliance, as described in [Configuring the FireEye appliance](#) on page 1257.
5. Test your RPZ configuration and verify that RPZ is functioning properly by viewing the syslog and the **Last Updated** column in the **Response Policy Zones** tab. For more information, see [Testing RPZ Feed Rules](#) on page 1253.

After you have set up your RPZs, RPZ feeds, and RPZ rules, you can do the following:

- Manage local RPZs such as viewing a list of RPZs, modifying, reordering, and deleting RPZs. You can also lock or unlock RPZs. For more information, see [Managing RPZs](#) on page 1261.
- Verify RPZs are functioning properly by viewing the syslog and the last updated RPZ. For more information, see [Managing RPZ Rules](#) on page 1264.
- Manage Local RPZ rules such as viewing, modifying, and deleting RPZ rules. You can also copy and import RPZ rules. For more information, see [Managing RPZ Rules](#) on page 1264.

- Generate the *DNS Top RPZ Hits* report, if you have a reporting server set up in the Grid. For more information, see [DNS Top RPZ Hits](#) on page 1166.

LICENSE REQUIREMENTS AND ADMIN PERMISSIONS

You must install required licenses before you can use the RPZ feature. An RPZ license is required to configure local RPZs and RPZ feeds. For more information, see [For Local RPZs and RPZ Feeds](#). For FireEye integrated RPZs, you must first install an RPZ license, and then a FireEye license. For more information, see [For FireEye Integrated RPZs](#). For all RPZ related licenses, you can install either a temporary or a permanent license on the NIOS appliance. The temporary license provides a 60-day free trial, which can be upgraded to a permanent license. After the license expires, the RPZs will remain intact, but you cannot delete existing or add new entries to it. Infoblox provides RPZ licenses that are compatible with each product model.

For Local RPZs and RPZ Feeds

Before you install an RPZ license, ensure that the following are completed:

- The entire Grid is running NIOS 6.6 or later.
- Grid members are properly configured and DNS is enabled on the members.

Note: Install RPZ licenses only on Infoblox members that have DNS recursion enabled.

Superusers can configure RPZs and RPZ rules by default. You can also assign global permissions for all RPZs and RPZ rules to specific admin groups and roles. For more information, see [Administrative Permissions for Zones](#) on page 201.

For FireEye Integrated RPZs

You can enable FireEye integrated RPZs on the appliances that have both the RPZ and FireEye licenses installed. Note that you must install an RPZ license prior to installing the FireEye license. For more information about how to manage licenses, see [License Requirements and Admin Permissions](#) on page 1235.

NIOS appliance creates a new group, **fireeye-group**, when you add the first FireEye zone. The FireEye admin group is read-only and you cannot assign permissions to it. It will not have any superuser privileges and you cannot modify or delete this group. You can add users to the **fireeye-group** admin group, and FireEye users can only send alerts to the NIOS appliance. They cannot access the Infoblox GUI, CLI, API, or RESTful API. These users are authenticated based on the usernames and passwords you configure in the FireEye admin group. Only admin users who belong to the FireEye admin group can publish FireEye alerts. Other admin users cannot do so. For information about how to configure the FireEye appliance, see [Configuring the FireEye appliance](#) on page 1257.

Note: The **fireeye-group** is created automatically. Infoblox recommends that you do not add a group with the same name. In addition, The “force password change at next login” feature does not apply to admin users in the **fireeye-group**. These users will not be prompted to change their passwords at the next login. Their original passwords continue to work. For more information, see [Managing Passwords](#) on page 170.

To add users to the **fireeye-group**, complete the following:

1. From the **Administration** tab, select **Administrators**, and then click **Admins**.
2. Click **Add** and enter the usernames and passwords. For more information on how to add users to an admin group, see [Creating Local Admins](#) on page 169. Select **fireeye-group** for the admin group and add users to this group.

Note: Ensure that you save the usernames and passwords. You must use these credentials when configuring FireEye alerts to enable the alerts to be received by NIOS.

Uninstalling the FireEye License

When you uninstall the FireEye license, new FireEye alerts will not be processed. However, the FireEye integrated RPZs and the rules in those zones will not be deleted. Note the following when you uninstall the FireEye license:

- New FireEye alerts will not be processed
- FireEye RPZ zones that were created before uninstalling the license will remain
- You cannot create new FireEye RPZ zones
- RPZ rules created from the alert will remain
- Note that if the DNS Firewall and FireEye licenses are installed, then you must first remove the FireEye license to remove the DNS Firewall license.
- The fireeye-group and the FireEye zones will remain even after you delete the FireEye license.

BEST PRACTICES FOR CONFIGURING RPZs

Before configuring RPZs, observe the following best practices to ensure a successful configuration:

General RPZ Best Practices

- When you enable Infoblox DNS Firewall, DNS performance for all queries, recursive or authoritative, will be affected.
- For performance reasons, Infoblox recommends that you maintain a reasonable number of zones.
- Do not enable RPZ on multiple layers, such as on DNS client facing servers and forwarders.
- If you have multiple DNS servers in a Grid, ensure that you configure RPZs on the recursive server that is closest to your DNS clients. If you configure RPZs on second level DNS caching servers, you will not be able to identify the DNS clients because only the IP addresses of the forwarding name servers can be identified.
- Infoblox recommends that you preview your RPZ rules to ensure ruleset integrity and to avoid unexpected results. You can preview your rules by selecting **Log Only (Disabled)** when you configure **Policy Override** for an RPZ, RPZ feed, or FireEye integrated RPZ. For information about how to configure this, see [Configuring Local RPZs](#) on page 1238, [Configuring RPZ Feeds](#) on page 1248, and [About FireEye Integrated RPZs](#) on page 1254.
- The appliance logs all matching and disabled rules for all queries in the syslog. You can view the syslog to ensure that the rules are set up correctly before they take effect. Ensure that you enable **rpz** in the **Logging Category** of *Grid DNS Properties* editor to log these events. For information about how to set logging categories, see [Setting DNS Logging Categories](#) on page 1015.
- You can use the standard TSIG mechanism to ensure that feed zones come from the correct servers. Grid members can function either as a primary or secondary servers for the RPZ. As with hosting any zone as a secondary, please ensure that the appliance is sized properly to hold the zone contents in memory.
- You can only export or import the RPZ local zones using the CSV export or import feature, but you cannot import or export FireEye zones using this feature.
- Note that the NIOS blacklist and NXDOMAIN features take precedence over RPZs.

Best Practices For FireEye Integrated RPZs

Before you configure a FireEye integrated RPZ, consider the following:

- FireEye integrated RPZs inherit default values from local RPZs. You can create, edit and delete rules using the Infoblox GUI, API, and RESTful API.
- To avoid false positives, Infoblox recommends that you create a whitelist of allowed zones using a local RPZ that is sorted above the FireEye RPZ and add your own domain to the whitelist RPZ. For example, you can add your company domain name, such as corp100.com. This list must contain popular domains, such as Alexa 250, and other desired domains.

- Note that there will be an impact on the storage capacity when you create a new FireEye alert and map it with an RPZ rule. The processing of alerts will consume a few CPU cycles, which will have some impact on the system.
- You must properly configure the settings on a FireEye appliance. NIOS supports only **Per Event** delivery mechanism and **JSON Normal** message format. To ensure that the NIOS appliance process alerts properly, configure the FireEye appliance accordingly. For more information about alerts, see [Handling Alerts from the FireEye appliance](#) on page 1259.
- You cannot add a FireEye integrated RPZ during a scheduled full upgrade. However, updates to the CNAME record are processed during a full upgrade. NIOS updates CNAME records in the database to store information that is specific to FireEye alerts.
- The rules created due to insertion of alerts will be visible through the FireEye RPZ viewer. Infoblox recommends that you do not modify any internal objects. For more information, see [Viewing RPZs](#) on page 1261.
- Note that SSL certificate validation is not supported.
- You must verify the following after you configure the FireEye and NIOS appliances:
 - The URL configured on the FireEye appliance matches the URL in the FireEye integrated RPZ on NIOS.
 - Verify the username and password for FireEye admin on the FireEye appliance.
 - Ensure that the settings are properly configured on the FireEye appliance.
 - Verify the state of the FireEye appliance.

For more information about configuring the FireEye appliance to send alerts to the NIOS appliance, see [Configuring the FireEye appliance](#) on page 1257.
- Note that the frequency of alerts received from FireEye can be minimal. A very small number of alerts are generated on a weekly basis. For example, the FireEye appliance may generate only tens of alerts per day.

ENABLING RECURSION FOR RPZs

For RPZ rules to function properly, you must enable DNS recursion. You can enable DNS recursion at the Grid, member, or DNS view level. To enable recursion:

- For the Grid or member, see [Enabling Recursive Queries](#) on page 571.
- For a DNS view, see [Managing Recursive DNS Views](#) on page 609.

Configuring RPZs for All Recursive Servers

When you configure a local or FireEye integrated RPZ, you must define an internal primary name server. The primary name server can be either recursive or non-recursive, depending on its usage. When you configure an RPZ feed, you must define an external primary name server. You can associate a name server or a name server group with the local RPZ, RPZ feed, or FireEye RPZ. You can also configure RPZs and RPZ feeds for all recursive servers in the Grid.

A local RPZ can have one or more secondary name servers associated with it. For an RPZ feed, you must create an external primary name server.

To configure a local RPZ, or RPZ feed, or FireEye RPZ for all recursive servers, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. Enter the *Response Policy Zone* details and click **Next** to associate an RPZ with at least one name server. For information about creating a local RPZ, see [Configuring Local RPZs](#) on page 1238. For information about creating an RPZ feed, see [Configuring RPZ Feeds](#) on page 1248. For information about creating a FireEye integrated RPZ, see [Configuring FireEye RPZs](#) on page 1254.
3. Select **All Recursive Name Servers** from the list to add all the recursive name servers in the Grid as the secondary name servers for the corresponding zone.
4. Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information about extensible attributes, see [Using Extensible Attributes](#) on page 332.

CONFIGURING LOCAL RPZs

You can define local RPZs to match responses for recursive queries. Each RPZ can have various rules associated with it. The response of a recursive query is modified if it matches any of the RPZ rules. The responses are first matched with the RPZ rules, and if there is a match, the rule defined at the RPZ level override is used. You can create multiple local RPZs and define multiple rules for a local RPZ. Note that override depends on the order of the zones. The zones on top will override the zones below. You can change the order of the RPZs. For more information, see [Reordering RPZs](#) on page 1263. You can also configure FireEye integrated RPZs on the NIOS appliance to detect persistent threats and malwares. The NIOS appliance considers the FireEye integrated RPZ as a local RPZ. For more information, see [About FireEye Integrated RPZs](#) on page 1254.

Note: When using IDN (Internationalized Domain Name) in a local RPZ or RPZ feed, you must manually convert the IDN to punycode. For information about IDN, see [Support for Internationalized Domain Names](#) on page 93.

To configure local RPZs:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
 - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed if you have not created additional DNS views and only have the default view.
3. If you have configured multiple DNS views, you must drill-down to the corresponding view to assign a local RPZ. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your local RPZ, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information, see [Adding a DNS View](#) on page 605. For information on modifying an existing view, see [Modifying DNS Views](#) on page 611.
 - In the *Add Response Policy Zone Wizard*, select **Add Local Response Policy Zone**, click **Next** and specify the following:
 - **Name:** Enter the name of the local RPZ. It can be a combination of alphanumeric characters. You can enter up to 256 characters.
 - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the local RPZ.

Note: The local RPZ must have a primary Grid name server before you can configure it.

- **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.
 - **Log Only (Disabled)**—Select this if you want to disable an RPZ rewrite using rules in the RPZ. If the response to the recursive query matches any RPZ rule, then the rule is logged, but the response will not be altered. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence. Select this option to preview the rules in the syslog before they take effect.
 - **None (Given)**—Select this if you want to use the policy from the rule level.
 - **Block (No Data)**—Select this to send a response that contains no data in it.
 - **Block (No Such Domain)**—Select this if you want the user to receive a DNS response that indicates there is no domain. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
 - **Passthru**—Select this if you want to send an actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
 - **Substitute (Domain Name)**—Select this if you want to replace all the policy actions in an RPZ with the specified substitution action.
 - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
- **Comment:** Optionally, enter additional information about the local RPZ.

- **Disable:** Select the check box to disable a local RPZ without deleting its configuration. Clear the check box to enable the local RPZ. For information, see [Enabling and Disabling Zones](#) on page 621.
 - **Lock:** Select the check box to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking RPZs](#) on page 1263.
4. Click **Next** to associate the local RPZ with at least one primary name server:
 - Define the name servers for the local RPZ. A Grid name server must be recursive when primary Grid name server is used as an RPZ source. A local RPZ may or may not have a recursive server. For example, there could be a Grid that has only primary Grid name server for a local RPZ to act as an RPZ source for an external set of name servers. A local RPZ must have only one primary Grid name server and it can have one or more secondary Grid name servers. When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid are added as secondary servers for the zone. For information on specifying primary or secondary name servers, see [Assigning Zone Authority to Name Servers](#) on page 623. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629. For information about all recursive name servers, see [Configuring RPZs for All Recursive Servers](#) on page 1237.
 5. Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information, see [Using Extensible Attributes](#) on page 332.

Note: You cannot convert a local RPZ to an RPZ feed or vice versa.

CONFIGURING RULES FOR RPZs

You can define a list of rules based on which DNS server determines its response to recursive queries. Based on the rules defined, responses to clients are either altered or forwarded without any changes. Each rule consists of a domain name or IP address, specification or pattern, and an associated action.

These rules are applicable to local RPZs, including FireEye integrated RPZs. For RPZ feeds, rules are imported from external servers. You cannot change the content of an RPZ feed, but you can override the actions in an RPZ feed.

To configure RPZ rules:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click *DNS_View* -> *Zone* and then click **Add** -> select a **Rule**.
2. The rules are classified as follows:
 - **Passthru Rule:** For information, see [Managing Passthru Rules](#) on page 1239.
 - **Block (No Such Domain) Rule:** For information, see [Managing Block \(No Such Domain\) Rules](#) on page 1240.
 - **Block (No Data) Rule:** For information, see [Managing Block \(No Data\) Rules](#) on page 1241.
 - **Substitute (Domain Name) Rule:** For information, see [Managing Substitute \(Domain Name\) Rules](#) on page 1242.
 - **Substitute (Record) Rule:** For information, see [Managing Substitute \(Record\) Rules](#) on page 1243.
3. Complete the details in the corresponding editor.
4. Save the configuration and click **Next** to define extensible attributes. For information about extensible attributes, see [Using Extensible Attributes](#) on page 332.

You cannot define the above rules for an RPZ feed. An RPZ feed uses rules defined by external servers. When you click on an RPZ feed, the appliance displays a dialog box that provides various options to export the rules of the configured external servers in .CSV format.

Managing Passthru Rules

You can define passthru rules if you do not want to modify the actual responses of the recursive queries. The response received for a query is not modified, if there is a matching passthru rule and the actual response is forwarded to the user.

Adding Passthru Rules for Domain Names

To define passthru rules for domains:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View*-> *Zone*, and then click **Add** -> select **Passthru Rule** -> **Passthru Domain Name Rule**.
2. The following fields are displayed in the *Add a Passthru Domain Name Rule* wizard:
 - **Name:** Enter the domain name for which you want to define the passthru rule. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the passthru rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Adding Passthru Rules for IP Addresses or Networks

To define passthru rules for IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View*-> *Zone*, and then click **Add** -> select **Passthru Rule** -> **Passthru IP Address Rule**.
2. The following fields are displayed in the *Add a Passthru IP Address Rule* wizard:
 - **IP Address or Network:** Enter the IP address or specify the address in CIDR format for which you want to define the passthru rule. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the passthru rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Managing Block (No Such Domain) Rules

You can define rules to block certain domain names or IP addresses. When you choose this option to block a domain name, the query name is matched with the RPZ rule. If the query name matches the RPZ rule, the DNS client receives a DNS response that indicates the domain does not exist.

When you block an IP address or network using this option, the A and AAAA records are matched with the RPZ rule. If the records match an RPZ rule, the DNS client receives a DNS response that indicates the domain does not exist.

Defining Block (No Such Domain) Rules for Domain Names

To define block rules for domains:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View*-> *Zone*, and then click **Add** -> select **Block (No Such Domain) Rule** -> **Block Domain Name (No Such Domain) Rule**.
2. The following fields are displayed in the *Add a Block Domain Name (No Such Domain) Rule* wizard:
 - **Name:** Enter the domain name which you want to be blocked from being resolved by the DNS. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Comment:** Optionally, enter additional information.

- **Disable:** Clear the check box to enable the record. Select the check box to disable it.
- 3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
- 4. Save the configuration.

Defining Block (No Such Domain) Rules for IP Addresses or Networks

To define block rules for IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Block (No Such Domain) Rule** → **Block IP Address (No Such Domain) Rule**.
2. The following fields are displayed in the *Add a Block IP Address (No Such Domain) Rule* wizard:
 - **IP Address or Network:** Enter the IP address or specify the address in CIDR format which you want to block. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the block rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Managing Block (No Data) Rules

You can define rules to block certain domain names or IP addresses. When you choose this option to block a domain name, the query name is matched with the RPZ rule. If the query name matches the RPZ rule, the DNS client receives a DNS response that indicates there is no data for the requested record type.

When you block an IP address or network using this option, the A and AAAA records are matched with the RPZ rules. If the records match an RPZ rule, the DNS client receives a DNS response that indicates there is no data for the requested record type.

Defining Block (No Data) Rules for Domain Names

To define block rules for domains:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Block (No Data) Rule** → **Block Domain Name (No Data) Rule**.
2. The following fields are displayed in the *Add a Block Domain Name (No Data) Rule* wizard:
 - **Name:** Enter the domain name which you want to block. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the block rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Defining Block (No Data) Rules for IP Addresses or Networks

To define block rules for IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Block (No Data) Rule** → **Block IP address (No Data) Rule**.
2. The following fields are displayed in the *Add a Block IP Address (No Data) Rule* wizard:

- **IP Address or Network:** Enter the IP address or specify the address in CIDR format which you want to block. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the block rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Save the configuration.

Managing Substitute (Domain Name) Rules

You can define an alternate IP address or a domain name to redirect a domain name or an IP address, which is malicious or unauthorized. When the response to the client query matches an RPZ rule, the actual domain name or IP address is substituted with the alternative domain name or IP address. The client will receive the substituted value instead of the actual response.

Defining Substitute Domain Name (Based on Domain Name) Rules

To define substitutes for domain names:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Substitute (Domain Name) Rule** → **Substitute Domain Name (Based on Domain Name) Rule**.
2. The following fields are displayed in the *Add a Substitute (Domain Name) Rule* wizard:
 - **Name:** Enter the domain name for which you want to define a substitute. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Substituted Name:** Enter an alternative domain name or IP address that has to be substituted with the actual domain name. Click **Select Zone** to select a different zone.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the substitute rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Defining Substitute Domain Name (Based on IP address) Rules

To define substitutes for IP addresses:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Substitute (Domain Name) Rule** → **Substitute Domain Name (Based on IP Address) Rule**.
2. The following fields are displayed in the *Add a Substitute (Domain Name) Rule* wizard:
 - **IP address or Network:** Enter the IP address or network for which you want to define a substitute. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Substituted Name:** Enter an alternative domain name or IP address that has to be substituted with the actual IP address. Click **Select Zone** to select a different zone.
 - **Comment:** Optionally, enter additional information.
 - **Disable:** Clear the check box to enable the substitute rule. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

4. Save the configuration.

Managing Substitute (Record) Rules

You can define a substitute record for a domain name, which is considered malicious. You can define substitutes for the following in a zone:

- **A records:** For information about defining substitutes for A records, see [Defining Substitute Rules for A Records](#) on page 1243.
- **AAAA records:** For information about defining substitutes for AAAA records, see [Defining Substitute Rules for AAAA Records](#) on page 1243.
- **MX records:** For information about defining substitutes for MX records, see [Defining Substitute Rules for MX Records](#) on page 1244.
- **NAPTR records:** For information about defining substitutes for NAPTR records, see [Defining Substitute Rules for NAPTR Records](#) on page 1244.
- **PTR records:** For information about defining substitutes for PTR records, see [Defining Substitute Rules for PTR Records](#) on page 1245.
- **SRV records:** For information about defining substitutes for SRV records, see [Defining Substitute Rules for SRV Records](#) on page 1246.
- **TXT records:** For information about defining substitutes for TXT records, see [Defining Substitute Rules for TXT Records](#) on page 1247.
- **IPv4 address:** For information about defining substitutes for IPv4 addresses, see [Defining Substitute Rules for IPv4 Addresses or Networks](#) on page 1247.
- **IPv6 address:** For information about defining substitutes for IPv6 addresses, see [Defining Substitute Rules for IPv6 Addresses or Networks](#) on page 1247.

You can define a substitute for a certain owner name and record type. When you substitute a record for a certain owner name and record type, then responses to queries for that owner name and type are modified to contain the substituted value(s).

Defining Substitute Rules for A Records

An RPZ A (address) record maps a domain name to a substitute IPv4 address. To define a specific name-to-address mapping, add an A record to a previously defined RPZ.

To define substitute rules for A records:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Substitute (Record) Rule** → **Substitute (A Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (A Record) Rule* wizard:
 - **Name:** Enter the domain name that you want to map to an IP address. The name that you specify, irrespective of the RPZ name, is used to determine a match for the RPZ rule. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **IP Address:** Enter the IPv4 address to which you want the domain name to map.
 - **Comment:** Optionally, enter additional information about the A record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Defining Substitute Rules for AAAA Records

An RPZ AAAA (address) record maps a domain name to a substitute IPv6 address. To define a specific name-to-address mapping, add an RPZ AAAA record to a previously defined RPZ.

To define substitute rules for AAAA records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View*-> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (AAAA Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (AAAA Record) Rule* wizard:
 - **Name:** Enter the domain name that you want to map to an IP address. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **IP Address:** Enter the IPv6 address to which you want the domain name to map.
 - **Comment:** Optionally, enter additional information about the AAAA record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Defining Substitute Rules for MX Records

An RPZ MX (mail exchanger) record maps a domain name to a mail exchanger. A mail exchanger is a server that either delivers or forwards mail. A wildcard MX record applies to an RPZ and all its subdomains of the owner name.

To define substitute rules for MX records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View*-> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (MX Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (MX Record) Rule* wizard:
 - **Mail Destination:** Enter the owner name of the MX record you want to substitute.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Host Name Policy:** Displays the hostname policy of the selected zone. Ensure that the hostname you enter complies with the hostname restriction policy defined for the zone.
 - **Mail Exchanger:** Enter the fully qualified domain name of the mail exchanger.
 - **Preference:** Select an integer from 10 to 100. The preference determines the order in which a client attempts to contact the target mail exchanger.
 - **Comment:** Optionally, enter additional information about the MX record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Defining Substitute Rules for NAPTR Records

A DNS NAPTR object represents a Naming Authority Pointer (NAPTR) resource record. This resource record specifies a regular expression-based rewrite rule that, when applied to an existing string, produces a new RPZ name or URI.

To define substitute rules for NAPTR records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View*-> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (NAPTR Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (NAPTR Record) Rule* wizard:
 - **Domain:** Enter the domain name to which this resource record refers. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Service:** Select a service from the drop-down list. This field specifies the service and protocol that are used to communicate with the host at the domain name.

- **Flags:** The Flag field indicates whether the current lookup is terminal; that is, the current NAPTR record is the last NAPTR record for the lookup. It also provides information about the next step in the lookup process. The flags that are currently used are:
 - **U:** Indicates that the output maps to a URI (Uniform Record Identifier).
 - **S:** Indicates that the output is a domain name that has at least one SRV record. The DNS client must then send a query for the SRV record of the resulting domain name.
 - **A:** Indicates that the output is a domain name that has at least one A or AAAA record. The DNS client must then send a query for the A or AAAA record of the resulting domain name.
 - **P:** Indicates that the protocol specified in the Service field defines the next step or phase.
 - **Order:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. This value indicates the order in which the NAPTR records must be processed. It processes the record with the lowest value first.
 - **Preference:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. Similar to the Preference field in MX records, this value indicates which NAPTR record the DNS client should process first when the records have the same Order values. It processes the record with the lowest value first.
 - **REGEX:** The regular expression that is used to rewrite the original string from the client into a domain name. RFC 2915 specifies the syntax of the regular expression. Note that the appliance validates the regular expression syntax between the first and second delimiter against the Python re module, which is not 100% compatible with POSIX Extended Regular Expression as specified in the RFC. For information about the Python re module, refer to <http://docs.python.org/release/2.5.1/lib/module-re.html>.
 - **Replacement:** This specifies the domain name for the next lookup. The default is a dot (.), which indicates that the regular expression in the REGEX field provides the replacement value. Alternatively, you can enter the replacement value in FQDN format.
 - **Comment:** Optionally, enter additional information about the NAPTR record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Save the configuration.

Defining Substitute Rules for PTR Records

In a forward-mapping zone, a PTR (pointer) record maps a domain name to another domain name. In an RPZ, a PTR (pointer) record maps an address to a domain name. To define a specific address-to-name mapping, add an RPZ PTR record to a previously defined RPZ.

To define substitute rules for PTR records:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, select *DNS_View* → *Zone*, and then click **Add** → select **Substitute (Record) Rule** → **Substitute (PTR Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (PTR Record) Rule* wizard:

You can select either **Name** or **IP address** from the drop-down list.

 - **Name:** Enter a domain name for which you want to create a pointer to another domain. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone. The name should be in the following format for RPZ:


```
ipaddress.in-addr.arpa.
```

Note that the IP address should be in the reverse format. For example, if the IP address is 10.2.1.4, then the name format for RPZ is `4.1.2.10.in-addr.arpa`. The following fields are displayed when you select **Name** from the drop-down list:

 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Domain Name:** Enter the domain name to which you want the PTR record to point.
 - **IP Address:** Enter an IP address for which you want to create a pointer to a domain. The following fields are displayed when you select **IP Address** from the drop-down list:
 - **Zone:** Displays the RPZ you have selected. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.

- **Domain Name:** Enter the domain name to which you want the PTR record to point.
 - **Comment:** Optionally, enter additional information about the PTR record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Save the configuration.

Defining Substitute Rules for SRV Records

A DNS RPZ SRV object represents an SRV resource record, which is also known as a service record. You can define a substitute for an SRV record. When the response to a user's query matches with an RPZ rule, then the combination of actual service, protocol, domain name and the zone is substituted with a combination of priority, weight, port and target details that you specify.

To define substitute rules for SRV records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (SRV Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (SRV Record) Rule* wizard:
 - **Display input as:** Select the format in which you want the SRV record to be displayed. When you select **RFC 2782 format**, the appliance follows the *_service._protocol.name* format as defined in RFC 2782. When you select **Free format**, enter the entire name in the **Domain** field.
 - **Service:** Specify the service that the host provides. You can either select a service from the list or type in a service, if it is not on the list. For example, if you are creating a record for a host that provides FTP service, select **_ftp**. To distinguish the service name labels from the domain name, the service name is prefixed with an underscore. If the name of the service is defined in RFC 1700, Assigned Numbers, use that name. Otherwise, you can use a locally-defined name. This field is disabled when you select **Free Format** as the display input.
 - **Protocol:** Specify the protocol that the host uses. You can either select a protocol from the list or type in a protocol, if it is not on the list. For example, if it uses TCP, select **_tcp**. To distinguish the protocol name labels from the domain name, the protocol name is prefixed with an underscore. This field is disabled when you select **Free Format** as the display input.
 - **Domain:** If Grid Manager displays a zone name, enter the name here to define an SRV record for a host or subdomain. The displayed zone name can either be the last selected zone or the zone from which you are adding the SRV record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name to define the SRV record. The SRV record name is used to determine the substitute.
 - **Preview:** After you enter all the information, this field displays the FQDN.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Priority:** Select or enter an integer from 0 to 65535. The priority determines the order in which a client attempts to contact the target host; the domain name host with the lowest number has the highest priority and is queried first. Target hosts with the same priority are attempted in the order defined in the **Weight** field.
 - **Weight:** Select or enter an integer from 0 to 65535. The weight allows you to distribute the load between target hosts. The higher the number, the more that host handles the load (compared to other target hosts). Larger weights give a target host a proportionately higher probability of being selected.
 - **Port:** Specify the appropriate port number for the service running on the target host. You can use standard or nonstandard port numbers, depending on the requirements of your network. You can select a port number from the list or enter an integer from 0 to 65535.
 - **Target:** Enter the canonical domain name of the host (not an alias); for example, *www2.corp100.com*.
 - **Comment:** Optionally, enter additional information about the SRV record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.

4. Save the configuration.

Defining Substitute Rules for TXT Records

A TXT (text) record contains supplemental information for a host. SPF (Sender Policy Framework) records are specialized RPZ TXT records that identify the servers that send mail from a domain.

To define substitute rules for TXT records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (TXT Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (TXT Record) Rule* wizard:
 - **Name:** Enter the name to define a TXT record for a host or subdomain. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Text:** Enter the text that you want to associate with the record. It can contain substrings of up to 255 bytes, up to a total of 512 bytes. Additionally, if you enter leading, trailing, or embedded spaces in the text, add quotes around the text to preserve the spaces. For example: " v=spf1 include:corp200.com -all ".
 - **Comment:** Optionally, enter additional information about the TXT record.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

Defining Substitute Rules for IPv4 Addresses or Networks

You can define a substitute for an IPv4 address or a network address. When a client queries for A records of a domain name, if the IP address in A records in the response match the specified address or network, then the response is modified to instead contain the substituted address.

To define substitute rules for IPv4 addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (IPv4 Address) Rule**.
2. The following fields are displayed in the *Add a Substitute (IPv4 Address) Rule* wizard:
 - **IP Address or Network:** Enter the IPv4 address which you want to substitute with another IPv4 address. Click **Select Zone** to select a different zone.

Note: You cannot define a substitute rule for the same IP address or a network address for which you have already defined a passthru rule.

- **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Substituted IP Address:** Enter the IPv4 address that must be returned to the user when the response matches the A records.
 - **Comment:** Optionally, enter additional information about the IPv4 address.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 4. Save the configuration.

Defining Substitute Rules for IPv6 Addresses or Networks

You can restrict access to specific IPv6 addresses or networks by providing a substitute IP address. When a client queries for AAAA records of a domain name if the IP addresses in AAAA records in the response match the specified address or network, then the response is modified to instead contain the substituted address.

To define substitute rules for IPv6 addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (IPv6 Address) Rule**.
2. The following fields are displayed in the *Add a Substitute (IPv6 Address) Rule* wizard:
 - **IP Address or Network:** Enter the IPv6 address or the network address which you want to substitute with another IP address. Click **Select Zone** to select a different zone.
 - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
 - **Policy:** Displays the selected policy.
 - **Substituted IP Address:** Enter the IPv6 address that must be returned to the user when the response matches the AAAA records.
 - **Comment:** Optionally, enter additional information about the IPv6 address.
 - **Disable:** Clear the check box to enable the record. Select the check box to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
4. Save the configuration.

CONFIGURING RPZ FEEDS

An RPZ feed receives response policies from an external source that produces reputation RPZ data and transfers the data to Grid name servers through zone transfers with or without a TSIG key. To ensure proper authentication and integrity of the RPZ feed zone transfers, using a TSIG key is recommended.

Note: To enter IDNS (Internationalized Domain Name) in an RPZ feed, you can use the punycode representation of the IDN.

You can use third party RPZ sources or create your own RPZ sources. You can configure an RPZ feed with multiple RPZ sources.

To configure RPZ feeds:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
 - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed, if you have not created additional *DNS views* and only have the *default view*.
 - If you have configured multiple DNS views, you must drill-down to the corresponding *DNS_View* to assign an RPZ feed. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your RPZ feed, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information, see [Adding a DNS View](#) on page 605. For information on modifying an existing view, see [Modifying DNS Views](#) on page 611.
3. In the *Add Response Policy Zone Wizard*, select **Add Response Policy Zone Feed**, click **Next** and specify the following:
 - **Name:** Enter the name of the RPZ feed. It can be a combination of alphanumeric characters. You can enter up to 256 characters.
 - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the RPZ feed.
 - **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.

- **Log Only (Disabled)**—Select this if you want to disable an RPZ rewrite using rules in the RPZ zone. If the response to the recursive query matches any RPZ rule, the rule is logged, but the response will not be altered. You cannot overwrite the response to the user. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence.
 - **None (Given)**—Select this if you want to use the policy from the rule level.
 - **Block (No Data)**—Select this if you want the user to receive a response, which indicates that there is no data.
 - **Block (No Such Domain)**—Select this if you want the user to receive a NXDOMAIN in the response. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
 - **Passthru**—Select this if you want the user to see the actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
 - **Substitute (Domain Name)**—Select this if you want to replace all the policy actions in an RPZ with the substitution action that is specified.
 - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
 - **Comment:** Optionally, enter additional information about the RPZ feed.
 - **Disable:** Select the check box to disable an RPZ feed without deleting its configuration. Clear the check box to enable the RPZ feed. For information, see [Enabling and Disabling Zones](#) on page 621.
 - **Lock:** Select the check box to lock the RPZ feed so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking RPZs](#) on page 1263.
4. Click **Next** to associate the RPZ feed with at least one external primary name server and a secondary name server:
- Define name servers for the RPZ feed. Note that either the Grid name server or the DNS view must be recursive for RPZ feed. An RPZ feed must have at least one RPZ source as an external primary name server and at least one Grid secondary name server. You can associate a lead secondary with an RPZ feed. Note that a lead secondary may or may not have recursion enabled when it is used only for an RPZ feed. Infoblox appliance includes an ACL configured for RPZ actions, which defines a list of addresses to bypass RPZ actions for any server in the Grid. For information on specifying primary and secondary, see [Assigning Zone Authority to Name Servers](#) on page 623. When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid with RPZ licenses are added as secondary servers for the zone. For more information, see [Configuring RPZs for All Recursive Servers](#) on page 1237. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629.

Note: The RPZ feed must have an external primary name server before you can configure it.

5. Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information, see [Using Extensible Attributes](#) on page 332.

CONFIGURING THE INFOBLOX RPZ FEED

Infoblox has validated a reputable source for RPZ updates. You can configure this RPZ feed and receive reputation RPZ updates on a regular basis. Infoblox offers subscription services for RPZ updates. Contact your sales representative for pricing and availability information.

To propagate RPZs as quickly as possible, the secondary DNS server needs an address to which the RPZ source feed can send NOTIFY messages. For example, if the secondary DNS server is configured behind a NAT, you may want to establish a one-to-one NAT for the lead secondary DNS server so it can receive NOTIFY messages from the RPZ source feed. Otherwise, the lead secondary DNS server will need to periodically poll the RPZ source feed, which could take longer than expected.

Note: The RPZ feed must have an external primary name server before you can configure it.

To configure the Infoblox RPZ feed:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
 - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed, if you have not created additional *DNS views* and only have the *default view*.
 - If you have configured multiple DNS views, you must drill-down to the corresponding *DNS_View* to assign an RPZ feed. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your RPZ feed, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information, see [Adding a DNS View](#) on page 605. For information on modifying an existing view, see [Modifying DNS Views](#) on page 611.
3. In the *Add Response Policy Zone Wizard*, select **Add Response Policy Zone Feed**, click **Next** and specify the following:
 - **Name:** Enter the name of the Infoblox RPZ feed.
Infoblox offers seven feeds. Three of the feeds consist of malicious devices/domains. The other four consist of the largest malicious feed in combination with various geographic data. The malicious feeds are listed in the order of size, and the smaller feeds are subsets of the larger ones. For more information, see [Infoblox RPZ feeds](#) on page 1251. You should configure only one of the seven available feeds.
 - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the RPZ feed.
 - **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.
 - **Log Only (Disabled)**—Select this if you want to disable an RPZ rewrite using rules in the RPZ zone. If the response to the recursive query matches any RPZ rule, the rule is logged, but the response will not be altered. You cannot overwrite the response to the user. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence.
 - **None (Given)**—Select this if you want to use the policy from the rule level.
 - **Block (No Data)**—Select this if you want the user to receive a response that indicates that there is no data.
 - **Block (No Such Domain)**—Select this if you want the user to receive a NXDOMAIN as the DNS response. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
 - **Passthru**—Select this if you want the user to see the actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
 - **Substitute (Domain Name)**—Select this if you want to replace all the policy actions in an RPZ with the substitution action that is specified.
 - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
 - **Comment:** Optionally, enter additional information about the Infoblox RPZ feed.
 - **Disable:** Select the check box to disable the RPZ feed without deleting its configuration. Clear the check box to enable the RPZ feed. For information, see [Enabling and Disabling Zones](#) on page 621.
 - **Lock:** Select the check box to lock the RPZ feed so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking RPZs](#) on page 1263.
4. Click **Next** to associate the RPZ feed with at least one external primary name server and a secondary name server:
 - Define name servers for the RPZ feed. An RPZ feed must have at least one RPZ source as an external primary name server and at least one Grid secondary name server. For external primary servers, specify the following:
 - **Name:** Enter the zone name of the primary name server.
 - **Address:** Enter the name server IP address provided by Infoblox for the RPZ feed.
 - **Use TSIG:** Select the check box to specify TSIG settings.

- **Key Name:** Enter the TSIG Key Name provided by Infoblox.
- **Key Algorithm:** Select `hmac-md5`.
- **Key Data:** Enter the TSIG string provided by Infoblox.
 Note that either the Grid name server or the DNS view must be recursive for the RPZ feed. You can associate a lead secondary with an RPZ feed. For information on specifying primary and secondary, see [Assigning Zone Authority to Name Servers](#) on page 623. When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid are added as secondary servers for the zone. For information about all recursive name servers, see [Configuring RPZs for All Recursive Servers](#) on page 1237. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629.

Note: The RPZ feed must have an external primary name server before you can configure it.

5. Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information, see [Using Extensible Attributes](#) on page 332.

Infoblox RPZ feeds

Infoblox RPZ feeds are categorized into pure malicious feeds and combination feeds. All the feeds listed below are set to return NXDOMAIN for items in the feed. Threat data changes are pushed every 20 minutes from the DNS servers and significant changes are typically made every two hours. The following tables list the Infoblox RPZ feeds:

Table 42.1 Pure Malicious Feeds

Name	Description
cnc.rpz.infoblox.local	Contains known botnet C&C domains/IPs and dropboxes as well as name servers that are known to be used solely by malicious entities. In addition to active botnets, this also includes resources used to sinkhole contact attempts by botnets that have been taken down by law enforcement and/or security researchers (e.g. conficker).
cnc.rpz.cn.infoblox.local	This data feed is available for servers located in Mainland China. Contains known botnet C&C domains/IPs and dropboxes as well as name servers that are known to be used solely by malicious entities. In addition to active botnets, this also includes resources used to sinkhole contact attempts by botnets that have been taken down by law enforcement and/or security researchers (e.g. conficker).
cnc-driveby.rpz.infoblox.local	In addition to the contents in the “cnc.rpz.infoblox.local” feed, this also includes the sites (IPs/domains/name servers) for known malware dropper sites and other places that can infect a computer that visits it. It includes networks and autonomous systems that are on the “Do not Route Or Peer” (DROP) list.
cnc-driveby.rpz.cn.infoblox.local	This data feed is available for servers located in Mainland China. In addition to the contents in the “cnc.rpz.cn.infoblox.local” feed, this also includes the sites (IPs/domains/name servers) for known malware dropper sites and other places that can infect a computer that visits it. It includes networks and autonomous systems that are on the “Do not Route Or Peer” (DROP) list.
malware.rpz.infoblox.local	A comprehensive list of malware hosts/domains/name servers. In addition to the contents in the “cnc-driveby.rpz.infoblox.local” feed, this includes known active phishing sites and other threats.

Name	Description
malware.rpz.cn.infoblox.local	This data feed is available for servers located in Mainland China. A comprehensive list of malware hosts/domains/name servers. In addition to the contents in the “cnc-driveby.rpz.cn.infoblox.local” feed, this contains known active phishing sites and other threats.

Table 42.2 Combination Feeds

Name	Description
malware-prc.rpz.infoblox.local	Contains the malware data feed as well as the IP subnets, the ccTLD domain and name servers for the People’s Republic of China.
malware-ee.rpz.infoblox.local	Contains the malware data feed as well as the IP subnets, the ccTLD domains and name servers for countries in Eastern Europe that are major hosts of malware: Russia, Ukraine, Latvia, Moldova, Romania.
malware-ee.rpz.cn.infoblox.local	This data feed is available for servers located in Mainland China. Contains the malware data feed as well as the IP subnets, the ccTLD domains and name servers for countries in Eastern Europe that are major hosts of malware: Russia, Ukraine, Latvia, Moldova, Romania.
malware-ee-prc.rpz.infoblox.local	Contains the malware data feed as well as the IP subnets, the ccTLD domains and name servers for the People’s Republic of China and countries in Eastern Europe that are major hosts of malware: Russia, Ukraine, Latvia, Moldova, Romania.
malware-sanction.rpz.infoblox.local	Contains the malware data feed as well as the IP subnets, the ccTLD domains and name servers for the countries on the OFAC Embargo and ITAR lists maintained by the US government. Currently the countries included are: Afghanistan, Belarus, Burma (Myanmar), China, Cote d'Ivoire, Cuba, Cyprus, Congo (Dem Rep), Eritrea, Haiti, Iran, Iraq, Lebanon, Liberia, Libya, North Korea, Sierra Leone, Somalia, Sri Lanka, Sudan, Syria, Venezuela, Vietnam, Yemen, Zimbabwe. See http://www.treasury.gov/about/organizational-structure/offices/Pages/Office-of-Foreign-Assets-Control.aspx and http://pmddtc.state.gov/regulations_laws/itar_official.html

DOWNLOADING RULES OF AN RPZ FEED

You can perform a zone transfer to transfer the rules from an external primary name server to the RPZ feed. You cannot modify these rules, but you can override the entire ruleset or an individual rule. However, if you import a zone to a local zone, you can edit the rules within a local zone. The feed zone supports NSIP and NSDNAME rules; however local RPZs do not support these rules. To download rules from an external primary name server:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, and then click the corresponding *RPZ Feed*.
2. In the **Export** dialog box, complete the following:
 - **Separator:** Select the separator used in the data file. The default value is **Comma**.
 - Click **Export**.

All the rules are transferred. You can download rules only if the lead secondary has completed at least one zone transfer from the external primary. You can either open the data file or save it to your computer. The rules are displayed for the selected RPZ feed in the *Rule*wizard.

After you have downloaded rules from an RPZ feed, you can test RPZ feed policies, as described in [Testing RPZ Feed Rules](#) on page 1253.

TESTING RPZ FEED RULES

After you have downloaded rules from an RPZ feed, you can test the downloaded policies by using the `dig` command and observing log messages that contain redirect or rewrite responses in the syslog. The NIOS appliance supports generation of RPZ log messages in CEF (Common Event Format). Note that non-RPZ messages cannot be generated in CEF.

You must enable the **rpz** option in the **Logging Category** of the *Grid DNS Properties* editor to receive RPZ related messages in the syslog. For information about configuring the logging category, see [Setting DNS Logging Categories](#) on page 1015.

To view RPZ log messages in the syslog, you can use the system filter **RPZ Logs** from the **Quick Filter** to filter the messages. Note that only messages in CEF are displayed.

To view RPZ log messages:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right corner, select the Grid member on which you want to view the syslog.
3. Click **Show Filters** to enable the filters. Select **RPZ Logs** from the **Quick Filter** drop-down list to narrow down the system messages you want to view.

The name server recursive cache makes a syslog entry when an RPZ functionality fails. The syslog message log format is as follows:

```
rpz <TYPE> rewrite <QUERY> via <RPZ_RECORD><ERROR_MESSAGE>
  where: <TYPE> is one of following RPZ action types: QNAME, IP, NSIP, NSDNAME;
  <QUERY> is a query record to process;
  <RPZ_RECORD> is an RPZ record that is used to perform an action to the query;
  <ERROR_MESSAGE> is a message with error details. Example: NS address rewrite rrset failed:,
  concatenate() failed:, NS db_find() failed:, stop on qresult in rpz_rewrite() failed:,
  stop on unrecognized qresult in rpz_rewrite() failed:, etc.
```

To test RPZ feed policies:

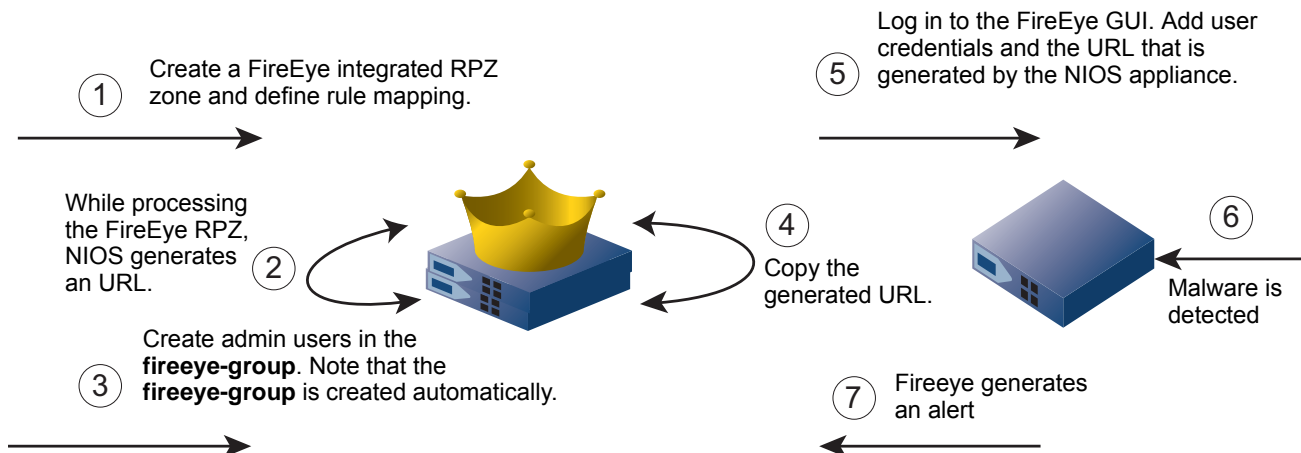
1. Open a terminal console on your computer.
2. Type the command `dig @<your DNS server IP> <queried domain>`.
3. Go to the **Administration** tab -> **Logs** tab -> **Syslog** tab to view CEF log messages.

ABOUT FIREEYE INTEGRATED RPZS

Infoblox DNS firewall provides a mechanism to further protect your network from malware and APTs (Advanced Persistent Threats) through the integration of FireEye appliances. When your NIOS appliance is properly integrated with a FireEye appliance, it receives periodic alerts and APTs from the FireEye appliance when it identifies such threats. Based on your configuration, the NIOS appliance translates these alerts into RPZ rules that not only further protect your network from malicious attacks, but also aid in identifying clients that have been compromised.

As illustrated in *Figure 21.2*, after installing the required RPZ and FireEye licenses on the NIOS appliance, you can configure a FireEye integrated RPZ in which you map RPZ rules to FireEye alert types. While creating the FireEye RPZ, the appliance generates a URL to which the FireEye appliance sends alerts. Ensure that you enter this URL when configuring the FireEye appliance. The NIOS appliance also creates the **fireeye-group** admin group after you define the first FireEye RPZ. You can add multiple admin users to this admin group. Note that users in the **fireeye-group** can only send alerts to the NIOS appliance; they cannot access the Infoblox GUI, CLI, API and RESTful API. They also do not have permissions to perform other tasks on the appliance. Ensure that you record the usernames and passwords for all user accounts so you can enter them correctly when you configure the FireEye appliance. You can map a single or multiple FireEye appliances to a NIOS appliance where multiple users or zones exist.

Figure 42.2 FireEye Integrated RPZ



To configure a FireEye integrated RPZ, complete the following:

1. Create a new FireEye integrated RPZ, as described in [Configuring FireEye RPZs](#) on page 1254.
2. Create FireEye admin users, as described in [For FireEye Integrated RPZs](#) on page 1235.
3. Add URL and user credentials on the FireEye appliance, as described in [Configuring the FireEye appliance](#) on page 1257.
4. When a malware or threat is detected, the FireEye appliance sends an alert message to the NIOS appliance, which is stored in the syslog. For more information, see [Handling Alerts from the FireEye appliance](#) on page 1259.

Configuring FireEye RPZs

You must create an RPZ zone and map the FireEye alerts with an RPZ rule to receive alerts from FireEye. These alerts will then be translated into appropriate RPZ rules that are added to the FireEye RPZ. You can also define a time limit for a specific alert type or set the alert type to live forever. When you define a lifetime, the alert type will be active for the specified number of days or weeks in the NIOS appliance, and will then expire after the specified time. After you configure the FireEye integrated RPZ, the NIOS Grid receives alerts from the FireEye appliance and creates RPZ rules for some of the alerts received. FireEye appliance sends alert messages with basic authentication. You must configure a username and password on the NIOS appliance prior to receiving any alerts from the FireEye appliance.

Note: The NIOS appliance treats the FireEye integrated RPZ as a local RPZ. Thus, you cannot assign an external primary name server to the zone.

An alert contains the malware URL along with a valid FQDN. The NIOS appliance can only map an alert to a RPZ rule if the FQDN is present. If an alert doesn't contain the FQDN, then the alert is ignored by the NIOS appliance. For more information about alerts, see [Handling Alerts from the FireEye appliance](#) on page 1259. Once the alert is processed and properly mapped to an RPZ rule, it remains in the database until you delete it manually. You can get more information about the alerts, which are sent by the FireEye appliance, from the syslog.

Note: You can configure feeds from multiple FireEye appliances. To enable or disable FireEye integration module feeds from individual appliances, you must enable or disable user access of the particular FireEye appliance. Note that the FireEye feeds will not be in the RPZ format, but when you configure a FireEye integrated RPZ, the NIOS appliance creates a new URL through which the FireEye appliance sends alerts.

To configure a FireEye integrated RPZ:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
 - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed if you have not created additional DNS views and only have the default view.
3. If you have configured multiple DNS views, you must drill-down to the corresponding view to assign a FireEye Integrated RPZ. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your FireEye integrated RPZ, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information, see [Adding a DNS View](#) on page 605. For information on modifying an existing view, see [Modifying DNS Views](#) on page 611.
4. In the *Add Response Policy Zone Wizard*, select **Add FireEye-Integrated Response Policy Zone**, click **Next** and specify the following:
 - **Name:** Enter the name of the FireEye integrated RPZ. It can be a combination of alphanumeric characters. You can enter up to 256 characters.
 - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the FireEye integrated RPZ.
 - **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.
 - **Log Only (Disabled)**—Select this if you want to disable an RPZ rewrite using rules in the RPZ. If the response to the recursive query matches any RPZ rule, then the rule is logged, but the response will not be altered. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence. Select this option to preview the rules in the syslog before they take effect.
 - **None (Given)**—Select this if you want to use the policy from the rule level.
 - **Block (No Data)**—Select this to send a response that contains no data in it.
 - **Block (No Such Domain)**—Select this if you want the user to receive a DNS response that indicates there is no domain. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
 - **Passthru**—Select this if you want to send an actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
 - **Substitute (Domain Name)**—Select this if you want to replace all the policy actions in an FireEye integrated RPZ with the specified substitution action.
 - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
 - **Comment:** Optionally, enter additional information about the FireEye integrated RPZ.

- **Disable:** Select the check box to disable the FireEye integrated RPZ without deleting its configuration. Clear the check box to enable the FireEye integrated RPZ. For information, see [Enabling and Disabling Zones](#) on page 621.
5. Click **Next** to define rule mapping:
- **Server URL:** The appliance displays the URL that you use when configuring the FireEye appliance. This URL is used to handle alerts, which is sent by the FireEye appliance. It handles alerts based on the standard authentication. The URL generated by the NIOS appliance consists of the Grid Manager IP address, network view, and DNS view of the FireEye zone. If you change the IP address, network view, zone or DNS view after you have configured a FireEye RPZ, the URL will change accordingly. Thus FireEye will not be able to send alerts to the updated URL. You must update the URL in the FireEye appliance to send alerts to the NIOS appliance. The **Server URL** is generated in this format:

```
https://<host address>/alert/feye/<network view>/<dns view>/<zone>
```
 - **Rule Mapping:** You can map a **FireEye alert type** with an RPZ policy. Select an **RPZ policy type** from the drop-down list. Note that the **FireEye alert type** is read-only. The NIOS appliance applies corresponding RPZ policy type when the FireEye appliance sends an alert to the NIOS appliance. You can also specify a time limit for each FireEye RPZ rule depending on the FireEye alert type. NIOS displays default lifetime value for each alert type. You can change the default lifetime of the alert type. When you define a value, the value must be greater than zero. When you select **Live Forever** from the drop-down list, the alert type will never expire and will be stored in the database until further notice. The NIOS appliance will use the default time if you do not specify a value. You can specify the expiration time in days or weeks only. The following table lists the FireEye alerts, RPZ policy types, and the time limit for a specific FireEye alert:

Table 42.3 FireEye Rule Mapping

FireEye Alert Type	RPZ Policy Type	Lifetime
Domain Match	Select a value from the drop-down list for a FireEye alert when a malware object is detected: None , Passthru , Block (No Such Domain) , Block (No Data) , and Substitute (Domain Name) . The drop-down list displays Passthru , by default. For more information about the RPZ Policy Types , see Configuring Rules for RPZs on page 1239.	Specify a lifetime for each FireEye alert in Days , Weeks , or select Live Forever from the drop-down list. The following are the default values for different alert types: <ul style="list-style-type: none"> Domain Match - 1 week Infection Events - 1 day Callback Events - 1 week Malware Object - 1 day Web Infection - 1 day Click on the default value to change the lifetime value.
Infection Events		
Callback Events		
Malware Object		
Web Infection		

When you edit the lifetime of an existing alert type, NIOS deletes the alert type based on the new lifetime setting. It also updates the expiration time for the corresponding alert type. Note that there might be an impact on the performance when you delete expired FireEye RPZ rules.

- **Override rule mapping for APT events:** Select a value from the drop-down list to override rule mapping for Advanced Persistent Threats. Events that are marked as APT events by FireEye override rules that are set for other event types. The values in the drop-down list are:
 - **No Override**—Select this if you want to use the policy from the rule level and do not want to override the rule mapping settings. This value is displayed in the drop-down list, by default.
 - **Passthru**—Select this if you want the user to see the actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
 - **Block (No Such Domain)**—Select this if you want the user to receive a NXDOMAIN as the DNS response. All the policy actions in an RPZ are replaced with a NXDOMAIN block.

- **Block (No Data)**—Select this if you want the user to receive a response that indicates that there is no data.
 - **Substitute (Domain Name)**—Select this if you want to replace all the policy actions in an RPZ with the substitution action that is specified.
 - **Substituted Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list either for APT events or for FireEye alerts. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
6. Click **Next** to associate the FireEye integrated RPZ with at least one primary name server:
 - Define the name servers for the FireEye integrated RPZ. A Grid name server must be recursive when primary Grid name server is used as an RPZ source. A FireEye integrated RPZ may or may not have a recursive server. For example, there could be a Grid that has only primary Grid name server for a FireEye integrated RPZ to act as an RPZ source for an external set of name servers. A FireEye integrated RPZ must have only one primary Grid name server and it can have one or more secondary Grid name servers. When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid are added as secondary servers for the zone. For information on specifying primary or secondary name servers, see [Assigning Zone Authority to Name Servers](#) on page 623. For information on specifying name server groups, see [Using Name Server Groups](#) on page 629. For information about all recursive name servers, see [Configuring RPZs for All Recursive Servers](#) on page 1237.
 7. Save the configuration and click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#) on page 332.
 8. Click **Restart** if it appears at the top of the screen.

Configuring Rules for FireEye RPZs

You can define a list of rules based on how the DNS server determines its response to recursive queries. Based on the rules defined, responses to clients are either manipulated or forwarded without any changes. To configure rules for FireEye RPZs:

1. From the **Data Management** tab, select the **DNS** tab → **Response Policy Zones** tab, click *DNS_View* → *Zone* and then click **Add** → select a **Rule**.
2. The rules are classified as follows:
 - **Passthru Rule:** For information, see [Managing Passthru Rules](#) on page 1239.
 - **Block (No Such Domain) Rule:** For information, see [Managing Block \(No Such Domain\) Rules](#) on page 1240.
 - **Block (No Data) Rule:** For information, see [Managing Block \(No Data\) Rules](#) on page 1241.
 - **Substitute (Domain Name) Rule:** For information, see [Managing Substitute \(Domain Name\) Rules](#) on page 1242.
 - **Substitute (Record) Rule:** For information, see [Managing Substitute \(Record\) Rules](#) on page 1243.
3. Complete the details in the corresponding editor.
4. Save the configuration and click **Next** to define extensible attributes. For information about extensible attributes, see [Using Extensible Attributes](#) on page 332.

Configuring the FireEye appliance

You must configure the FireEye appliance to send alerts to the NIOS appliance. Ensure that the following are complete before you configure the FireEye appliance:

1. Install required license on the NIOS appliance. For more information about license, see [License Requirements and Admin Permissions](#) on page 1235.
2. Create a new FireEye RPZ zone. For more information, see [Configuring FireEye RPZs](#) on page 1254.
3. Create FireEye admin users. For more information, see [For FireEye Integrated RPZs](#) on page 1235.

4. Get the URL from the NIOS appliance and record it. You need this to configure the FireEye appliance. For more information about the Server URL, see [Configuring FireEye RPZs](#) on page 1254. If you have already configured a FireEye integrated RPZ, then you can retrieve the URL through the **FireEye** tab of the corresponding FireEye RPZ zone. For more information about managing and retrieving the URL, see [Modifying RPZs](#) on page 1262.
5. Record the usernames and passwords on the NIOS appliance. You must use these credentials when configuring FireEye alerts to enable the alerts to be received by NIOS. For more information, see [Configuring the FireEye appliance to send alerts to NIOS](#) on page 1258.

Configuring the FireEye appliance to send alerts to NIOS

You must configure the NIOS generated URL, usernames and passwords on the FireEye appliance. FireEye appliance embeds the configured usernames and passwords in the alerts for authentication. When an alert is received, the NIOS appliance verifies the FireEye username prior to processing the alert. Note that the NIOS appliance accepts alerts sent by the FireEye appliance in **JSON Normal** format only.

To configure a FireEye appliance:

1. Login to the FireEye appliance with your username and password.
2. In the FireEye GUI, click **Settings** tab and then click the **Notifications** tab on the left panel.
3. In the **Notification Settings** page, click the **http** link and then enter the name of the HTTP server you want to add. Click **Add HTTP Server** and complete the following:
 - **Name:** When you click add, the HTTP server name that you specified is listed in this column.
 - **Enabled:** Select the check box to enable alerts and notifications for the HTTP server.
 - **Server Url:** Enter the URL you received on the NIOS appliance. The alerts and notifications are sent using this URL by the FireEye appliance.
 - **Auth:** Select this check box if authentication is required for the server.
 - **Username and Password:** Enter the Username and Password of the user that you have configured for the **fireeye-group** on the NIOS appliance. For more information, see [For FireEye Integrated RPZs](#) on page 1235.
 - **Notification:** Select a notification from the drop-down list. You can choose to include notifications for all events or only events of a selected type. The FireEye appliance will send an alert to the NIOS appliance only when selected event is encountered. When you select **All Events**, alerts are sent when each event is encountered by the FireEye appliance.
 - **Delivery:** Select **Per Event** from the drop-down list. Note that the NIOS appliance supports only **Per Event** selection. The FireEye appliance sends an alert each time it encounters an event.
 - **Account:** You can specify a user account name for this notification.
 - **SSL Enable:** Select this check box to enable SSL for secure transmission of alerts from the FireEye appliance to NIOS.
 - **Default Provider:** Select a default provider from the list.
 - **Message Format:** Select **JSON Normal** from the drop-down list. Note that the NIOS appliance supports only this message format.
4. Click **Update** at the bottom of the page.

Note: You can also click **Test-Fire** to test the configuration. If the configuration is successful, FireEye sends a confirmation message to the NIOS appliance and the NIOS appliance logs this message in the syslog. It generally takes a few seconds for the NIOS appliance to receive alerts. You must verify the configuration, if there is no entry in the syslog.

Handling Alerts from the FireEye appliance

The NIOS appliance processes each alert that it receives from the FireEye appliance. The alert contains the malware URL along with a valid FQDN. NIOS appliance can only map an alert to a RPZ rule if the FQDN is present. Once the alert is processed and properly mapped to an RPZ rule, it remains in the database until you delete it manually. When the RPZ rule is different from the existing rules, the new RPZ rule gains precedence over the existing RPZ rule in the FireEye integrated RPZ. Note that you cannot retrieve alerts that are ignored. You can get more information about the alerts, which are sent by the FireEye appliance, from the syslog. An alert will not be processed and will be ignored:

- when there are changes to the URL or if the alert does not have the malware URL or FQDN in them.
- if the zone is not found.
- if the alert is sent without any username in it or if the username does not belong to the fireeye-group.
- if a FireEye admin user is deleted. NIOS will neither authenticate the deleted user credentials nor process any future alerts with deleted user credentials.
- if the search mapping fields contain IP addresses other than FQDNs.
- if alerts contain domain names in an IPv4 or IPv6 address format.

Logging FireEye Integrated RPZ messages

The NIOS appliance logs FireEye events and alerts in the syslog and audit log. Each FireEye feed event is logged every time an alert is sent to NIOS by the FireEye appliance. When you create a new rule or update an existing rule, then those are also logged in the syslog. You can use messages logged in syslog to verify events that are related to communication between the FireEye and NIOS appliances. It also enables the admin to monitor alerts and verify how the alerts are processed. Details about alerts that are received and processed are also logged. Syslog messages are logged when:

- an alert is received from the FireEye appliance.
- syslog messages contain required information for reporting.
- an alert is successfully mapped to an RPZ rule. The message format is as follows:
 - <FireEye: Found an APT alert>
- the NIOS appliance cannot process alerts. For example, alert structure mismatch, unrecognizable data, etc. The messages will have the following format:
 - <FireEye: Cannot parse FQDN due to missing field"cnc-services">
 - <FireEye: Cannot determine if it is an APT alert..>
 - <FireEye: Invalid Alert Type>
 - <FireEye: Couldn't find the required field...>
 - <FireEye: No mapping rule has been set for alert type.....>
- a duplicate alert is sent by the FireEye appliance for which the same RPZ rule already exists.

Note: For debugging purposes, alert messages will be displayed in the infoblox.log file.

NIOS periodically scans the syslog of a member that has RPZ license installed to generate recent hits data for the *RPZ Recent Hits* tab. This might cause a performance impact as CPU cycles will be used on the member. For more information about *RPZ Recent Hits* tab, see [Response Policy Zone \(RPZ\) Statistics](#) on page 134.

Configuration Examples

This section illustrates some of the examples of local and FireEye integrated RPZs.

Local RPZ Examples

Following is an example of an IP related rule. For example, execute the following command:

```
dig @10.35.104.19 abc.net
```


If the above command returns 18.58.20.1, then define an IPv4 substitute rule 18.0.0.0/8, Substitute (IPv4) 8.8.8.89.

Execute the command `dig @10.35.104.19 abc.net` again. You will receive the substituted address instead of the actual domain name.

Following is an example of values in CEF for the above substitution example:

```
2012-11-06T19:04:02+00:00 daemon (none) named[25193]: info
CEF:0|Infoblox|NIOs|6.6.0-185622|RPZ-IP|records|4|app=DNS dst=10.35.104.19 src=10.32.0.242
spt=50035 view=_default qtype=A msg="rpz IP records rewrite abc.net [A] via
8.0.0.0.18.rpz-ip.localrpz"
```

Following is an example of values in the CEF for Block (No Data) rules:

```
2012-11-06T19:00:01+00:00 daemon (none) named[25193]: info
CEF:0|Infoblox|NIOs|6.6.0-185622|RPZ-QNAME|NODATA|4|app=DNS dst=10.35.104.19
src=10.32.0.242 spt=50035 view=_default qtype=A msg="rpz QNAME NODATA rewrite nodata.net [A]
via nodata.net.localrpz"
```

You can view the NIOS version, name of the view, source, and destination.

FireEye Integrated RPZ Examples

Following is an example of a syslog message when an alert gets converted to an RPZ rule:

```
013-09-11T10:59:55-07:00 user (none) httpd[]: info fireeye-rpt:
'79167', 'infection-match', 'minr' 'eng-lab-249.inca.infoblox.com'
2013-09-11T10:59:55-07:00 user (none) httpd[]: info FireEye: Create an RPZ rule for
'd.bnksw.com' with 'SUBSTITUTE' rule in RPZ zone 'com.lock'
```

Note that the domain name lock.com is displayed in the reverse format.

Following is an example of a syslog message when an alert is ignored by the NIOS appliance:

```
2013-09-11T11:04:01-07:00 user (none) httpd[]: info fireeye-rpt:
'114488', 'malware-object', 'majr' 'eng-lab-249.inca.infoblox.com'
2013-09-11T11:04:01-07:00 user (none) httpd[]: info FireEye: Cannot parse FQDN due to
missing field 'cnc-services'
```

Example of a basic RPZ Workflow

Following is an example of a basic RPZ workflow:

1. Install the RPZ license. For more information, see [License Requirements and Admin Permissions](#) on page 1235.
2. Enable recursive queries for a DNS view, member, or Grid, as described in [Enabling Recursion for RPZs](#) on page 1237.
3. Enable RPZ logging in the *Grid DNS Properties* editor to view syslog entries for RPZ queries. For more information, see [Setting DNS Logging Categories](#) on page 1015.
4. Create a local RPZ. For more information, see [Configuring Local RPZs](#) on page 1238.
5. Define a **Substitute (PTR Record) Rule** for domain name 3.3.3.5.in-addr.arpa, which is substituted with the domain name ptr1.com. For more information, see [Defining Substitute Rules for PTR Records](#) on page 1245.
6. Execute the dig command to view output. The output contains the substituted domain name ptr1.com. Following is the output of an RPZ query for **Substitute (PTR Record) Rule**:

```
$ dig @10.36.2.73 3.3.3.5.in-addr.arpa in ptr
; <<>> DiG 9.6.2-P2-RedHat-9.6.2-5.P2.fc12 <<>> @10.36.2.73 3.3.3.5.in-addr.arpa in ptr
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 7351
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0
```

```
;; QUESTION SECTION:
3.3.3.5.in-addr.arpa.      IN      PTR

;; ANSWER SECTION:
3.3.3.5.in-addr.arpa.      7200   IN      PTR      ptr1.com.

;; Query time: 3 msec
;; SERVER: 10.36.2.73#53(10.36.2.73)
;; WHEN: Thu Sep 26 23:27:10 2013
;; MSG SIZE rcvd: 60
```

7. Following is the syslog entry for the query mentioned above:

```
2013-09-27T02:26:46-04:00 daemon (none) named[21737]: info
CEF:0|Infoblox|NIOS|6.9.0-218052|RPZ-QNAME|Local-Data|4|app=DNS dst=10.36.2.73
src=10.120.20.194 spt=40518 view=2 qtype=PTR msg="rpz QNAME Local-Data rewrite
3.3.3.5.in-addr.arpa [PTR] via 3.3.3.5.in-addr.arpa.local1.com"
```

For more information about syslog, see [Viewing RPZ in the Syslog](#) on page 1266.

MANAGING RPZS

You can manage RPZs that you defined earlier and modify their information. You can do the following:

- View RPZs, as described in [Viewing RPZs](#) on page 1261.
- Modify RPZs, as described in [Modifying RPZs](#) on page 1262.
- Reorder RPZs, as described in [Reordering RPZs](#) on page 1263.
- Lock and unlock RPZs, as described in [Locking and Unlocking RPZs](#) on page 1263.
- Delete RPZs, as described in [Deleting RPZs](#) on page 1263.

Viewing RPZs

You can view the list of RPZs, local, feed, or FireEye integrated RPZs, which are currently listed in the Grid. To view RPZs:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab.
2. Grid Manager displays the following:
 - **Order:** Displays the order of RPZs. The order value is empty if you do not assign a primary name server when configuring a local RPZ, or if the local RPZ or the service is disabled.
 - **Name:** Displays the name of the RPZs. Click *Zone* to view the following details:
 - **Name or Address:** Displays the domain name or the IP address.
 - **Policy:** Defines the policy defined for the corresponding domain name or IP address.
 - **Data:** Displays the target data of the rule.
 - **Comment:** Displays the comment specified when an RPZ is defined.
 - **Disabled:** Displays **Yes** if the RPZ rule is disabled.
 - **Site:** Displays extensible attributes that are associated with the domain name or IP address.
 - **Type:** Displays the type of RPZs, that is, **Local**, **Feed**, or **FireEye**.
 - **Primary Name Server:** Displays the primary name server that is associated with an RPZ.
 - **Last Updated:** Displays the last updated time. For RPZ feed, it indicates if the RPZ feed has stalled and when the last zone transfer happened. For a local and FireEye integrated RPZ, it indicates the last time the zone or data was modified.

- The last updated time is empty, if:
 - A local RPZ is not associated with a primary Grid name server.
 - A zone, either a local RPZ or an RPZ feed, is not enabled.
 - An inbound zone transfer has not occurred for an RPZ feed.
 - Member's DNS service is disabled.
- **Comment:** Displays the comment recorded when creating the zone. You can double-click on a row to edit the comment. Click **Save** after modification
 For FireEye integrated RPZs, this column displays the comment recorded when creating the FireEye integrated RPZ. The rules that are created from the FireEye alerts will have alert information in this column. This differentiates between fireeye alert created rules and user created rules. You can double-click on a row to edit the comment. Click **Save** after modification. Infoblox recommends that you do not modify any internal objects. For example, the **Comment** column has alert related information, if you modify the data, then the actual alert data will be compromised.
- **Disabled:** Displays **Yes** if the RPZ is disabled. Otherwise, this field displays **No**.
- **Locked:** Displays **Yes** when a zone is locked by an admin, and displays **No** when the zone is unlocked.
- **Site:** Displays the values that were entered for this pre-defined attribute. You can double-click on a row to edit the Site. Click **Save** after modification.

You can also do the following:

- Use **Quick Filter** and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. Select a value from the drop-down list to filter the RPZs.
 - **None:** Select this to display all the RPZs that you have configured.
 - **All Local Response Policy Zones:** Select this to list only the local RPZs.
 - **All Feed Response Policy Zones:** Select this to list only the RPZ feeds.
 - **All FireEye Response Policy Zones:** Select this to list only the FireEye RPZs.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.
- You can create a bookmark for the RPZs. For information, see [Using Bookmarks](#) on page 63.
- You can modify some of the data in the table. Double-click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- To export the list of RPZs to a .csv file, click the *Export* icon. For information on the export options, see [About CSV Import](#) on page 86.
- Click the *Print* icon to print the list of RPZs. For more information, see [Printing from Grid Manager](#) on page 91.

Modifying RPZs

You can modify the name servers or name server groups, update policy override details and permissions, or edit extensible attributes that are associated with an RPZ.

To modify RPZs:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab -> *Zone* check box and then click the *Edit* icon.
2. The RPZ editor provides the following tabs from which you can modify data:
 - In the **General** tab, you can change the information you previously entered through the wizard, as described in [Configuring Local RPZs](#) on page 1238. For FireEye integrated RPZs, you can update the policy type, comments, enable or disable, or lock the zone. For more information, see [Configuring FireEye RPZs](#) on page 1254.

- For a FireEye integrated RPZ, the **FireEye** tab is displayed. This tab is displayed only after you install the FireEye license. You can modify or override the rule mapping for FireEye alerts or APT events. For more information, see [Configuring FireEye RPZs](#) on page 1254.
 - You can also enter or edit information in the **Name Servers**, **Extensible Attributes**, **Settings** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) on page 679.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Reordering RPZs

You can change the order of RPZs, local, feed, or FireEye integrated RPZ, in each view. When you add a new local RPZ, it is added to the top of the zone list and an RPZ feed is automatically added to the bottom of the zone list. You can change the order of each through the re-ordering process.

The policy override works based on zone ordering. The zone at top has the highest priority and it overrides the lower priority zone. To override an RPZ feed with a local RPZ, place the RPZ feed at the top before a local RPZ. You cannot reorder zones, if they are disabled or do not have any primary name server assigned.

To reorder RPZs:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click **Order Response Policy Zones** from the **Toolbar**.
2. The following are displayed in the *Order Response Policy Zones* wizard:
 - **Ordering:** Use the up and down arrows to move the RPZ to the desired order.
 - **Response Policy Zone:** Displays all the RPZs.
 - **Priority:** Displays the order of RPZs.
3. Click **OK** to save the changes.

Locking and Unlocking RPZs

You can lock an RPZ so you can make changes to it and prevent others from making conflicting changes. When you lock an RPZ, Grid Manager displays LOCKED beside the RPZ. When other administrators try to make changes to a locked RPZ, the system displays a warning message that the RPZ is locked and the name of the admin who locked the RPZ.

Only a superuser or the administrator who locked the RPZ can unlock it. RPZ locks do not expire; you must manually unlock a locked RPZ. To lock or unlock RPZs:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select the *Zone* -> *Ruleset*.
2. You can do the following:
 - **To Lock:** Click the *Lock* icon to lock the zone.
 - **To Unlock:** Click the *Unlock* icon to unlock the zone.

Deleting RPZs

You can delete RPZs or schedule them for deletion for a later date. The NIOS appliance moves the deleted RPZs to the Recycle Bin, if enabled. When you restore the zone from the Recycle Bin, it will be restored to the bottom of the zone list.

To delete RPZs:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> *Zone* check box.
2. To delete an RPZ immediately, click the *Delete* icon, and then click **Yes** to confirm the delete request. To schedule the deletion, click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#) on page 76.

Grid Manager moves the RPZ to the Recycle Bin, from which you can restore or permanently delete it.

MANAGING RPZ RULES

You can manage local RPZ, including FireEye integrated RPZ rules that you defined earlier and modify their information. You can do the following:

- View RPZ rules, as described in [Viewing RPZ Rules](#) on page 1264.
- Modify RPZ rules, as described in [Modifying RPZ Rules](#) on page 1265.
- Delete RPZ rules, as described in [Deleting RPZ Rules](#) on page 1265.
- Copy RPZ rules, as described in [Copying RPZ Rules](#) on page 1265.
- Import RPZ rules, as described in [Importing RPZ Rules](#) on page 1266.

Note: You cannot modify the rules of an RPZ feed. However, you can override the entire ruleset or each rule using local RPZs.

Viewing RPZ Rules

You can view and edit the rules that are defined for each local RPZ, including FireEye integrated RPZs.

To view RPZ rules:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab -> click *Zone*.
2. You can view the following:
 - **Name or Address:** Displays the domain name or the IP address on which the rule is defined.
 - **Policy:** Displays the rule applied on the domain name or the IP address.
 - **Data:** Displays the target data of the rule.
 - **Comment:** Displays the comment specified when the rule is defined.
 - **Disabled:** Displays **Yes** if the RPZ rule is disabled. Otherwise, this field displays **No**.
 - **Site:** Displays an extensible attribute, **Site**.
 - **Expiration:** Displays the expiration time for the corresponding FireEye integrated RPZ rule. Note that NIOS updates the expiration time when you change the lifetime of the FireEye integrated RPZ rule, or if the last updated time of the rule changes, or if the alert type that generates the rule changes. This time is estimated based on the following:
 Expiration Time = Lifetime of an alert type + Last updated time of the rule
 NIOS runs a scheduler every 10 minutes to identify FireEye integrated RPZ rules whose expiration time is less than the current time. If there are rules whose expiration time is less than the current time, then such rules will be deleted. NIOS logs all deletion activities in the syslog. You can view the syslog to verify expired rules. For more information, see [Viewing RPZ in the Syslog](#) on page 1266.
 - **FireEye Alert Type:** Displays the type of FireEye alert.
 - **Last Updated:** Displays the time when the RPZ rule was last updated.

Note: The columns, **Expiration**, **FireEye Alert Type**, and **Last Updated**, are displayed only for FireEye integrated RPZ rules. These columns are not displayed for non-FireEye RPZ rules.

You can also do the following:

- Use **Quick Filter** and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) on page 68.

- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. You can edit **Comments** and **Extensible Attributes**. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#) on page 62.
- To export the list of RPZ rules to a .csv file, click the *Export* icon. For information on the export options, see [About CSV Import](#) on page 86.
- Click the *Print* icon to print the list of RPZ rules. For more information, see [Printing from Grid Manager](#) on page 91.

Modifying RPZ Rules

You can modify the name of a local or FireEye integrated RPZ rule, IP address, network address, substituted name, and the comment recorded for the corresponding rule. You can also update the TTL settings or the extensible attributes that are associated with an RPZ rule.

To modify RPZ rules:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab -> click *Zone* -> *Name or Address* check box, and then click the *Edit* icon.
2. The RPZ rules editor provides the following tabs from which you can modify data:
 - In the **General** tab, you can change the information you previously entered through the wizard. For more information, see [Configuring Rules for RPZs](#) on page 1239.
 - You can also enter or edit information in the **TTL and Extensible Attributes** tabs. For information about TTL settings, see [About Time To Live Settings](#) on page 557. For information about extensible attributes, see [Using Extensible Attributes](#) on page 332.
3. Save the configuration.

Deleting RPZ Rules

You can delete local RPZ rules, including FireEye integrated RPZ rules, or schedule them for deletion for a later date. When you remove an RPZ rule, the NIOS appliance moves it to the Recycle Bin, if enabled.

To delete RPZ rules:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> *Zone* -> *Ruleset*.
2. To delete an RPZ rule immediately, click the *Delete* icon, and then click **Yes** to confirm the delete request. To schedule the deletion, click **Schedule Deletion** and in the **Schedule Change** panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#) on page 76.

Grid Manager moves the RPZ rule to the Recycle Bin, from which you can restore or permanently delete it.

Copying RPZ Rules

You can copy rules from one local RPZ to another local RPZ or from one FireEye integrated RPZ to another FireEye RPZ. You can also copy rules from a local RPZ to a FireEye integrated RPZ or vice-versa. Different views of the same RPZ may have a number of rules in common. If this is the case, you can copy rules between views and zones.

To copy RPZs between DNS zones and views:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click **Copy Rules** from the **Toolbar**.
2. In the *Copy Rules* dialog box, Grid Manager displays the last selected zone or the zone from which you are copying rules in the **Source** field. The following fields are displayed:
 - **Source:** Grid Manager displays the last selected zone or the zone from which you are copying rules. It also displays the associated DNS view.
 - **Destination:** Click **Select Zone** to select the destination zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one. After you select the zone, Grid Manager displays the associated DNS view.

- **Copy All Rules:** Select this option to copy all the rules.
- **Copy Specific Rules:** Select this option to copy specific rules only. Select a rule from the **Available** column and click the right arrow to move it to the **Selected** column.
- **Copy Options:** Select one of the following:
 - **Delete all rules in the destination before copying the rules:** Select to delete all rules in the destination zone before the records are copied.
 - **Overwrite existing rules:** Select to overwrite existing rules that have the same domain name owners as the rules being copied.

3. Click **Copy & Close**.

Importing RPZ Rules

You can import rules from an RPZ zone to a local zone. To import, you must enable zone transfer on the external server. The rules of the existing zone are overwritten when you import rules from an external server.

To import RPZ rules:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> **Zones** -> **Rules**, click **Import Zone** from the **Toolbar**.
2. In the *Import Zone* dialog box, the following fields are displayed:
 - **Zone:** The RPZ that you have selected is displayed.
 - **DNS View:** The DNS view that you have selected is displayed.
 - **Address:** Enter the address of the external server from where you want to import rules.
3. Click **Import**.

VERIFYING RPZ CONFIGURATION

After you have set up and configured RPZs and RPZ rules, you can verify whether the RPZ zone transfers are functioning properly by doing the following:

- View the syslog for zone transfer confirmation, as described in [Viewing RPZ in the Syslog](#) on page 1266.
- Verify the last RPZ updates, as described in [Viewing the Last Updated RPZs](#) on page 1267.

Viewing RPZ in the Syslog

To verify RPZ zone transfers:

1. Go to the **Administration** tab -> **Logs** tab -> **Syslog** tab.
2. Select **RPZ Logs** from the **Quick Filter** drop-down list.
3. Review the syslog for zone transfer confirmation, as shown in [Figure 42.3](#).

Figure 42.3 The syslog

2013-02-20 10:27:36 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: Transfer completed: 0 messages, 1 record
2013-02-20 10:27:36 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: connected using 172.16.209.10#45424
2013-02-20 10:27:36 PST	daemon	INFO	named[7583]	zone naughty-list.rpz.infoblox.local/IN: Transfer started.
2013-02-20 10:18:43 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: Transfer completed: 0 messages, 1 record
2013-02-20 10:18:43 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: connected using 172.16.209.10#52922
2013-02-20 10:18:43 PST	daemon	INFO	named[7583]	zone naughty-list.rpz.infoblox.local/IN: Transfer started.
2013-02-20 10:10:55 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: Transfer completed: 0 messages, 1 record
2013-02-20 10:10:55 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: connected using 172.16.209.10#54448
2013-02-20 10:10:55 PST	daemon	INFO	named[7583]	zone naughty-list.rpz.infoblox.local/IN: Transfer started.
2013-02-20 10:02:24 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: Transfer completed: 10 messages, 9812 r
2013-02-20 10:02:24 PST	daemon	INFO	named[7583]	zone naughty-list.rpz.infoblox.local/IN: transferred serial 1360180898: TSIG 'rpz-naughty-list'
2013-02-20 10:02:23 PST	daemon	INFO	named[7583]	transfer of 'naughty-list.rpz.infoblox.local/IN' from 10.60.66.5#53: connected using 172.16.209.10#49333
2013-02-20 10:02:23 PST	daemon	INFO	named[7583]	zone naughty-list.rpz.infoblox.local/IN: Transfer started.
2013-02-12 06:53:39 PST	kern	INFO	kernel[]	pci 0000:00:00:0: Limiting direct PCI/PCI transfers

Viewing the Last Updated RPZs

To view the last updated RPZs:

1. Go to the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab.
2. Review the **Last Updated** column and confirm the time when an RPZ was last updated, as shown in [Figure 42.4](#).

Note: It may take up to 10 minutes before the updated information is displayed.

Figure 42.4 Last Updated RPZ

IPAM DHCP DNS File Distribution						
Zones Members Name Server Groups Shared Record Groups Response Policy Zones Blacklist Rulesets DNS64 Groups						
default						
Quick Filter None Off Filter On Show Filter						
Go to						
	Order	Zone	Type	Primary Name Server	Last Updated	Comment
	0	override_policy	Local	infoblox.localdomain	2013-03-01 22:18:13 PST	
<input checked="" type="checkbox"/>	1	naughty-list.rpz.infoblox.local	Feed	rpz-server.edu.tme.infoblox.com	2013-03-01 22:44:36 PST	



PART 10 REFERENCE

This section provides reference information in the following appendices:

- [Appendix A, "Glossary of Terms"](#), on page 1271
- [Appendix B, "Grid Manager Icons"](#), on page 1279
- [Appendix C, "Guidance Documentation Supplement"](#), on page 1285
- [Appendix D, "Regular Expressions"](#), on page 1303
- [Appendix E, "vNIOS Appliance Limitations"](#), on page 1305
- [Appendix F, "Product Compliance"](#), on page 1309
- [Appendix G, "Open Source Copyright and License Statements"](#), on page 1317
- [Appendix H, "Threat Protection Rules"](#), on page 1371



Appendix A Glossary of Terms

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

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

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

Term	Description
Extensible Attribute	Metadata you define to capture additional information about an object managed by the Infoblox NIOS appliance. You can use predefined attributes or create your own. You can also specify required attributes and restrict the values that users can enter for each attribute.
Filters	Criteria the Infoblox NIOS appliance uses to request specific information in the database. You can use filters to control the amount and the kind of data displayed in a panel or table in Infoblox Multi-Grid Manager, Grid Manager, and System Manager.
FQDN (fully qualified domain name)	A complete domain name that specifies its exact location in the hierarchy of the DNS. It specifies all the domain levels, including the top-level domain and the root domain.
FTP (File Transfer Protocol)	A standard network protocol used to transfer files from one network device to another over a TCP-based network, such as the Internet. FTP is built on a client-server architecture and utilizes separate control and data connections between the client and server.
Gateway	The default router for the immediate network segment of an interface.
Grid™ Technology	Infoblox's unique and patented high availability Grid technology ensures network reliability. The Infoblox Grid provides resilient network services, failover, recovery, and seamless maintenance for an Infoblox deployment inside a single building, across a networked campus, or between remote locations. The Infoblox Grid establishes a distributed relationship between individual or paired appliances to remove single points of failure and other operational risks inherent in legacy DNS, DHCP, and IP address management infrastructure.
Grid Manager	The NIOS web interface that provides access to your Grid for performing IPAM, DNS, and DHCP management and other administration tasks.
Grid Master	The Grid member in an Infoblox Grid that maintains the NIOS database that is distributed among all members of the Grid. You connect to the Grid Master to configure and monitor the entire Grid.
Grid Member	Any single Infoblox NIOS appliance or HA pair that belongs to a Grid. Each member can use the data and services of the Grid. You can also modify settings so that a Grid member can use unique data and member-specific services.
HA Pair	Two physical Infoblox NIOS appliances that are linked to perform as a single virtual appliance in an HA (high availability) configuration. The HA configuration provides hardware redundancy to minimize service outages. In this configuration, one appliance is the active node and the other is the passive node.
Host Record	On Infoblox appliances, host records provide a unique approach that enables you to manage multiple DNS records and DHCP and IPAM data collectively, as one object on the appliance.
IBOS (Infoblox Orchestration Server)	IBOS is the Infoblox IF-MAP (Interface to Metadata Access Points) server that contains a searchable database for storing state information about network resources. It is the central point with which IF-MAP clients communicate to send and retrieve real-time information defined in the IF-MAP data format.
IF-MAP (Interface for Metadata Access Points)	An open standard client-server protocol developed by the Trusted Computing Group as one of the core protocols of the TNC (Trusted Network Connect) open architecture. IF-MAP allows network resources to share real-time information.
IP Map	In Infoblox Grid Manager or System Manager, this is a graphical representation of all IPv4 addresses in a given subnet.

Term	Description
IPAM (IP Address Management)	Infoblox IPAM provides a means of planning, tracking, and managing IP address space in a network. It glues DNS and DHCP services together so that each service is aware of changes in the other. The Infoblox IPAM implementation offers an IP address-centric approach so you can manage your networks and IP addresses through a centralized GUI.
Leaf Network	On Infoblox appliances, a network that does not contain any subnets.
Lease Logging Member	An Infoblox Grid member that is designated to collect DHCP lease events.
Limited-Access User	An admin user account that has specific roles and permissions assigned. Limited-access users have restricted access to Infoblox Multi-Grid Manager, Grid Manager, and System Manager, and can only perform certain tasks based on their assigned roles and permissions.
Lite Upgrade	On Infoblox appliances, a lite upgrade occurs when there are incremental changes to the NIOS software that do not require any change to the database. The appliance can perform a lite upgrade only if the format of the database between the existing NIOS version and the upgrade version is the same. In general, when you upgrade from a major release to a patch release or a patch release to another patch release, you are performing a lite upgrade.
Loopback Interface	On Infoblox appliances, the virtual network interface on which you can consolidate DNS servers for migration purposes, add anycast addresses to improve the performance of the DNS service, and separate DNS traffic.
Managing Member	An Infoblox Grid member that is configured to manage Microsoft DNS and DHCP servers.
Master Candidate	An Infoblox Grid member that is designated to assume the role of the Grid Master as a disaster recovery measure.
Master Grid	A group of Infoblox appliances that are connected to provide a single point of administration for multiple Grids and network management of these Grids.
Master Grid Member	Any single Infoblox appliance or HA pair that belongs to the Master Grid. All Master Grid members serve as Master Candidates.
Multi-Grid Manager	The NIOS web interface that provides access to the Master Grid, from which you can manage multiple Grids and their networks.
Multi-Grid Master	The Infoblox Master Grid member that maintains the NIOS database that is distributed among all Master Grid members. You connect to Multi-Grid Manager to configure and monitor the Master Grid.
Multi-Grid Master Candidate	An Infoblox Master Grid member that is designated to assume the role of the Multi-Grid Master as a disaster recovery measure.
Name Server Group	On Infoblox appliances, a server group that contains one primary DNS server and/or one or more secondary DNS servers. Specifying a single name server group can simplify DNS zone creation.
NAT (Network Address Translation) Group	A group of Infoblox Grid members that are configured on the same side of a NAT appliance. In a Grid configuration where the Grid Master is configured behind a NAT appliance and there are Grid members on both sides of the NAT appliance, it is necessary to create a NAT group to ensure that the Grid Master and Grid members use the correct NAT and interface addresses for Grid communications.
Network Block	On Infoblox appliances, an IP address space that is defined in the Master Grid. A network block can consist of other network blocks, network containers, and leaf networks.

Term	Description
Network Container	On Infoblox appliances, an automatically created container of multiple networks that are subnets of the IP address space configured for the network container. A network container cannot be assigned to a Grid member or be directly created.
Network Discovery	A set of tools provided by the Infoblox NIOS appliance for detecting active hosts on specified networks and specified VMware vSphere servers.
Network Map	In Infoblox Grid Manager and System Manager, Network Map presents a complete view of your network space, including the different types of networks that are in it and its unused address space. You can use Network Map to design and plan your network infrastructure, configure and manage individual networks, and evaluate their utilization.
Network Mask or Netmask	A numeric representation of the bits that are used to split an IP address into the network portion and the host portion. In Infoblox products, this is represented by either quad-dotted decimal representation or CIDR notation for IPv4 network masks, or by CIDR notation for IPv6 network masks.
Network View	On Infoblox appliances, a single routing domain with its own networks and shared networks. A network view can contain both IPv4 and IPv6 networks. All networks must belong to a network view on the Infoblox appliance.
NIOS	An Infoblox proprietary system that powers Infoblox solutions with an embedded processor that delivers core network services. It is the operating system that runs on the NIOS appliances—a security-hardened, real-time set of appliances built to ensure the non-stop operation of network infrastructure. NIOS automates the error-prone and time-consuming manual tasks associated with deploying and managing IPAM, DNS, and DHCP required for continuous IP network availability and business uptime.
NIOS Virtual Appliance	Any Infoblox supported platform, such as the Riverbed Steelhead appliances or VMware appliances, that runs the vNIOS software. These appliances are also known as the vNIOS appliances.
Node	A single Infoblox appliance of an HA (high availability) pair. An HA pair consists of an active node and a passive node.
NTP (Network Time Protocol)	A protocol for synchronizing the clocks of computer systems over packet-switched, variable latency data networks; it essentially keeps network devices on a common clock by resisting the effects of variable latency by means of a jitter buffer.
Passive Node	The Infoblox NIOS appliance in an HA pair that constantly keeps its database synchronized with that of the active node, so it can take over core network services when an HA failover occurs. When an HA failover occurs, the passive node becomes the active node in the HA pair.
PortIQ	An Infoblox switch port appliance that enables quick discovery of the Ethernet switch ports. PortIQ identifies ports that are not fully utilized and those that exceed their capacity. You can use PortIQ to troubleshoot LAN environments.
Quick Filter	A filter that stores specific filter criteria for requesting information displayed in a specific panel in Infoblox Multi-Grid Manager, Grid Manager, and System Manager. For more information, see “Filter.”
Overlapping Network	On Infoblox appliances, a network that exists in multiple locations, which can be multiple Grids in the Master Grid or within various network views in a Grid.
Replication	Database distribution among the Infoblox Grid Master and Grid members as well as among the Multi-Grid Master and Master Grid members.

Term	Description
Reservation	On Infoblox appliances, a static IP address that you create for future use. A reservation is a pre-provisioned fixed address. You can reserve this static IP address on the NIOS appliance and assign it to a client in the future.
Resource Records	A collection of data in the DNS server database. Each resource record specifies information about a DNS object. For example, an A (address mapping) record maps a host name to an IP address, and a PTR (reverse-lookup pointer) record maps an IP address to a host name. The DNS server uses these records to answer queries.
Roaming Host	On Infoblox appliances, a host with a dynamically assigned IP address and a specific set of properties and DHCP options. When you create a roaming host for a network device, the device can receive any dynamically assigned address from the network to which it belongs.
Scope	A DHCP address range on a Microsoft server. Microsoft scope information is converted to equivalent DHCP range information after Microsoft data is synchronized with the NIOS appliance.
Shared Network	On Infoblox appliances, a network segment to which you assign two or more subnets. When subnets in a shared network contain IP addresses that are available for dynamic allocation, the addresses are put into a common pool for allocation when client requests arise.
Shared Record Group	On Infoblox appliances, a set of resource records that you add to multiple DNS zones. You can create resource records in a group and share the group among multiple zones. The zones handle the shared resource records as any other resource record.
SSO (Single Sign On)	An Infoblox feature that allows you to automatically sign in to selected Grids from the Master Grid, without having to log in to each individual Grid each time you sign on.
Smart Folder	On Infoblox appliances, a virtual folder in which you place the results of filter criteria that you select to request specific data in the NIOS database. Once you set up a smart folder, the appliance displays up-to-date information based on your filter and grouping criteria each time you access the folder.
Subnet (or network)	A logical division of an IP network. A subnet of network may also be called a network. For example, 10.1.0.0/16 is a subnet of 10.0.0.0/8, and fc80:8:8:16::/64 is a subnet of fc80:8:8::/48.
Superscope	On a Microsoft server, superscope comprises multiple scopes or DHCP address ranges created on a single physical network segment. Microsoft superscope information is converted to equivalent network information after Microsoft data is synchronized with the NIOS appliance.
Superuser	An admin user account that has unrestricted access to Infoblox Multi-Grid Manager, Grid Manager, or System Manager.
Support Bundle	A tar.gz file that contains configuration files and system files of the Infoblox NIOS appliance. You can download a support bundle for an independent appliance and for each member in a Grid.
System Manager	The NIOS web interface that provides access to an independent appliance (single or HA) for performing IPAM, DNS, and DHCP management and other administration tasks.
TFTP (Trivial File Transfer Protocol)	A data transfer service that provides devices—such as phones, RFID readers, IP cameras, and other devices—with up-to-date software and configuration data.

Term	Description
Traffic Capture	An Infoblox tool that captures the traffic on one or all of the ports on a NIOS appliance. The NIOS appliance saves all captured traffic in a .cap file and compresses it into a .tar.gz file.
Upgrade Group	On Infoblox appliances, a group of Grid members that you put together so you can perform software distribution and upgrade at the same time.
VIP (Virtual IP)	On Infoblox appliances, the shared IP address of an HA pair. A VIP address links to the HA port on the active node of an HA pair.
VRID (Virtual Router ID)	VRID identifies the VRRP (Virtual Router Redundancy Protocol) HA pair to which the Infoblox appliance belongs. Through VRID, two HA nodes identify each other as belonging to the same HA pair, and they obtain a virtual MAC address to share with a VIP. A VRID can be any number between 1 and 255, and it must be unique on the local LAN so that it does not conflict with any other Infoblox appliances using VRRP on the same subnet.
vNIOS	The virtual version of NIOS. You can install Infoblox vNIOS software on any supported virtual platform and configure the system as a vNIOS virtual appliance.
VRRP (Virtual Router Redundancy Protocol)	An industry standard MAC address level HA failover mechanism.

























Appendix B Grid Manager Icons

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
















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













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

Icon	Icon Name	Description
	Add	Adds an object
	Add Bookmark	Adds a bookmark for an object and displays it in the Bookmarks panel
	Arrow (Down)	Moves an object down in a list
	Arrow (Up)	Moves an object up in a list
	Clear	Clears the status of an object
	Clock	Displays a drop-down list for time
	Delete	Deletes an object
	Disabled	Indicates a disabled object
	Download	Downloads a file or data
	Edit	Displays the corresponding editor for modifying object configurations
	Edit	Displays the corresponding editor for modifying object configurations
	Execute Now	Executes a scheduled task immediately




Icon	Icon Name	Description
	Export	Exports data in the current panel
	Extensible Attribute	Configures extensible attributes for the selected object
	Flat View	Displays a list of objects in a flat view
	Help	Displays information about an object
	Hierarchy	Displays objects in a hierarchical view
	Import	Imports a file or data
	Import Job Manager	Imports CSV data
	Information	Displays informational data about an object
	Locked	Indicates a locked object
	Microsoft Server	Indicates a Microsoft server
	Pause	Pauses a function
	Print	Prints the information in the current panel
	Refresh	Refreshes the current page or table
	Report	Displays a report, such as the capacity report
	Search	Searches for specific objects
	Selected object	Selects an object in a table for a specific function
	Start	Starts a process
	Stop	Stops a process
	Unlocked	Indicates an unlocked object
	User Profile	Configures a user profile
	View	Lists data in the current panel or lists detailed status about an object
	Warning	Indicates a warning message

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







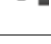

Icon	Icon Name	Description	Tab/Table/Panel
	Configure	<ul style="list-style-type: none"> Configures DHCP properties Configures File Distribution properties Configures Licenses 	<ul style="list-style-type: none"> Data Management tab -> DHCP tab -> Toolbar Data Management tab -> DHCP tab -> Toolbar Grid tab-> Grid Manager tab -> Toolbar
	Conflict	Indicates an IP address conflict	Data Management tab -> IPAM tab -> Net Map
	Convert	Converts an object	Data Management tab -> IPAM tab -> <i>network</i> -> IP Map -> Toolbar
	Discovery	Performs a network discovery	Data Management tab -> IPAM tab -> Toolbar
	Force HA Failover	Forces an HA failover	Data Management tab -> DHCP tab -> Toolbar
	Force Recovery	Forces a recovery	Data Management tab -> DHCP tab -> Members tab -> Failover Associations tab -> Toolbar
	Grid Manager	Indicates the Grid Master	Data Management tab -> DHCP tab -> Members tab -> Data Management tab -> IPAM tab
	Grid Manager Candidate	Indicates the Grid Master candidate	Data Management tab -> DHCP tab -> Members tab -> Data Management tab -> IPAM tab
	Grid Member	Indicates the Grid member	Data Management tab -> DHCP tab -> Members tab -> Data Management tab -> IPAM tab
	Join	Joins networks	Data Management tab -> IPAM tab -> <i>network</i> -> Toolbar
	Key-signing Key Rollover	Indicates the key-signing key that is due to rollover	Data Management tab -> DNS tab
	Leaf Network	Indicates a leaf network	Data Management tab -> IPAM tab or DHCP tab
	Disabled Leaf Network	Indicates a disabled leaf network	Data Management tab -> IPAM tab or DHCP tab
	Microsoft Server	Indicates a Microsoft server	Data Management tab -> DHCP tab -> Members tab -> Data Management tab -> IPAM tab
	Multi-Ping	Pings all the addresses in a network	Data Management tab -> IPAM tab -> IP Map -> Toolbar

Icon	Icon Name	Description	Tab/Table/Panel
	Network	Indicates a network	Data Management tab -> IPAM tab or DHCP tab
	Network Container	Indicates a network container	Data Management tab -> IPAM tab or DHCP tab
	Network (Disabled)	Indicates a disabled network	Data Management tab -> IPAM tab or DHCP tab
	Microsoft Network	Indicates a network with Microsoft servers	Data Management tab -> IPAM tab or DHCP tab
	Infoblox Network	Indicates a network with Infoblox appliances	Data Management tab -> IPAM tab or DHCP tab
	Ping	Pings an IP address	Data Management tab -> IPAM tab -> IP Map -> Toolbar
	Properties	Configures Grid DNS properties	Data Management tab -> DNS tab -> Toolbar
	Reclaim	Reclaims an IP address	Data Management tab -> IPAM tab -> IP Map -> Toolbar
	Resize	Resizes a network	Data Management tab -> IPAM tab -> <i>network</i> -> Toolbar
	Resolve Conflict	Resolves an IP address conflict	Data Management tab -> IPAM tab -> IP Map -> Toolbar
	Set Partner Down	Sets partner down	Data Management tab -> DHCP tab -> Members tab -> Failover Associations tab -> Toolbar
	Split Network	Splits a network	Data Management tab -> IPAM tab -> <i>network</i> -> Toolbar
	DNSSEC status	Displays status for DNSSEC	Data Management tab -> DNS tab -> Toolbar
	Secondary Zone Status	Displays status for the secondary zone	Data Management tab -> DNS tab
	Zoom In	Zooms in to the selected network	Data Management tab -> IPAM tab -> Net Map
	Zoom Out	Zooms out from the selected network	Data Management tab -> IPAM tab -> Net Map
	Directory	Indicates a directory	Data Management tab -> File Distribution tab





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

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





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

Icon	Icon Name	Description	Tab/Table/Panel
	Backup	Backs up the configuration file and database	Grid tab-> Grid Manager tab -> Toolbar
	Restore	Restores the configuration file and database	Grid tab-> Grid Manager tab -> Toolbar
	bloxTools	Performs bloxTools services	Grid tab-> Grid Manager tab -> Toolbar
	Certificate	Creates, generates, uploads, or downloads an HTTPS certificate	Grid tab-> Grid Manager tab -> Toolbar
	Control	Restarts, reboots, or shuts down a member	Grid tab-> Grid Manager tab -> Members tab -> <i>member</i> -> Toolbar
	Manage Services	Manages member services	Grid tab-> Grid Manager tab -> Members tab -> <i>member</i>
	Syslog	Displays the syslog file	Grid tab-> Grid Manager tab -> Members tab -> <i>member</i> -> Toolbar
	Traffic Capture	Captures the traffic report on a member	Grid tab-> Grid Manager tab -> Members tab -> <i>member</i> -> Toolbar



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

Icon	Icon Name	Description	Tab/Table/Panel
	Execute Now	Executes a scheduled task immediately	Administration tab -> Scheduling tab -> Toolbar
	Overlap	Shows overlapping permissions	Administration tab -> Permissions tab
	Reschedule	Reschedules a task	Administration tab -> Scheduling tab -> Toolbar
	Schedule Delete	Schedules a deletion for a task	Administration tab -> Scheduling tab -> Toolbar




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

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

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

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

The following icons appear in Multi-Grid Manager:

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



Appendix C Guidance Documentation Supplement

Common Criteria provides an independent and objective evaluation of the security of Information Technology (IT) products. It gives assurance that the product satisfies a set of internationally recognized security standards.

This document provides additional guidance on the secure installation of the Target of Evaluation (TOE) for Common Criteria Evaluation Assurance Level (EAL) 2. The TOE is the TrinziC 810 and 820, TrinziC 1410 and 1420, TrinziC 2210 and 2220, and IB-4010 and 4030 with NIOS version 6.3 (hereafter referred to as Infoblox appliances), which are network appliances that provide delivery of IP network services and management.

To ensure that your appliance is Common Criteria compliant, make sure that your hardware and software settings match the evaluated configuration that was certified for Common Criteria.

This document provides clarifications and changes to the Infoblox Administrator Guide and Infoblox CLI Guide, and should be used as the guiding document for installation of the TOE in the Common Criteria evaluated configuration.

This appendix contains the following sections:

- [Pre-Requisites](#) on page 1286
- [Verifying the Hardware](#) on page 1286
- [Security Guidelines](#) on page 1286
 - [Installation and Configuration](#) on page 1286
- [Administration](#) on page 1287
 - [Setting Password Restrictions for Local Admins](#) on page 1287
- [Enabling/Disabling Common Criteria Mode](#) on page 1288
 - [Using the CLI](#) on page 1288
- [Licenses and Services](#) on page 1289
- [WebUI Settings](#) on page 1290
 - [Creating a Login Banner](#) on page 1290
 - [Modifying the Session Timeout Setting](#) on page 1290
 - [Managing Certificates](#) on page 1290
- [DNS](#) on page 1290
 - [DNSSEC](#) on page 1291
- [Backing Up and Restoring the Database](#) on page 1291
- [Audit Log](#) on page 1292
- [Syslog](#) on page 1296

PRE-REQUISITES

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

- A Trinzic 810 and 820, Trinzic 1410 and 1420, Trinzic 2210 and 2220, and IB-4010 and 4030 running NIOS version 6.3.
- A management station or computer from which you configure and manage the NIOS appliance. See [Management System Requirements](#) on page 46 for the system and browser requirements.
- The IP address of the appliance on your network.

VERIFYING THE HARDWARE

To verify the secure delivery of the hardware:

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

SECURITY GUIDELINES

Following are security assumptions to ensure that the TOE is administered in a secure manner after it is delivered:

- The environment ensures the physical security of the TOE, commensurate with its value and the value of the data that it contains.
- Administrators are non-hostile, properly trained and trusted to apply all administrator guidance.
- Administrators will take appropriate measures to prevent unauthorized individuals from accessing the TOE.

Installation and Configuration

To ensure the security of the installation and configuration of the TOE:

- Administrators must install the appliance according to the procedures in the installation guides.
- The TOE contains an option for upgrading the system. This is available only for security administrators. The security administrator will be able to upgrade to a validated release package only. The security administrator can verify the TOE by the version number included in the file name as well as through the administrative interface before and after the upgrade.

When upgrading, ensure that the .bin2 file is uploaded, and not the .bin file. Refer to the Release Notes of the NIOS version to which the TOE is upgrading for additional upgrade instructions.

- Users' access to the TOE is controlled by security mechanisms and unauthorized users are denied access to the TOE. For more information, see [Administration](#) on page 1287.
- The TOE provides external authentication mechanisms for remote users using SSL with Active Directory. For more information, see [Authenticating Admins Using Active Directory](#) on page 177.

ADMINISTRATION

A user must have an admin account to log in to the TOE. Each admin account belongs to an admin group, which contains roles and permissions that determine the tasks a user can perform.

The TOE provides a default superuser admin group, called **admin-group**, with one superuser administrator, **admin**. The default superuser admin can log in to the TOE, using the default user name **admin** and password **infoblox**. Superuser admins are the security admins and have full access and control of all the operations of a TOE. Note that you must change the default user name and password of the default superuser admin to prevent unauthorized access to the TOE.

Only superusers can do the following:

- Create admin accounts and groups. For more information, see [Chapter 4, Managing Administrators](#), on page 149.
- Set password parameters. For more information, see [Managing Passwords](#) on page 170.
- Create the login banner. For more information, see [Creating a Login Banner](#) on page 1290.
- Set the session timeout. For more information, see [Modifying the Session Timeout Setting](#) on page 1290.

Limited-access admin groups provide their members with read-only or read/write access to specific resources. These admin groups can access the appliance through the GUI, API, or both. They cannot access the appliance through the console. In addition, limited-access admins are not allowed to perform the following tasks:

- Download the support bundle.
- Enable SNMP on Grid members.
- Upload files that are larger than 100 MB. If the file size is greater than the maximum size allowed, the **Upload** dialog box closes and an error message is displayed in the feedback panel. The attempt to upload a file that exceeded the maximum will be logged to syslog. non-superusers only are able to upload files for file distribution and do CSV import

Setting Password Restrictions for Local Admins

All admins are required to enter a username and password when they log in to Grid Manager or the CLI. The password is always obscured when an admin logs in. The TOE defaults to locking out the user after three consecutive failed logins.

A superuser must define a password policy that is consistent with the security policy of the organization. The password policy specifies the minimum password length and character types, such as lowercase or uppercase characters, that are allowed in the password. In addition, the policy specifies the number of required character changes from the previous password, whether passwords expire and their duration. Additionally, you can require admins to change their passwords when they first log in or after their passwords are reset. For information about defining the password policy, see [Managing Passwords](#) on page 170.

Local admins must change their passwords according to the defined password policy. A password can be changed as follows:

- By the local admin in the User Profile page. For more information, see [Changing the Password and Email Address](#) on page 50.
- By the local admin when a password expires or when the admin first logs in. Note that this applies to logging in to the CLI or WebUI.
- By a superuser admin.

ENABLING/DISABLING COMMON CRITERIA MODE

Note: Infoblox recommends that you do not change the Common Criteria setting of a Grid that is in a production environment.

Before you enable Common Criteria mode, you must reset a NIOS appliance to its original factory settings. This removes the database, network settings, logs, and configuration files. Then, it reboots with its factory settings, which are the default user name and password, and default network settings. If you do not reset the appliance to its original factory settings, the appliance will not be Common Criteria compliant, even if you enable Common Criteria mode.

To reset the NIOS appliance to its factory settings:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command:

reset all

You can enable and disable Common Criteria mode from the Infoblox CLI only. In a Grid, you enable or clear Common Criteria Mode on the Grid Master only. After you execute the CLI command, the setting is propagated to all Grid members. To change the setting on an HA Grid Master, you must manually shut down the passive node and then change the setting on the active node.

Do the following to set Common Criteria mode on the appliance:

1. Log in to the Infoblox CLI. After executing the **reset all** command, you can log in to the TOE only by using the default superuser admin name **admin** and password **infoblox**.
2. Type the following command:

set cc_mode

The TOE reboots and goes through boot time self tests. If the test fails, the TOE goes into a loop and displays an error message on the serial console and the LCD. Otherwise, it displays the Login prompt after the self tests.

To clear Common Criteria mode on an appliance, log in to the Infoblox CLI and execute the command: **reset all**.

Using the CLI

Only superusers can access the CLI. To ensure security, access to the CLI is permitted through a direct console connection only. Note that activating the option **Enable Remote Console Access** in the Grid or Member Properties editor will result in a non-compliant system.

To access the Infoblox CLI through the console port:

1. Connect a serial cable from the console port on your management system to the console port on the appliance. The appliance has a male DB-9 console port on its front panel.
2. Use the following connection settings to launch an emulation session through a serial terminal emulation program such as Hilgraeve Hyperterminal® (provided with the Windows® operating systems):
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: Xon/Xoff
3. Use the following default user name and password to log in to the Infoblox appliance:

admin

infoblox

Note: After you log in, change the default user name and password of the default superuser admin to prevent unauthorized access to the TOE. For more information on changing passwords, see [Changing the Password and Email Address](#) on page 50.

For more information about the Infoblox CLI, refer to the *Infoblox CLI Guide*. By default, the **set network** and **set membership** commands display the Grid shared secret. In Common Criteria mode, these commands display **(not shown)** in the Grid Shared Secret field.

LICENSES AND SERVICES

The TOE does not have general computing capabilities, other than the services required for the operation, administration and support of the TOE. In the evaluated configuration, the TOE has only the following licenses installed: DNS, DHCP, Grid, Microsoft Management, DNS Cache Acceleration and Query Redirection. It does not have the following licenses installed: IF-MAP service, IF-MAP Federation, and Multi-Grid Management. For more information about licenses, see [Managing Licenses](#) on page 377.

The following services are disabled by default in the Common Criteria evaluated configuration and no claims are made regarding their function:

- bloxTools
- MGM (Multi-Grid Management)
- HSM (Hardware Security Module) Signing
- Support access
- Remote console access
- Remote authentication using RADIUS and TACACS+

Installing additional licenses or enabling any of the listed services may result in a non-compliant system.

WEBUI SETTINGS

This section describes the properties that you can set to ensure the security of the Grid Manager web interface

Creating a Login Banner

Before establishing a user session via the WebUI, the TOE displays an initial banner regarding unauthorized use. The message is displayed before the session is established. You can change this message to your organization's specific advisory notice and warning message regarding unauthorized use of the system. For information about defining the login banner, see [Creating a Login Banner](#) on page 49.

Modifying the Session Timeout Setting

You can set the length of idle time before an administrative session to the WebUI times out. The default timeout value is 600 seconds (10 minutes). If an admin does not interact with the application for the specified time, the TOE displays a message that a timeout has occurred. The admin is then required to log back in to Grid Manager. For information about setting the session timeout, see [Modifying the Session Timeout Setting](#) on page 344.

Managing Certificates

The TOE generates a self-signed certificate when it first starts. Because the default certificate is self-signed, your browser does not have a trusted CA certificate or a cached NIOS appliance server certificate (saved from an earlier connection) to authenticate the NIOS appliance certificate. Also, the hostname in the default certificate is `www.infoblox.com`, which is unlikely to match the hostname of your NIOS appliance. Consequently, a message appears warning that the certificate is not from a trusted certifying authority and that the hostname on the certificate is either invalid or does not match the name of the site that sent the certificate. To eliminate certificate warnings, you can replace the default self-signed certificate with a different certificate.

After the initial login, you can do one of the following:

- Generate another self-signed certificate with the correct hostname and save it to the certificate store of your browser.
- Request a CA-signed certificate with the correct hostname by generating a Certificate Signing Request (CSR) and sending it to your trusted Certificate Authority (CA). Then when you receive the certificate from the CA, import it to the appliance.

For information about these tasks, see [Managing Certificates](#) on page 53.

For Common Criteria compliance, superusers must not use CSRs or certificates with keys smaller than 2048 bits. Limited access users are not allowed to upload a certificate with a key that is smaller than 2048 bits, or create a certificate signing request or self-signed certificate with a key size that is smaller than 2048 bits.

DNS

The TOE provides DNS service. There are two basic methods used to protect DNS communication: TSIG and GSS-TSIG. The TSIG (transaction signature) method signs communications using either HMAC-MD5 or HMAC-SHA256. Both end points must be configured with the key. The GSS-TSIG method (based on the GSS API) uses a Kerberos server to retrieve the key, and is only available in Microsoft environments.

When you configure the TOE to use TSIG and GSS-TSIG keys, you must select HMAC-SHA256 as the key algorithm.

For information about using TSIG keys to ensure security in several DNS operations, see the following:

- To control access to DNS views. For more information, see [Defining Match Clients Lists](#) on page 605 and [Defining a Match Destinations List](#) on page 607.

- To control to which recursive and non-recursive queriers the TOE is allowed to respond. For more information, see [Specifying Queriers](#) on page 570 and [Enabling Recursion](#) on page 571
- To authenticate zone transfer requests and replies. For more information, see [Configuring Zone Transfers](#) on page 584.
- To authenticate and verify dynamic DNS updates from DHCP servers. For more information, see [Enabling DNS Servers to Accept DDNS Updates](#) on page 706.
- When a secondary DNS server receives DDNS updates, it must forward the updates to the primary server because it cannot update zone data itself. To specify the source of DDNS updates. For more information, see [Forwarding Updates](#) on page 707.

For information about using GSS-TSIG, see [About GSS-TSIG](#) on page 710.

DNSSEC

DNSSEC (DNS Security Extensions) provides mechanisms for authenticating the source of DNS data and ensuring its integrity. For more information, see [Chapter 21, DNSSEC](#), on page 733. When an authoritative name server digitally signs a zone, it typically generates two key pairs, a zone-signing key (ZSK) pair and a key-signing key (KSK) pair. You select the cryptographic algorithm that the Grid Master uses when it generates the KSK and ZSK. The TOE supports only the following algorithms:

- RSA/SHA-1
- RSA/SHA-1/NSEC-3
- RSA/SHA-256
- RSA/SHA-512

For information about selecting the cryptographic algorithm for the keys, see [Setting DNSSEC Parameters](#) on page 743.

BACKING UP AND RESTORING THE DATABASE

You must log in with a superuser account to back up files. The administrator must back up system files to the local appliance.

You can restore a backup up file to an appliance running the same NIOS version as that of the appliance from which the backup file originates. You can also restore a backup file from an appliance running a NIOS version to an appliance running a later NIOS version as long as the upgrade from the earlier NIOS version to the later version is supported. Note that if you need to restore a backup file to an appliance, ensure that the backup file that you are restoring is from an appliance that was Common Criteria compliant as well.

For more information about backing up and restoring the database, see [Backing Up and Restoring Configuration Files](#) on page 423.

AUDIT LOG

The audit log contains a record of all TOE administrative activities. The stored audit records in the audit trail are protected from unauthorized modifications and deletion. For more information about the audit log, see [Using the Audit Log](#) on page 1018.

Following are the events that are logged and examples of their corresponding audit log messages:

Identification and Authentication

Event: Invalid password when logging in to the WebUI.

Message: "2011-10-19 14:02:32.750Z [admin]: Login_Denied - - to=Serial\040Console apparently_via=Direct error=invalid\040login\040or\040password"

Event: Number of attempts exceeds the limit when logging in to the WebUI.

Message: "2011-10-19 14:05:23.217Z [admin]: Login_Denied - - to=Serial\040Console apparently_via=Direct error=failed\040logins\040exceed\040limit"

Event: Invalid password when logging in to the CLI.

Message: "2011-10-19 14:02:32.750Z [admin]: Login_Denied - - to=Serial\040Console apparently_via=Direct error=invalid\040login\040or\040password"

Event: Number of attempts exceeds the limit when logging in to the CLI.

Message: "2011-10-19 14:05:23.217Z [admin]: Login_Denied - - to=Serial\040Console apparently_via=Direct error=failed\040logins\040exceed\040limit"

Event: Enable Common Criteria mode:

Message: 2011-10-19 19:48:37.299Z [admin]: Login_Allowed - - to=Serial\040Console apparently_via=Direct auth=Local group=.admin-group

Message: 2011-10-19 19:48:48.705Z [admin]: Called - set_cc_mode: Args cc_mode_enabled="true"

Event: Disable Common Criteria mode:

Message: 2011-10-19 19:48:37.299Z [admin]: Login_Allowed - - to=Serial\040Console apparently_via=Direct auth=Local group=.admin-group

Message: 2011-10-19 19:48:48.705Z [admin]: Called - set_cc_mode: Args cc_mode_enabled="false"

Event: Login successful

Message: 2011-10-19 19:48:48.706Z [USER\040admin]: rebooted the system

2011-11-01 17:09:21.696Z [admin]: Login_Allowed - - to=Serial\040Console apparently_via=Direct auth=Local group=.admin-group

Event: First login

Message: 2011-10-19 12:43:47.375Z [user]: First_Login - - to=AdminConnector ip=127.0.0.1 auth=LOCAL group=admin-group apparently_via=GUI first login

Event: Password expired

Message: 2011-10-20 13:17:29.257Z [user]: Password_Expired - - to=AdminConnector ip=127.0.0.1 auth=LOCAL group=admin-group apparently_via=GUI

Event: Password was successfully reset.

Message: 2011-10-19 12:44:45.962Z [user]: Password_Reset - - to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI

Event: New password did not conform to the rule.

Message: 2011-10-19 13:07:33.343Z [user]: Password_Reset_Error - - to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI

Quotas

Event: Upload file limit reached.

Message: user manojk-vm httpd[]: err User {0} tried to upload the file. File {1} with size 272629904 kBytes is greater than maximum size allowed. Maximum size is 102400 kBytes.

LDAP

Event: Establishment of session

Message: 2011-10-27T07:50:59-04:00 user epbyminw0065t2 python[]: notice Connection established:success

Event: Failure to establish a session

Message: 2011-10-27T07:50:38-04:00 user epbyminw0065t2 python[]: err 10.6.11.249: AD user authentication timed out

Message: 2011-10-27T07:51:02-04:00 user epbyminw0065t2 python[]: err Connection timed out

Event: Crypto Failure (Type and name of crypto algorithm that failed cannot be logged, since openldap uses SSL/TLS protocol functions from OpenSSL and did not use crypto functions directly.)

Message: 2011-10-27T07:51:00-04:00 user epbyminw0065t2 python[]: err SSL handshake failed.

Message: 2011-10-27T07:51:02-04:00 user epbyminw0065t2 python[]: err SSL handshake failed. Cannot verify server certificate.

GSS-TSIG

Event: Invalid size specified for algorithm HMAC-SHA256

Message: 2011-10-19T17:57:12-04:00 user EPBYMINW2856 httpd[]: err TSIG key generation failure: Size 512 can not be used with algorithm HMAC-SHA256

Event: dnssec-keygen error

Message: 2011-10-19T17:57:13-04:00 user EPBYMINW2856 httpd[]: err Failed to execute dnssec-keygen: errno = 2, keylen = 256, algnam = HMAC-MD5

Message: 2011-10-19T17:57:13-04:00 user EPBYMINW2856 httpd[]: err Unable to generate TSIG key

Event: Invalid algorithm specified in Common Criteria mode

Message: 2011-10-19T18:12:22-04:00 user EPBYMINW2856 httpd[]: err TSIG key (keylen = 256, algnam = HMAC-MD5) generation error : Only HMAC-SHA256 available in CC mode.

Event: Algorithm restriction

Message: Only AES128_CTS_HMAC_SHA1_96 or AES256_CTS_HMAC_SHA1_96 algorithms are allowed in CC mode. Current algorithm is DES_CBC_CRC.

TSIG CSV Import/Export

Event: Import error (TSIG algorithm is not allowed in Common Criteria mode)

Message: [2011/10/20 09:38:42.496] (24473 /usr/bin/python) /infoblox/common/lib/python/infoblox/one/csv_import_function.py:601 write_to_error_file(): Import Error: authzone,zone.com,FORWARD,,,,,,,,False,False,False,,1.2.3.4/1.2.3.4/False/False/True/ext_sec_key/ut29ROLajwty6a%2Fhsgg0wA==,infoblox.localdomain,False,,,,,,,,,2,,default,Authoritative-Line 2: Insertion aborted due to IBDataError?: IB.Data:TSIG algorithm used for TSIG key name 'ext_sec_key' is not allowed in CC mode.

“set” commands

Message: 2011-10-19 13:14:04.030Z [admin]: Called - set_snmptrap: Args variable="sysName.0", address="10.120.20.31"

Message: 2011-10-19 13:16:16.545Z [admin]: Called - set_scheduled: Args task_restarts="0 from 60"

Message: 2011-10-19 13:17:19.391Z [admin]: Called - set_mld_version_1: MLD version set to 1

Message: 2011-10-19 13:18:28.171Z [admin]: updated grid security

Message: 2011-10-19 13:18:28.171Z [admin]: Called - set_support_access: Args support_access="true from false"

Message: 2011-10-19 13:19:46.668Z [admin]: updated grid security

Message: 2011-10-19 13:19:46.669Z [admin]: Called - set_session_timeout: Args session_timeout="650 from 600"

Message: 2011-10-19 13:23:11.596Z [admin]: Called - set_phonehome: Args phonehome_disabled="true from false"

Message: 2011-10-19 13:24:02.372Z [admin]: updated grid security

Message: 2011-10-19 13:24:02.372Z [admin]: Called - set_remote_console: Args remote_console="true from false"

Message: 2011-10-19 13:25:31.696Z [admin]: updated grid security

Message: 2011-10-19 13:25:31.704Z [admin]: Called - set_security: Args address="10.120.20.31",netmask="255.255.255.0"

Message: 2011-10-19 13:26:12.673Z [admin]: Called - set_safemode

Message: 2011-10-19 13:28:12.302Z [admin]: Called - set_prompt: Args prompt=ip

Message: 2011-10-19 13:30:22.221Z [admin]: Called - set BGP: Args log_level="debugging"

Message: 2011-10-19 13:31:20.142Z [admin]: Called - set OSPF: Args log_level="informational"

Message: 2011-10-19 13:32:10.319Z [admin]: Called - set_nosafemode

Message: 2011-10-19 13:38:42.998Z [admin]: Called - set_network: Args ip_address="10.120.20.34 from 10.120.20.31",netmask="255.255.255.0 from 255.255.255.0",gateway_address="10.120.20.1 from 10.120.20.1"

Message: 2011-10-19 13:41:56.178Z [admin]: Called - set_ip_rate_limit: Args ip_rate_limit="on from off"

Message: 2011-10-19 13:43:42.828Z [admin]: Called - set_monitor_dns_alert: Args dns_alert="on from off"

Message: 2011-10-19 13:46:34.647Z [admin]: updated physical node 0

Message: 2011-10-19 13:46:34.648Z [admin]: Called - set_interface: Args interface="LAN", speed="100M", duplex="half"

Message: 2011-10-19 13:48:03.066Z [admin]: Called - set_dns: Args dns="flush all "

Message: 2011-10-19 13:49:35.527Z [admin]: Called - set_debug: Args all="on from off"

Message: 2011-10-19 09:53:53.595Z [admin]: Called - set_ibtrap: Args ibtrap="DNS", snmp="true", email="true"

Message: 2011-10-19 09:57:00.747Z [admin]: Called - set_thresholdtrap: Args thresholdtrap="CpuUsage", trigger="60", reset="50"

Message: 2011-10-19 10:32:50.183Z [admin]: Called - set_maintenancemode: Args maintenancemode="on from off"

Message: 2011-10-19 14:05:20.132Z [admin]: Called - set_dhcp_expert_mode: Args dhcp_expert_mode="true from false"

Message: 2011-10-19 14:07:02.082Z [admin]: Called - set_dhcp_release_delay: Args delay_time=40 secs

Message: 2011-10-19 14:08:16.395Z [admin]: Called - set_lower_case_ptr_dname: Args grid="false from true"

Message: 2011-10-19 14:09:24.285Z [admin]: Called - set_gsstskey_expiration_time: Args gsstskey_expiration_time="3000 from 3600"

Message: 2011-10-19 14:10:19.906Z [admin]: Called - set_named_worker_threads: Args named_worker_threads="20 from 0"

Message: 2011-10-19 14:11:04.731Z [admin]: Called set_recursion_log_interval: Args recursion_log_interval="60"

Message: 2011-10-19 14:11:54.147Z [admin]: Called - set_fixed_address_obeyes_mac_filter: Args grid="false from false"

Message: 2011-10-19 14:12:57.589Z [admin]: Called - set_transfers_out: Args grid="use_default from NULL"

Message: 2011-10-19 14:14:12.170Z [admin]: Called - set_partial_replication: Args partial_replication="off from on"

Message: 2011-10-19 14:15:33.978Z [admin]: Called - set_rep_queue_ixfr_limit: Args rep_queue_ixfr_limit="60 from 1000"

Message: 2011-10-19 14:16:16.797Z [admin]: Called - set_watchdog: Args watchdog_enabled="true from false"

Message: 2011-10-19 14:17:14.605Z [admin]: Called - set_fsck

Message: 2011-10-19 14:19:25.282Z [admin]: Called - set_host_consistency_check: Args host_consistency_check="on from off"

Message: 2011-10-19 14:21:00.202Z [admin]: Called - set_internal_apache_http_port: Args internal_apache_http_port="2000 from 9000"

Message: 2011-10-19 14:22:18.682Z [admin]: Called - set_internal_jetty_http_port: Args internal_apache_http_port="6060 from 8080"

Message: 2011-10-19 14:23:32.571Z [admin]: Called - set_device_status_log_interval: Args grid="400 from 10800"

Message: 2011-10-19 14:25:58.704Z [admin]: Called - set_always_ret_nxdomain_for_fmz_ptr: Args always_ret_nxdomain_for_fmz_ptr="true from false"

Message: 2011-10-19 14:28:18.046Z [admin]: Called - set_debug_tools: Args debug_tools="db_binary_dump"

Message: 2011-10-19 14:29:06.511Z [admin]: Called - set_dns_autogen: Args dns_auto_gen="check"

Message: 2011-10-19 14:30:54.628Z [admin]: Called - set_named_rcv_sock_buf_size: Args udp_so_rcvbuf="122 from (null)"

CLI Top Level Commands

Message: 2011-10-19 10:33:29.664Z [admin]: Called - delete_cores_all

Message: 2011-10-19 10:38:12.356Z [admin]: Called - delete_cores: Args filename="core.8295.gz"

Message: 2011-10-19 10:58:28.064Z [admin]: Called - delete_backup_all

Message: 2011-10-19 11:00:17.917Z [admin]: Called - delete_backup: Args filename="BACKUP_6.bkp"

Message: 2011-10-19 12:41:47.707Z [admin]: Called - rotate_log: Args log="syslog"

Message: 2011-10-19 12:58:11.738Z [admin]: Called - rotate_log: Args log="audit"

Message: 2011-10-19 12:58:11.738Z [USER\040admin]: rotated the previous audit log to audit.log.0.gz

Message: 2011-10-19 13:51:36.982Z [admin]: Called - reset_database

Message: 2011-10-19 13:54:14.023Z [admin]: Called - debug_webui_restart

Message: 2011-10-19 13:57:39.407Z [USER\040admin]: rebooted the system

Message: 2011-10-19 14:03:41.124Z [admin]: Called - delete_file: Args groupname="bloxtools", filename="/storage/web-portal/udata/logs/access.log"

CLI Emergency Commands

Message: 2011-10-19 14:32:31.927Z [Emergency\040User]: Called - set_safemode

Message: 2011-10-19 14:33:23.591Z [Emergency\040User]: Called - set_nosafemode

Message: 2011-10-19 14:33:41.286Z [Emergency\040User]: Called set_repsafe_mode: Args repsafe_mode = on

Message: 2011-10-19 14:34:47.321Z [Emergency\040User]: Called - set_weak

Message: 2011-10-19 14:35:25.969Z [Emergency\040User]: Called - set_fsck

Message: 2011-10-19 14:35:46.604Z [Emergency\040User]: Called - set_watchdog: Args watchdog_enabled="true from true"

Message: 2011-10-19 14:41:13.727Z [Emergency\040User]: Called - reset_database

SYSLOG

NIOS appliances generate syslog messages that you can view through the Syslog viewer and download to a directory on your management station. For more information about syslog, see [Using a Syslog Server](#) on page 1012.

Following are the events that are logged and examples of their corresponding syslog messages:

Establishment/Termination of an HTTPS Session

Event: Generation of RSA key failed.

Message: "Oct 19 09:15:01 EPBYMINW0065T1 httpd[2115]: cryptographic key generation failed"

Event: Session is terminated.

Message: "Oct 19 09:15:01 EPBYMINW0065T1 httpd[2115]: Session terminated (remote address: 10.6.11.249)"

Event: Failed to establish a session.

Message: "Oct 19 08:50:21 EPBYMINW0065T1 httpd[2314]: Failed to establish a session (remote address: 10.6.11.249), error 1115 (SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)"

Event: Session is established.

Message: "Oct 19 08:54:42 EPBYMINW0065T1 httpd[2314]: Session has been established (remote address: 10.6.11.249)"

Establishment/Termination of an SSH Session

Event: Establishment of session.

Message: sshd[15926]: info Session has been established (remote ip - 127.0.0.1 , user - root).

Event: Termination of session.

Message: sshd[14727]: info Received disconnect from 127.0.0.1: 11: disconnected by user

Event: Failure to establish a session.

Message: sshd[17671]: info Failed password for root from 127.0.0.1 port 50566 ssh2

Message: sshd[18358]: info Connection closed by 127.0.0.1

Establishment/Termination of a TLS Session

Event: Generation of RSA key failed.

Message: "Oct 19 08:38:08 EPBYMINW0065T1 openvpn[1415]: cryptographic key generation failed"

Event: Session has been established.

Message: "Oct 19 08:38:08 EPBYMINW0065T1 openvpn[1552]: Session has been established (remote address: 10.6.11.249)"

Event: HMAC failure:

Message: "Oct 19 08:41:01 EPBYMINW0065T1 openvpn[1567]: cryptographic key generation failed: HMAC"

Event: Signing failure (constructed message, it is not trivial to obtain it into the syslog).

Message: "Oct 19 08:45:01 EPBYMINW0065T1 openvpn[1582]: cryptographic operation failed: signature"

Event: Encryption failure.

Message: "Oct 19 08:46:41 EPBYMINW0065T1 openvpn[1612]: cryptographic operation failed: encryption"

Event: Decryption failure.

Message: "Oct 19 08:46:41 EPBYMINW0065T1 openvpn[1612]: cryptographic operation failed: decryption"

Event: Session was not established.

Message: "Oct 19 08:50:21 EPBYMINW0065T1 openvpn[1701]: Failed to establish a session (remote address: 10.6.11.249), error 1115 (SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)"

Event: Packet was not verified.

Message: "Oct 19 08:55:25 EPBYMINW0065T1 openvpn[1815]: Packet verification fails (remote address: 10.6.11.249)"

Random Number Generation Process

[2011/10/19 10:13:46.282] (26360 /infoblox/one/bin/ib_prngd_control) : ib_prngd daemon is not running while CC mode is enabled

[2011/10/19 10:13:46.324] (26368 /infoblox/one/bin/ib_prngd) main.c:202 main(): ib_prngd daemon starting up...

[2011/10/19 10:13:46.700] (26368 /infoblox/one/bin/ib_prngd) main.c:214 main(): Setting FIPS mode OK

[2011/10/19 10:13:48.400] (26368 /infoblox/one/bin/ib_prngd) main.c:214 main(): Setting FIPS mode FAILED

[2011/10/19 10:13:46.700] (26368 /infoblox/one/bin/ib_prngd) main.c:125 rename_rnd_dev(): Moving /dev/random to /dev/random_backup OK

[2011/10/19 10:13:46.700] (26368 /infoblox/one/bin/ib_prngd) main.c:127 rename_rnd_dev(): Moving /dev/urandom to /dev/urandom_backup OK

[2011/10/19 10:13:46.700] (26368 /infoblox/one/bin/ib_prngd) main.c:234 main(): Creating FIFO /dev/ib_random OK

[2011/10/19 10:13:46.700] (26368 /infoblox/one/bin/ib_prngd) main.c:158 symlink_rnd_dev(): Symlinking /dev/random to /dev/ib_random OK

[2011/10/19 10:13:46.700] (26368 /infoblox/one/bin/ib_prngd) main.c:160 symlink_rnd_dev(): Symlinking /dev/urandom to /dev/ib_random OK

[TIME NOT KNOWN] (26368) main.c:signal_handler(): ib_prngd received SIGTERM signal....exiting.

[TIME NOT KNOWN] (26368) main.c:signal_handler(): ib_prngd received SIGINT signal....exiting.

[TIME NOT KNOWN] (26368) main.c:signal_handler(): ib_prngd received SIGQUIT signal....exiting.

[TIME NOT KNOWN] (26368) main.c:signal_handler(): ib_prngd received an unknown signal....exiting.

[2011/10/19 10:13:49.205] (26368 /infoblox/one/bin/ib_prngd) main.c:135 rename_rnd_dev(): Renaming /dev/random back OK

[2011/10/19 10:13:49.205] (26368 /infoblox/one/bin/ib_prngd) main.c:141 rename_rnd_dev(): Renaming /dev/urandom back OK

[2011/10/19 10:13:49.205] (26368 /infoblox/one/bin/ib_prngd) main.c:255 main(): Removing custom FIFO /dev/ib_random OK

[2011/10/19 10:13:49.205] (26368 /infoblox/one/bin/ib_prngd) main.c:255 main(): Removing custom FIFO /dev/ib_random FAILED

[2011/10/19 10:13:49.205] (26368 /infoblox/one/bin/ib_prngd) main.c:141 rename_rnd_dev(): Renaming /dev/urandom back FAILED

[2011/10/19 10:13:49.205] (26368 /infoblox/one/bin/ib_prngd) main.c:135 rename_rnd_dev(): Renaming /dev/random back FAILED

[2011/10/19 10:25:22.931] (26557 /infoblox/one/bin/ib_prngd) main.c:189 main(): Error!
/infoblox/one/bin/ib_prngd is already running
[2011/10/19 10:26:58.107] (26560 /infoblox/one/bin/ib_prngd) main.c:52 self_test(): OpenSSL FIPS mode
functionality self test OK
[2011/10/19 10:26:58.107] (26560 /infoblox/one/bin/ib_prngd) main.c:52 self_test(): OpenSSL FIPS mode
functionality self test FAILED

Failures on Invoking Functionality

Event: Invalid size specified for algorithm HMAC-SHA256.

Message: 2011-10-19T17:57:12-04:00 user EPBYMINW2856 httpd[]: err TSIG key generation failure: Size 512 can not be used with algorithm HMAC-SHA256

Event: dnssec-keygen erro.

Message: 2011-10-19T17:57:13-04:00 user EPBYMINW2856 httpd[]: err Failed to execute dnssec-keygen: errno = 2, keylen = 256, alname = HMAC-MD5

Message: 2011-10-19T17:57:13-04:00 user EPBYMINW2856 httpd[]: err Unable to generate TSIG key

Event: Invalid algorithm specified in Common Criteria mode.

Message: 2011-10-19T18:12:22-04:00 user EPBYMINW2856 httpd[]: err TSIG key (keylen = 256, alname = HMAC-MD5) generation error : Only HMAC-SHA256 available in CC mode.

Open VPN

Event: Generation of RSA key failed

Message: Oct 19 08:38:08 EPBYMINW0065T1? openvpn[1415]: cryptographic key generation failed

Event: Session has been established

Message: Oct 19 08:38:08 EPBYMINW0065T1? openvpn[1552]: Session has been established (remote address: 10.6.11.249)

Event: HMAC failure

Message: Oct 19 08:41:01 EPBYMINW0065T1? openvpn[1567]: cryptographic key generation failed: HMAC

Event: Signing failure

Message: Oct 19 08:45:01 EPBYMINW0065T1? openvpn[1582]: cryptographic operation failed: signature

Event: Encryption failure

Message: Oct 19 08:46:41 EPBYMINW0065T1? openvpn[1612]: cryptographic operation failed: encryption

Event: Decryption failure

Message: Oct 19 08:46:41 EPBYMINW0065T1? openvpn[1612]: cryptographic operation failed: decryption

Event: Session was not established

Message: Oct 19 08:50:21 EPBYMINW0065T1? openvpn[1701]: Failed to establish a session (remote address: 10.6.11.249), error 1115 (SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)

Event: Packet was not verified

Message: Oct 19 08:55:25 EPBYMINW0065T1? openvpn[1815]: Packet verification fails (remote address: 10.6.11.249)

HTTPS

Event: Generation of RSA key failed

Message: Oct 19 09:15:01 EPBYMINW0065T1? httpd[2115]: cryptographic key generation failed

Event: Session is terminated

Message: Oct 19 09:15:01 EPBYMINW0065T1? httpd[2115]: Session terminated (remote address: 10.6.11.249)

Event: Failed to establish a session

Message: Oct 19 08:50:21 EPBYMINW0065T1? httpd[2314]: Failed to establish a session (remote address: 10.6.11.249), error 1115 (SSL routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)

Event: Session is established

Message: Oct 19 08:54:42 EPBYMINW0065T1? httpd[2314]: Session has been established (remote address: 10.6.11.249)

Event: HMAC failure

Message: Oct 19 08:55:56 EPBYMINW0065T1? httpd[2356]: cryptographic key generation failed: HMAC

SSH

Event: Establishment of session

Message: sshd[15926]: info Session has been established (remote ip - 127.0.0.1 , user - root).

Event: Termination of session

Message: sshd[14727]: info Received disconnect from 127.0.0.1: 11: disconnected by user

Event: Failure to establish a session

Message: sshd[17671]: info Failed password for root from 127.0.0.1 port 50566 ssh2

Message: sshd[18358]: info Connection closed by 127.0.0.1

Message: sshd[18358]: fatal Unable to negotiate a key exchange method

Event: Crypto Failure

Message: sshd[18358]: fatal ssh_rsa_sign: sign failed

Message: sshd[18358]: fatal ssh_rsa_sign: RSA_sign failed: "open SSL error"

Message: sshd[18358]: fatal ssh_rsa_verify: remaining bytes in signature 10

Message: sshd[18358]: fatal ssh_rsa_verify: RSA modulus too small: 100 < minimum 1024 bits

Message: sshd[18358]: fatal cipher_encrypt: bad plaintext length 3

Message: sshd[18358]: fatal evp_crypt: EVP_Cipher failed

Message: sshd[18358]: fatal rsa_generate_private_key: key generation failed.

Message: sshd[18358]: fatal cipher_init: EVP_CipherInit: set key failed for aes-256

Event: Detection of modification

csshd[18358]: fatal Disconnecting (remote ip - 127.0.0.1): Packet corrupt

Message: sshd[18358]: fatal Corrupted MAC on input (remote ip - 127.0.0.1).

DNS

Message: 2011-10-18T13:37:33+00:00 daemon (none) named[4456]: err client 10.32.2.108#47160: request has invalid signature: TSIG sha256cc: tsig verify failure (BADKEY) 2011-10-18T13:37:33+00:00 daemon (none) named[4456]: err client 10.32.2.108#47160: request has invalid signature: TSIG sha256cc: tsig verify failure (BADKEY)

DHCP

Message: 2011-10-18T11:18:38+00:00 daemon (none) dhcpd[20440]: err No tsec for use with key sha128cc

Message: 2011-10-31T18:32:17+00:00 daemon (none) dhcpd[20440]: err Invalid operation in ddns code.

DNSSEC

Message: 2011-10-27T07:25:49-04:00 user EPBYMINW2994t1 python[]: info Generated DNSSEC KSK (algorithm=5, key tag=50101, key size=2048): '.' in view 'default'.

Message: 2011-10-27T07:25:56-04:00 user EPBYMINW2994t1 python[]: info Generated DNSSEC ZSK (algorithm=7, key tag=27674, key size=2048): '.' in view 'default'.

Message: 2011-10-27T07:25:56-04:00 user EPBYMINW2994t1 python[]: err DNSSEC key generation failed (rc = -1, command = /usr/sbin/dnssec-keygen -q -r /dev/urandom -a NSEC3RSASHA1? -b 20480 -n ZONE . , error = nssec-keygen: fatal: RSA key size 20480 out of range

Message: dnskey 'key_string' failed to sign data: out of memory

Upgrade

Message: 2011-10-26T12:33:30-04:00 user EPBYMINW2994t1 infoblox_crypt[]: err cryptographic operation failed: decryption

Message: 2011-10-26T12:34:33-04:00 user EPBYMINW2994t1 infoblox_crypt[]: err cryptographic operation failed: encryption

Message: 2011-10-26T12:35:53-04:00 user EPBYMINW2994t1 infoblox_crypt[]: err cryptographic operation failed: RSA verify signature

Message: 2011-10-26T12:38:56-04:00 user EPBYMINW2994t1 infoblox_crypt[]: err cryptographic operation failed: RSA signing

Quotas

Event: When the administration backend is overloaded by too much combined GUI and API traffic, a message like this is logged to syslog (it is not associated with any user).

Message: 2011-10-31T23:42:21+00:00 user (none) httpd[]: warning Too many administration connections

Event: Disk space limit was changed and is below the disk usage.

Message: 2011-11-02T00:24:54+00:00 user manojk-vm httpd[]: err Storage Limit has been lowered and usage now exceeds the limit, Usage: 150 MB, Limit :100 MB

Event: Disk space limit reached.

Message: 2011-11-02T00:24:54+00:00 user manojk-vm httpd[]: err Exceed the TFTP Storage limit, User name:user1, Used Storage:2048 B, File name :a.zip, File size :272629904 B, Limit :102400 B

Open SSL

Event: FIPS self test failed.

Message: FIPS routines:EVP_DigestInit_ex:fips selftest failed:digest.c:18:

Event: Tried to use non-FIPS algorithm in FIPS mode.

Message: 140576691959464:error:140A9129:SSL routines:SSL_CTX_new:only tls allowed in fips mode:ssl_lib.c:1527:

Message: 139852903503528:error:0A07C06E:dsa routines:func(124):reason(110):dsa_key.c:131:

Event: Used DES-CBC-SHA cipher suite in FIPS mode.

Message: 140418599392936:error:1410D0B9:SSL routines:SSL_CTX_set_cipher_list:no cipher match:ssl_lib.c:1282:

Event: Error setting digest MD5.

Message: 140403566474920:error:060800A0:digital envelope routines:EVP_DigestInit_ex:unknown cipher:digest.c:248:

Replay Detection

Event: OpenVPN

Message: Mon Oct 22 22:30:00 2007 us=939054 Authenticate/Decrypt packet error: bad packet ID (may be a replay): [#0 / time = (4196958004) Wed Nov 23 16:11:48 1966] silence this warning with --mute-replay-warnings, error_prefix, packet_id_net_print (&pin, true, &gc)

Event: OpenVPN

Message: Mon Oct 22 22:30:00 2007 ACK reliable_can_send is a replay : [1] 0

Event: HTTPS

Message: Mon Oct 22 22:30:00 2007 Digest: Warning possible replay attack: nonce-count check failed: 12345678 = 123456789

GSS-TSIG

Message: 2011-10-18T13:37:33+00:00 named[4456]: err signature invalid: message integrity

Message: 2011-10-18T14:32:22+00:00 named[4456]: err authentication failed for aes128-cts-hmac-sha1-96: unknown principal

Message: 2011-10-18T14:42:12+00:00 named[4456]: err signature failed to verify(1)

Message: 2011-10-18T14:45:54+00:00 named[4456]: err signature is in the future

User Login

Message: 2011-10-19T08:27:23-04:00 user spradhan-vm serial_console[]: info User admin set_repsafe_mode: On

Message: 2011-10-19T08:29:54-04:00 user spradhan-vm serial_console[]: info User admin set_repsafe_mode: Off

Message: 2011-10-19T08:38:02-04:00 user spradhan-vm serial_console[]: info audit has been truncated to approximately 2011-10-19T08:29:00-04:00

Message: 2011-10-19T08:41:47-04:00 user spradhan-vm serial_console[]: info syslog has been truncated to approximately 2011-10-19T08:41:00-04:00

File Rotation

Event: Audit log is rotated.

Message: 2011-11-01T18:23:00-07:00 user manojk-vm perl[18990]:info audit has been truncated to approximately 2011-11-01T18:23:00-07:00

Event: Syslog is rotated.

Message: 2011-11-01T18:23:00-07:00 user manojk-vm perl[18990]:info syslog has been truncated to approximately 2011-11-01T18:23:00-07:00

Zeroization

Event: Logged in case of error

Message: 2011-11-01T15:32:59-04:00 daemon manojk-vm ntpd[18990]:err Error erasing /storage/etc/ntp.keys using shred

First Login

Message: [2011/10/19 08:44:45.866] (32289 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): [user] First_Login
to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI

Password Expired

Message: [2011/10/20 09:17:29.257] (15750 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): [user] Password_Expired
to=AdminConnector ip=127.0.0.1 auth=LOCAL group=admin-group apparently_via=GUI

Password Reset

Message: [2011/10/19 08:44:45.962] (32289 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): [user] Password_Reset
to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI

Failed Password Reset

Message: [2011/10/19 09:07:33.343] (32526 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): [user] Password_Reset_Error
to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI



Appendix D Regular Expressions

SUPPORTED EXPRESSIONS FOR SEARCH PARAMETERS

Regular expressions are text strings that you use to describe search patterns. You can use the following special characters to define regular expressions for search parameters.

Special character	Purpose	Example	Meaning
()	Defines the scope and precedence of the operator	gr(a e)y	Matches “gray” or “grey”.
	Matches either the regular expression before or after the vertical bar	a c	Matches “a” or “c”.
.	Matches any single character	.at	Matches any text string ending with “at”, such as “hat”, “cat”, and “bat”.
*	Matches the previous regular expression zero or more times	a*bc	Matches zero or multiple occurrences of “a” followed by “bc”, such as “bc”, “abc”, “aabc”, “aaabc”, and so on.
+	Matches the previous regular expression one or more times	a+bc	Matches one or more occurrences of “a”, followed by “bc”, such as “abc”, “aabc”, “aaabc”, and so on.
?	Matches the previous regular expression zero or one time	a?bc	Matches zero or one occurrence of “a”, followed by “bc”, such as “bc” or “abc”.
^	Matches the beginning of a text string	^c	Matches any string beginning with “c”, such as “cat”.
\$	Matches the end of a text string	com\$	Matches any string ending with “com”, such as “Infoblox.com”.
[]	Matches any character specified in the brackets	[03] [abcd] [15a-d]	Matches “0” or “3”. Matches “a”, “b”, “c”, or “d”. Matches “1”, “5”, “a”, “b”, “c”, or “d”.

Special character	Purpose	Example	Meaning
[<i>m-n</i>]	Matches single characters contained in the specified range, including the start and end points	[0-3] [a-f]	Matches 0, 1, 2, and 3. Matches a, b, c, d, e, and f.
\{m,n\}	Matches the preceding expression at least m but not more than n times.	a\{3,5\}	Matches “aaa”, “aaaa”, and “aaaaa”.

Note: You can change a special character—such as the period (.), asterisk (*), plus sign (+), or question mark (?)—into a literal character by prefixing it with a backslash (\). For example, to specify a literal period, asterisk, plus sign, or question mark, use the characters within the following parentheses: (\.), (*), (\+), (\?), (\^), (\\$).



Appendix E vNIOs Appliance Limitations

vNIOs appliances support most of the features of the Infoblox NIOS software, with some limitations. This appendix describes these limitations. [Table E.1](#) summarizes the supported Grid configurations on vNIOs appliances for Riverbed VMWare, and Hyper-V.

Note: VMWare Tools are automatically installed for each vNIOs appliance. Infoblox supports the control functions in Hyper-V Manager and VMWare Tools. For example, through the vSphere client, you can shut down the virtual appliance.

Table E.1 Supported vNIOs Appliance Configurations

vNIOs Appliance	vNIOs for Riverbed	vNIOs for VMWare (All IB-VM models)	vNIOs for Hyper-V
Single Independent Appliance	✗	✓	✓
Independent HA Pair	✗	✓	✓
Grid Master	✗	✓	✗
Grid Master Candidate	✗	✓	✗
HA Grid Member	✗	✓	✓
Single Grid Member	✓	✓	✓

For detailed information about the limitations on each vNIOs appliance, see the following:

- [vNIOs for Riverbed](#)
- [vNIOs for VMWare](#)
- [vNIOs for Hyper-V](#)

VNIOs FOR RIVERBED

vNIOs appliances on Riverbed have the following limitations:

- They can function as Grid members only. You cannot configure them as HA (high availability) pairs, Grid Masters, Grid Master candidates, or independent appliances.
- On a Grid with a vNIOs appliance on Riverbed as a Grid member, the maximum storage space for HTTP, FTP and TFTP is 1 GB (a Grid with only Infoblox appliances provides a maximum of 5 GB for these services), core files are 100 MB each, and syslog and infoblox.log files are 20 MB each. Scheduled backup file is 100 MB.
- The LAN interface is the only network interface available on the vNIOs appliance. You cannot configure the speed and transmission type (full or half duplex) of the network interface.
- You can use the traffic capture tool of the vNIOs software package to capture traffic only on the LAN port of the vNIOs appliance.
- vNIOs appliances on Riverbed do not support the following features:
 - Anycast addressing
 - Configuration as a DHCP lease history logging member
 - Dedicated MGMT port
 - NTP service
 - bloxTools environment
 - Configuration for managing Microsoft® Windows DNS servers
 - IF-MAP service

VNIOs FOR VMWARE

The Infoblox vNIOs for VMware can also run on Cisco SRE-V (Services Ready Engine Virtualization), which is part of the Cisco UCS (Unified Computing System) Express. For more information about vNIOs for VMware, refer to the *Infoblox Installation Guide for vNIOs Software on VMware*.

vNIOs for VMware appliances support most of the features of the Infoblox NIOS appliances, with the following limitations:

- You must have a vNIOs license installed on the appliance before you can access the Infoblox GUI.
- vNIOs appliances do not support the following features:
 - Configuration of port settings for MGMT, LAN, LAN2, and HA ports
 - The bloxTools environment
- The IB-BOB virtual appliance is supported on Cisco SRE-V and can function as a Grid member only. It does not support configuration as an independent appliance, an HA pair, a Grid Master, or a Grid Master candidate. It also does not support access to the Infoblox GUI.
- The IB-VM-250 virtual appliance supports all the services provided by vNIOs virtual appliances, but it is not recommended as a Grid Master or Grid Master candidate.
- The Captive Portal is supported only on IB-VM-1050 virtual appliances.
- When you configure an HA pair, both nodes in the HA pair must be vNIOs instances. You cannot configure a physical NIOS appliance and a vNIOs instance in an HA pair.
- vNIOs appliances run on virtual hardware. They do not have sensors to monitor the physical CPU temperature, fan speed, and system temperature.
- Changing the vNIOs appliance settings through the VMware vSphere or vCenter console may violate the terms of the vNIOs licensing and support models. The vNIOs appliance may not join the Grid or function properly.

vNIOs FOR HYPER-V

vNIOs for Microsoft Windows 2008 R2 server appliances support most of the features of the Infoblox NIOs appliances, with the following limitations:

- You must have a vNIOs license installed on the appliance before you can access the Infoblox GUI.
- vNIOs appliances do not support the following features:
 - Configuration of port settings for MGMT, LAN, LAN2, and HA ports
 - The bloxTools environment
- All the IB-VM appliance models support all the services provided by vNIOs virtual appliances, but Grid Master or Grid master candidate is not supported.
- Appliance model IB-VM-800 (Reporting) is not supported.
- The Captive Portal is supported only on IB-VM-1410 virtual appliances.
- vNIOs appliances run on virtual hardware. They do not have sensors to monitor the physical CPU temperature, fan speed, and system temperature.
- Changing the vNIOs appliance settings through the Hyper-V Manager or Virtual Machine Manager Administrator console may violate the terms of the vNIOs licensing and support models. The vNIOs appliance may not join the Grid or function properly.



Appendix F Product Compliance

This appendix describes the hardware components, requirements, and specifications, plus agency and RFC (Request for Comments) compliance for the Infoblox appliance. Topics in this appendix include:

- [Power Safety Information](#) on page 1310
 - [AC](#) on page 1310
 - [DC](#) on page 1310
- [Agency Compliance](#) on page 1311
 - [FCC](#) on page 1311
 - [Canadian Compliance](#) on page 1311
 - [VCCI](#) on page 1312
- [RFC Compliance](#) on page 1313
 - [DNS RFC Compliance](#) on page 1313
 - [DHCP RFC Compliance](#) on page 1315
 - [DHCPv6 RFC Compliance](#) on page 1316
 - [IDN \(Internationalized Domain Names\) RFC Compliance](#) on page 1316

POWER SAFETY INFORMATION

The main external power connector for the Infoblox appliance is located on the back of the system. Ensure power to the system is off before connecting the power cord into the power connector. Please read the following power safety statements for your AC- or DC-powered appliance:

AC

English

WARNING: *THIS PRODUCT RELIES ON THE BUILDING'S INSTALLATION FOR SHORT-CIRCUIT (OVERCURRENT) PROTECTION. ENSURE THAT A FUSE OR CIRCUIT BREAKER NO LARGER THAN 120VAC, 15AU.S. (240VAC, 10A INTERNATIONAL) IS USED ON THE PHASE CONDUCTORS (ALL CURRENT-CARRYING CONDUCTORS).*

French

WARNING: *POUR CE QUI EST DE LA PROTECTION CONTRE LES COURTS-CIRCUITS (SURTENSION), CE PRODUIT DÉPEND DE L'INSTALLATION ÉLECTRIQUE DU LOCAL. VÉRIFIER QU'UN FUSIBLE OU QU'UN DISJONCTEUR DE 120V ALT., 15A U.S. MAXIMUM (240V ALT., 10A INTERNATIONAL) EST UTILISÉ SUR LES CONDUCTEURS DE PHASE (CONDUCTEURS DE CHARGE).*

German

WARNING: *DIESES PRODUKT IST DARAUF ANGEWIESEN, DAß IM GEBÄUDE EIN KURZSCHLUß - BZW. ÜBERSTROMSCHUTZ INSTALLIERT IST. STELLEN SIE SICHER, DAß EINE SICHERUNG ODER EIN UNTERBRECHER VON NICHT MEHR ALS 240V WECHSELSTROM, 10A (BZW. IN DEN USA 120V WECHSELSTROM, 15A) AN DEN PHASENLEITERN (ALLEN STROMF, HRENDEN LEITERN) VERWENDET WIRD.*

DC

English

WARNING: *WHEN STRANDED WIRING IS REQUIRED, USE APPROVED WIRING TERMINATIONS, SUCH AS CLOSED-LOOP OR SPADE-TYPE WITH UPTURNED LUGS. THESE TERMINATIONS SHOULD BE THE APPROPRIATE SIZE FOR THE WIRES AND SHOULD CLAMP BOTH THE INSULATION AND CONDUCTOR.*

AGENCY COMPLIANCE

The Infoblox appliance is compliant with these EMI and safety agency regulations:

Table F.1 Agency Regulation Compliance

Standard	Agency	Marks
FCC Part 15	FCC	FCC
EN55022, EN55024, EN61000-3-2, EN61000-3-3	TUV	CE
UL60950/CSA60950	UL	cULus
EN60950	TUV	GS
CB Scheme	IECEE	Report and Certificate IEC 60950-1:2001
VCCI-A	VCCI	VCCI
AS/NZS 3548	ACMA	C-Tick

FCC

The FCC label on the back of the system indicates this network appliance is compliant with limits for a Class A digital device in accordance with Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when this equipment is operated in a commercial environment. Operation is subject to the following two conditions:

- This device might not cause harmful interference.
- This device must accept any interference received, including interference that may cause undesired operation.

This device generates, uses, and can radiate radio frequency energy if not installed and used in accordance with the instructions in this manual. Operating this equipment in a residential area is likely to cause harmful interference, and the customer will be required to rectify the interference at his or her own expense. This product requires the use of external shielded cables to maintain compliance pursuant to Part 15 of the FCC Rules.

Canadian Compliance

English

This Class A digital apparatus complies with Canadian ICES-003.

French

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

VCCI

The Infoblox appliance complies with this VCCI regulation (compliance statement follow by its translation):

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA 情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

This is a Class A product based on the Technical Requirements of the Voluntary Control Council for Interference Technology (VCCI). In a domestic environment this product may cause radio interference, in which case the user may be required to take corrective action.

Caution: Lithium battery included with this board. Do not puncture, mutilate, or dispose of battery in fire. Danger of explosion if battery is incorrectly replaced. Replace only with the same or equivalent type recommended by manufacturer. Dispose of used battery according to manufacturer instructions and in accordance with your local regulations.

RFC COMPLIANCE

The NIOS appliance is compliant with the following:

- Qualys and Nessus security requirements
- Joint Interoperability Test Command (JITC) certification for Internet Protocol version 6 capability
- RFCs (Request for Comments):
 - [DNS RFC Compliance](#) on page 1313
 - [DHCP RFC Compliance](#) on page 1315
 - [DHCPv6 RFC Compliance](#) on page 1316
 - [IDN \(Internationalized Domain Names\) RFC Compliance](#) on page 1316

DNS RFC Compliance

The NIOS appliance complies with the following DNS RFCs:

Table F.2 RFCs for DNS

RFC Number	RFC Title
805	Computer Mail Meeting Notes
811	Hostnames Server
819	The Domain Naming Convention for Internet User Applications
881	The Domain Names Plan and Schedule
882	Domain Names: Concepts and Facilities
883	Domain Names: Implementation Specification
897	Domain Name System Implementation Schedule
920	Domain Requirements
921	Domain Name System Implementation Schedule – Revised
973	Domain System Changes and Observations
974	Mail Routing and the Domain System
1032	Domain Administrators Guide
1033	Domain Administrators Operations Guide
1034	Domain Names – Concepts and Facilities
1035	Domain Names – Implementation and Specification
1101	DNS Encoding of Network Names and Other Types
1122	Requirements for Internet Hosts – Communication Layers
1123	Requirements for Internet Hosts – Application and Support
1178	Choosing a Name for Your Computer
1348	DNS NSAP RRs
1386	The US Domain

RFC Number	RFC Title
1464	Using the Domain Name System to Store Arbitrary String Attributes
1535	A Security Problem and Proposed Correction with Widely Deployed DNS Software
1536	Common DNS Implementation Errors and Suggested Fixes
1537	Common DNS Data File Configuration Errors
1591	Domain Name System Structure and Delegation
1611	DNS Server MIB Extensions
1612	DNS Resolver MIB Extensions
1637	DNS NSAP Resource Records
1664	Using the Internet DNS to Distribute RFC 1327 Mail Address Mapping Tables
1713	Tools for DNS debugging
1794	DNS Support for Load Balancing
1811	U.S. Government Internet Domain Names
1816	U.S. Government Internet Domain Names
1912	Common DNS Operational and Configuration Errors
1956	Registration in the MIL Domain
1982	Serial Number Arithmetic
1995	Incremental Zone Transfer in DNS
1996	A Mechanism for Prompt Notification of Zone Changes
2010	Operational Criteria for Root Name Servers
2052	A DNS RR for specifying the location of services (DNS SRV)
2053	The AM (Armenia) Domain
2136	Dynamic Updates in the Domain Name System (DNS UPDATE)
2142	Mailbox Names for Common Services, Roles and Functions
2146	U.S. Government Internet Domain Names
2168	Resolution of Uniform Resource Identifiers using the Domain Name System
2181	Clarifications to the DNS Specification
2182	Selection and Operation of Secondary DNS Servers
2219	Use of DNS Aliases for Network Services
2240	A Legal Basis for Domain Name Allocation
2308	Negative Caching of DNS Queries (DNS NCACHE)
2317	Classless IN-ADDR.ARPA Delegation
2352	A Convention for Using Legal Names as Domain Names
2537	RSA/MD5 KEYS and SIGs in the Domain Name System (DNS)
2606	Reserved Top Level DNS Names

RFC Number	RFC Title
2671	Extension Mechanisms for DNS (EDNS0)
2782	A DNS RR for Specifying the Location of Services (DNS SRV)
2845	Secret Key Transaction Authentication for DNS (TSIG)
2915	The Naming Authority Pointer (NAPTR) DNS Resource Record
3596	DNS Extensions to Support IP Version 6
3645	Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)
3768	Virtual Router Redundancy Protocol (VRRP)
4033	DNS Security Introduction and Requirements
4034	Resource Records for the DNS Security Extensions
4035	Protocol Modifications for the DNS Security Extensions
4641	DNSSEC Operational Practices
4956	DNS Security (DNSSEC) Opt-In
4986	Requirements Related to DNS Security (DNSSEC) Trust Anchor Rollover
5155	DNSSEC Hashed Authenticated Denial of Existence
5702	Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC
5936	DNS Zone Transfer Protocol (AXFR)
6147	DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers

DHCP RFC Compliance

The appliance complies with the following DHCP RFCs:

Table F.3 RFCs for DHCP

RFC Number	RFC Title
1531	Dynamic Host Configuration Protocol
1534	Interoperation Between DHCP and BOOTP
1542	Clarifications and Extensions for the Bootstrap Protocol
2131	Dynamic Host Configuration Protocol
2132	DHCP Options and BOOTP Vendor Extensions
3046	DHCP Relay Agent Information Option
3315	Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
3925	Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4)
4388	Dynamic Host Configuration Protocol (DHCP) Leasequery

DHCPv6 RFC Compliance

The appliance complies with the following DHCPv6 RFCs:

Table F.4 RFCs for DHCPv6

RFC Number	RFC Title
4075	Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6
3898	Network Information Service (NIS) Configuration Options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
3736	Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6
3646	DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
3633	IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6
3319	Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers

IDN (Internationalized Domain Names) RFC Compliance

The appliance complies with the following IDN RFCs:

Table F.5 RFCs for IDN

RFC Number	RFC Title
3492	Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)
5890	Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework
5891	Internationalized Domain Names for Applications (IDNA): Background, Explanation, and Rationale
5892	The Unicode code points and IDNA
5893	Right-to-left scripts for IDNA
5894	Internationalized Domain Names in Applications (IDNA): Protocol
5895	Mapping Characters in IDNA2008
6452	The Unicode Code Points and Internationalized Domain Names for Applications (IDNA) - Unicode 6.0



Appendix G Open Source Copyright and License Statements

Infoblox has made every attempt to adhere to the guidelines for use and contribution to the open source community. Please report back to Infoblox any suspected violations of the copyrights, use of open source contributions via the distribution of binaries and/or source from Infoblox. It is the intent of Infoblox to comply with the open source rules of use, and comply with the various copyrights found in the distribution of the products from Infoblox.

This appendix contains the copyright notices for the binary-only distribution from Infoblox. Source changes are contributed back to the open source community when the copyright holder states this is desired. As stated by the enclosed copyrights, a copy of open source files used in our binary-only distribution is available from Infoblox. There is a nominal cost to obtain a CD containing the source files, to cover our costs of duplication and distribution. To obtain a copy of the source, contact us via e-mail at info@infoblox.com, or call us at 1.408.625.4200. The sections in this appendix include:

- [*GNU General Public License*](#) on page 1319
- [*GNU Lesser General Public License*](#) on page 1322
- [*Apache Software License, Version 2.0*](#) on page 1328
- [*ISC BIND Copyright*](#) on page 1334
- [*ISC DHCP Copyright*](#) on page 1335
- [*Julian Seward Copyright*](#) on page 1336
- [*Carnegie Mellon University Copyright*](#) on page 1336
- [*Thai Open Source Software Center Copyright*](#) on page 1337
- [*Ian F. Darwin Copyright*](#) on page 1338
- [*Lawrence Berkeley Copyright*](#) on page 1339
- [*MIT Kerberos Copyright*](#) on page 1339
- [*BSD License*](#) on page 1340
- [*David L. Mills Copyright*](#) on page 1341
- [*OpenLDAP License*](#) on page 1341
- [*OpenSSL License*](#) on page 1342
- [*VIM License*](#) on page 1343
- [*ZLIB License*](#) on page 1345
- [*Wietse Venema Copyright*](#) on page 1345
- [*ECLIPSE SOFTWARE*](#) on page 1346
- [*Eclipse Public License - v 1.0*](#) on page 1346
- [*AOP Alliance \(Java/J2EE AOP standards\)*](#) on page 1350
- [*ASM*](#) on page 1350

- [*Distributed Computing Laboratory, Emory University*](#) on page 1351
- [*COMMON DEVELOPMENT AND DISTRIBUTION LICENSE \(CDDL\)*](#) on page 1351
- [*The FreeType Project LICENSE*](#) on page 1355
- [*The Independent JPEG Group's JPEG software*](#) on page 1359
- [*Net-SNMP*](#) on page 1361
- [*The PHP License, version 3.01*](#) on page 1367
- [*INFO-ZIP*](#) on page 1368
- [*MIT License*](#) on page 1370
- [*Ehcache*](#) on page 1370

GNU GENERAL PUBLIC LICENSE

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundations software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each authors protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyones free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Programs source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

GNU LESSER GENERAL PUBLIC LICENSE

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original authors reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the users freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an appropriate program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customers own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the users computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice. This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

APACHE SOFTWARE LICENSE, VERSION 2.0

Copyright (c) 2004 The Apache Software Foundation. All rights reserved.

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications,

including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
 - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
 - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and

do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all

other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

PERL ARTISTIC LICENSE

The "Artistic License"

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.

2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.

3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:

a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as unet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b) use the modified Package only within your corporation or organization.

c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:

a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b) accompany the distribution with the machine-readable source of the Package with your modifications.

c) give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Packages interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whoever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.

7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.

8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Packages interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

ISC BIND COPYRIGHT

Copyright (C) 1996-2002 Internet Software Consortium.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND INTERNET SOFTWARE CONSORTIUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Portions Copyright (C) 1996-2001 Nomimum, Inc.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NOMINUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR

ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

ISC DHCP COPYRIGHT

Copyright (c) 1995, 1996, 1997, 1998, 1999 Internet Software Consortium -DHCP. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Internet Software Consortium - DHCP nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY INTERNET SOFTWARE

CONSORTIUM AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

JULIAN SEWARD COPYRIGHT

This program, "bzip2" and associated library "libbzip2", are copyright (C) 1996-2002 Julian R Seward. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
3. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
4. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Julian Seward, Cambridge, UK.

jseward@acm.org

bzip2/libbzip2 version 1.0.2 of 30 December 2001

CARNEGIE MELLON UNIVERSITY COPYRIGHT

```
/*
 * Copyright (c) 2001 Carnegie Mellon University. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in
 *    the documentation and/or other materials provided with the
 *    distribution.
 *
 * 3. The name "Carnegie Mellon University" must not be used to
 *    endorse or promote products derived from this software without
 *    prior written permission. For permission or any other legal
 *    details, please contact
```

```

*      Office of Technology Transfer
*      Carnegie Mellon University
*      5000 Forbes Avenue
*      Pittsburgh, PA 15213-3890
*      (412) 268-4387, fax: (412) 268-7395
*      tech-transfer@andrew.cmu.edu
*
* 4. Redistributions of any form whatsoever must retain the following
*      acknowledgment:
*
*      "This product includes software developed by Computing Services
*      at Carnegie Mellon University (http://www.cmu.edu/computing/)."

```

THAI OPEN SOURCE SOFTWARE CENTER COPYRIGHT

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd
and Clark Cooper

Copyright (c) 2001, 2002 Expat maintainers.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

IAN F. DARWIN COPYRIGHT

Copyright (c) Ian F. Darwin 1986, 1987, 1989, 1990, 1991, 1992, 1994, 1995.
Software written by Ian F. Darwin and others;
maintained 1994-1999 Christos Zoulas.

This software is not subject to any export provision of the United States Department of Commerce, and may be exported to any country or planet.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice immediately at the beginning of the file, without modification, this list of conditions, and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. All advertising materials mentioning features or use of this software must display the following acknowledgement:

This product includes software developed by Ian F. Darwin and others.

4. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

LAWRENCE BERKELEY COPYRIGHT

Copyright (c) 1990 The Regents of the University of California.
All rights reserved.

This code is derived from software contributed to Berkeley by Vern Paxson.

The United States Government has rights in this work pursuant to contract no. DE-AC03-76SF00098 between the United States Department of Energy and the University of California.

Redistribution and use in source and binary forms with or without modification are permitted provided that: (1) source distributions retain this entire copyright notice and comment, and (2) distributions including binaries display the following acknowledgement: "This product includes software developed by the University of California, Berkeley and its contributors" in the documentation or other materials provided with the distribution and in all advertising materials mentioning features or use of this software. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

MIT KERBEROS COPYRIGHT

Copyright Notice and Legal Administrivia

Copyright (C) 1985-2002 by the Massachusetts Institute of Technology.

All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Furthermore if you modify this software you must label your software as modified software and not distribute it in such a fashion that it might be confused with the original MIT software.

M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Individual source code files are copyright MIT, Cygnus Support, OpenVision, Oracle, Sun Soft, FundsXpress, and others.

Project Athena, Athena, Athena MUSE, Discuss, Hesiod, Kerberos, Moira, and Zephyr are trademarks of the Massachusetts Institute of Technology (MIT). No commercial use of these trademarks may be made without prior written permission of MIT.

"Commercial use" means use of a name in a product or other for-profit manner. It does NOT prevent a commercial firm from referring to the MIT trademarks in order to convey information (although in doing so, recognition of their trademark status should be given).

The following copyright and permission notice applies to the OpenVision Kerberos Administration system located in kadmin/create, kadmin/dbutil, kadmin/passwd, kadmin/server, lib/kadm5, and portions of lib/rpc:

Copyright, OpenVision Technologies, Inc., 1996, All Rights Reserved

WARNING: Retrieving the OpenVision Kerberos Administration system source code, as described below, indicates your acceptance of the following terms. If you do not agree to the following terms, do not retrieve the OpenVision Kerberos administration system.

You may freely use and distribute the Source Code and Object Code compiled from it, with or without modification, but this Source Code is provided to you "AS IS" EXCLUSIVE OF ANY WARRANTY, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, OR ANY OTHER WARRANTY, WHETHER EXPRESS OR IMPLIED. IN NO EVENT WILL OPENVISION HAVE ANY LIABILITY FOR ANY LOST PROFITS, LOSS OF DATA OR COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, INCLUDING, WITHOUT LIMITATION, THOSE RESULTING FROM THE USE OF THE SOURCE CODE, OR THE FAILURE OF THE SOURCE CODE TO PERFORM, OR FOR ANY OTHER REASON.

OpenVision retains all copyrights in the donated Source Code. OpenVision also retains copyright to derivative works of the Source Code, whether created by OpenVision or by a third party. The OpenVision copyright notice must be preserved if derivative works are made based on the donated Source Code.

OpenVision Technologies, Inc. has donated this Kerberos Administration system to MIT for inclusion in the standard Kerberos 5 distribution. This donation underscores our commitment to continuing Kerberos technology development and our gratitude for the valuable work which has been performed by MIT and the Kerberos community.

Portions contributed by Matt Crawford <crawdada@fnal.gov> were work performed at Fermi National Accelerator Laboratory, which is operated by Universities Research Association, Inc., under contract DE-AC02-76CHO3000 with the U.S. Department of Energy.

BSD LICENSE

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

DAVID L. MILLS COPYRIGHT

```
*****
*
* Copyright (c) David L. Mills 1992-2003 *
*
* Permission to use, copy, modify, and distribute this software and *
* its documentation for any purpose and without fee is hereby *
* granted, provided that the above copyright notice appears in all *
* copies and that both the copyright notice and this permission *
* notice appear in supporting documentation, and that the name *
* University of Delaware not be used in advertising or publicity *
* pertaining to distribution of the software without specific, *
* written prior permission. The University of Delaware makes no *
* representations about the suitability this software for any *
* purpose. It is provided "as is" without express or implied *
* warranty. *
*
*****
```

OPENLDAP LICENSE

The OpenLDAP Public License
Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions in source form must retain copyright statements and notices,
2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

OPENSSL LICENSE

```

/* =====
* Copyright (c) 1998-2003 The OpenSSL Project. All rights reserved.
*
* Redistribution and use in source and binary forms, with or without
* modification, are permitted provided that the following conditions
* are met:
*
* 1. Redistributions of source code must retain the above copyright
*    notice, this list of conditions and the following disclaimer.
*
* 2. Redistributions in binary form must reproduce the above copyright
*    notice, this list of conditions and the following disclaimer in
*    the documentation and/or other materials provided with the
*    distribution.
*
* 3. All advertising materials mentioning features or use of this
*    software must display the following acknowledgment:
*    "This product includes software developed by the OpenSSL Project
*    for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
*
* 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
*    endorse or promote products derived from this software without
*    prior written permission. For written permission, please contact
*    openssl-core@openssl.org.
*
* 5. Products derived from this software may not be called "OpenSSL"
*    nor may "OpenSSL" appear in their names without prior written
*    permission of the OpenSSL Project.
*
* 6. Redistributions of any form whatsoever must retain the following
*    acknowledgment:
*    "This product includes software developed by the OpenSSL Project
*    for use in the OpenSSL Toolkit (http://www.openssl.org/)"
*
* THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY
* EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
* IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
* PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
* ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
* SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
* NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
* LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
* HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
* STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
* ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
* OF THE POSSIBILITY OF SUCH DAMAGE.
* =====
*
* This product includes cryptographic software written by Eric Young
* (eay@cryptsoft.com). This product includes software written by Tim
* Hudson (tjh@cryptsoft.com).
*
*/

```

VIM LICENSE

COPYING

Vim is Charityware. You can use and copy it as much as you like, but you are encouraged to make a donation to orphans in Uganda. Please read the file "runtime/doc/uganda.txt" for details.

There are no restrictions on distributing an unmodified copy of Vim. Parts of Vim may also be distributed, but this text must always be included. You are allowed to include executables that you made from the unmodified Vim sources, your own usage examples and Vim scripts.

If you distribute a modified version of Vim, you are encouraged to send the maintainer a copy, including the source code. Or make it available to the maintainer through ftp; let him know where it can be found. If the number of changes is small (e.g., a modified Makefile) e-mailing the diffs will do. When the maintainer asks for it (in any way) you must make your changes, including source code, available to him.

The maintainer reserves the right to include any changes in the official version of Vim. This is negotiable. You are not allowed to distribute a modified version of Vim when you are not willing to make the source code available to the maintainer.

The current maintainer is Bram Moolenaar <Bram@vim.org>. If this changes, it will be announced in appropriate places (most likely www.vim.org and comp.editors). When it is completely impossible to contact the maintainer, the obligation to send him modified source code ceases.

It is not allowed to remove these restrictions from the distribution of the Vim sources or parts of it. These restrictions may also be used for previous Vim releases instead of the text that was included with it.

Vim is Charityware. You can use and copy it as much as you like, but you are encouraged to make a donation for needy children in Uganda. Please see |kcc| below or visit the ICCF web site, available at these mirrors:

<http://iccf-holland.org/>
<http://www.vim.org/iccf/>
<http://www.iccf.nl/>

The Open Publication License applies to the Vim documentation, see |manual-copyright|.

=== begin of license ===

VIM LICENSE

There are no restrictions on distributing unmodified copies of Vim except that they must include this license text. You can also distribute unmodified parts of Vim, likewise unrestricted except that they must include this license text. You are also allowed to include executables that you made from the unmodified Vim sources, plus your own usage examples and Vim scripts.

It is allowed to distribute a modified (or extended) version of Vim, including executables and/or source code, when the following four conditions are met:

- 1) This license text must be included unmodified.
- 2) The modified Vim must be distributed in one of the following five

ways:

a) If you make changes to Vim yourself, you must clearly describe in the distribution how to contact you. When the maintainer asks you (in any way) for a copy of the modified Vim you distributed, you must make your changes, including source code, available to the maintainer without fee. The maintainer reserves the right to include your changes in the official version of Vim. What the maintainer will do with your changes and under what license they will be distributed is negotiable. If there has been no negotiation then this license, or a later version, also applies to your changes. The current maintainer is Bram Moolenaar <Bram@vim.org>. If this changes it will be announced in appropriate places (most likely vim.sf.net, www.vim.org and/or comp.editors). When it is completely impossible to contact the maintainer, the obligation to send him your changes ceases. Once the maintainer has confirmed that he has received your changes they will not have to be sent again.

b) If you have received a modified Vim that was distributed as mentioned under a) you are allowed to further distribute it unmodified, as mentioned at 1). If you make additional changes the text under a) applies to those changes.

c) Provide all the changes, including source code, with every copy of the modified Vim you distribute. This may be done in the form of a context diff. You can choose what license to use for new code you add. The changes and their license must not restrict others from making their own changes to the official version of Vim.

d) When you have a modified Vim which includes changes as mentioned under c), you can distribute it without the source code for the changes if the following three conditions are met:

- The license that applies to the changes permits you to distribute the changes to the Vim maintainer without fee or restriction, and permits the Vim maintainer to include the changes in the official version of Vim without fee or restriction.
- You keep the changes for at least three years after last distributing the corresponding modified Vim. When the maintainer or someone who you distributed the modified Vim to asks you (in any way) for the changes within this period, you must make them available to him.
- You clearly describe in the distribution how to contact you. This contact information must remain valid for at least three years after last distributing the corresponding modified Vim, or as long as possible.

e) When the GNU General Public License (GPL) applies to the changes, you can distribute the modified Vim under the GNU GPL version 2 or any later version.

3) A message must be added, at least in the output of the ":version" command and in the intro screen, such that the user of the modified Vim is able to see that it was modified. When distributing as mentioned under 2)e) adding the message is only required for as far as this does not conflict with the license used for the changes.

4) The contact information as required under 2)a) and 2)d) must not be removed or changed, except that the person himself can make corrections.

If you distribute a modified version of Vim, you are encouraged to use the Vim license for your changes and make them available to the maintainer, including the source code. The preferred way to do this is by e-mail or by uploading the files to a server and e-mailing the URL. If the number of changes is small (e.g., a modified Makefile) e-mailing a context diff will do. The e-mail address to be used is <maintainer@vim.org>

It is not allowed to remove this license from the distribution of the Vim sources, parts of it or from a modified version. You may use this license for previous Vim releases instead of the license that they came with, at your option.

=== end of license ===

ZLIB LICENSE

(C) 1995-2002 Jean-loup Gailly and Mark Adler

This software is provided as-is, without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.

This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler
jloup@gzip.org madler@alumni.caltech.edu

If you use the zlib library in a product, we would appreciate *not* receiving lengthy legal documents to sign. The sources are provided for free but without warranty of any kind. The library has been entirely written by Jean-loup Gailly and Mark Adler; it does not include third-party code.

If you redistribute modified sources, we would appreciate that you include in the file ChangeLog history information documenting your changes.

WIETSE VENEMA COPYRIGHT

```

/***** * Copyright
1995 by Wietse Venema. All rights reserved. Some individual * files may be covered by
other copyrights. * * This material was originally written and compiled by Wietse Venema
at * Eindhoven University of Technology, The Netherlands, in 1990, 1991, * 1992, 1993,
1994 and 1995. * * Redistribution and use in source and binary forms are permitted *
provided that this entire copyright notice is duplicated in all such * copies. * * This
software is provided "as is" and without any expressed or implied * warranties, including,
without limitation, the implied warranties of * merchantability and fitness for any
particular purpose.
*****/

```

ECLIPSE SOFTWARE

The product includes Eclipse software (the "Eclipse Program") provided by the Eclipse Foundation and licensed to Infoblox Inc. under the Eclipse Public License v1.0.

EXCEPT AS EXPRESSLY SET FORTH IN THE ECLIPSE PUBLIC LICENSE, THE ECLIPSE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

EXCEPT AS EXPRESSLY SET FORTH IN THE ECLIPSE PUBLIC LICENSE, NEITHER THE ECLIPSE FOUNDATION NOR ANY CONTRIBUTORS TO THE ECLIPSE PROGRAM SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THE ECLIPSE PROGRAM, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any provisions provided by Infoblox relating to the Eclipse Program which differ from the above terms or the Eclipse Public License are offered by Infoblox alone and not by any other party.

The source code for the Eclipse Program is available from Infoblox as described in the open source introduction.

ECLIPSE PUBLIC LICENSE - V 1.0

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS ECLIPSE PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

b) in the case of each subsequent Contributor:

i) changes to the Program, and

ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents " mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

a) it complies with the terms and conditions of this Agreement; and

b) its license agreement:

i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;

iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

a) it must be made available under this Agreement; and

b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must:

- a) promptly notify the Commercial Contributor in writing of such claim, and
- b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. The Eclipse Foundation is the initial Agreement Steward. The Eclipse Foundation may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses

to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

Related Links

- * EPL in plain HTML
- * The EPL on OSI's site
- * CPL to EPL conversion

AOP ALLIANCE (JAVA/J2EE AOP STANDARDS)

LICENCE: all the source code provided by AOP Alliance is Public Domain.

ASM

Copyright (c) 2000-2005 INRIA, France Telecom
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS

INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

DISTRIBUTED COMPUTING LABORATORY, EMORY UNIVERSITY

This software is released to the public domain, in the spirit of the original code written by Doug Lea. The code can be used for any purpose, modified, and redistributed without acknowledgment. No warranty is provided, either express or implied.

COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)

Version 1.0

1. Definitions.

1.1. Contributor means each individual or entity that creates or contributes to the creation of Modifications.

1.2. Contributor Version means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.

1.3. Covered Software means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.

1.4. Executable means the Covered Software in any form other than Source Code.

1.5. Initial Developer means the individual or entity that first makes Original Software available under this License.

1.6. Larger Work means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.

1.7. License means this document.

1.8. Licensable means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. Modifications means the Source Code and Executable form of any of the following:

A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;

B. Any new file that contains any part of the Original Software or previous Modification; or

C. Any new file that is contributed or otherwise made available under the terms of this License.

1.10. Original Software means the Source Code and Executable form of computer software code that is originally released under this License.

1.11. Patent Claims means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.12. Source Code means (a) the common form of computer software code in which 1.13. You (or Your) means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, You includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, control means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. License Grants.

2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer, to use, reproduce, modify, display, perform, sublicense and distribute the Original Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).

(c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distributes or otherwise makes the Original Software available to a third party under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: (1) for code that You delete from the Original Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof), either on an unmodified basis, with other Modifications, as Covered Software and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distributes or otherwise makes the Modifications available to a third party.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: (1) for any code that Contributor has deleted from the Contributor Version; (2) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License.

You must include a copy of this License with every copy of the Source Code form of the Covered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

3.3. Required Notices.

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contributor or the Initial Developer.

3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipients rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form under a different license, You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.6. Larger Works.

You may create a Larger Work by combining Covered Software with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Software.

4. Versions of the License.

4.1. New Versions.

Sun Microsystems, Inc. is the initial license steward and may publish revised and/or new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any

subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Software available under the terms of any subsequent version of the License published by the license steward.

4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any references to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE IS PROVIDED UNDER THIS LICENSE ON AN AS IS BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED SOFTWARE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

6. TERMINATION.

6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as Participant) alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Developer) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contributors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.

6.3. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been validly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

7. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED SOFTWARE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOST PROFITS, LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTYS NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

8. U.S. GOVERNMENT END USERS.

The Covered Software is a commercial item, as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of commercial computer software (as that term is defined at 48 C.F.R. 252.227-7014(a)(1)) and commercial computer software documentation as such terms

are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that addresses Government rights in computer software under this License.

9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice contained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdictions conflict-of-law provisions. Any litigation relating to this License shall be subject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export administration regulations (and the export control laws and regulation of any other countries) when You use, distribute or otherwise make available any Covered Software.

10. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

THE FREETYPE PROJECT LICENSE

2002-Apr-11

Copyright 1996-2002 by

David Turner, Robert Wilhelm, and Werner Lemberg

Introduction

=====

The FreeType Project is distributed in several archive packages; some of them may contain, in addition to the FreeType font engine, various tools and contributions which rely on, or relate to, the FreeType Project.

This license applies to all files found in such packages, and which do not fall under their own explicit license. The license

affects thus the FreeType font engine, the test programs, documentation and makefiles, at the very least.

This license was inspired by the BSD, Artistic, and IJG (Independent JPEG Group) licenses, which all encourage inclusion and use of free software in commercial and freeware products alike. As a consequence, its main points are that:

- o We don't promise that this software works. However, we will be interested in any kind of bug reports. ('as is' distribution)
- o You can use this software for whatever you want, in parts or full form, without having to pay us. ('royalty-free' usage)
- o You may not pretend that you wrote this software. If you use it, or only parts of it, in a program, you must acknowledge somewhere in your documentation that you have used the FreeType code. ('credits')

We specifically permit and encourage the inclusion of this software, with or without modifications, in commercial products. We disclaim all warranties covering The FreeType Project and assume no liability related to The FreeType Project.

Finally, many people asked us for a preferred form for a credit/disclaimer to use in compliance with this license. We thus encourage you to use the following text:

```
"""
Portions of this software are copyright © 1996-2002 The FreeType
Project (www.freetype.org). All rights reserved.
"""
```

Legal Terms

=====

0. Definitions

Throughout this license, the terms 'package', 'FreeType Project', and 'FreeType archive' refer to the set of files originally distributed by the authors (David Turner, Robert Wilhelm, and Werner Lemberg) as the 'FreeType Project', be they named as alpha,

beta or final release.

'You' refers to the licensee, or person using the project, where 'using' is a generic term including compiling the project's source code as well as linking it to form a 'program' or 'executable'. This program is referred to as 'a program using the FreeType engine'.

This license applies to all files distributed in the original FreeType Project, including all source code, binaries and documentation, unless otherwise stated in the file in its original, unmodified form as distributed in the original archive. If you are unsure whether or not a particular file is covered by this license, you must contact us to verify this.

The FreeType Project is copyright (C) 1996-2000 by David Turner, Robert Wilhelm, and Werner Lemberg. All rights reserved except as specified below.

1. No Warranty

THE FREETYPE PROJECT IS PROVIDED 'AS IS' WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL ANY OF THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY DAMAGES CAUSED BY THE USE OR THE INABILITY TO USE, OF THE FREETYPE PROJECT.

2. Redistribution

This license grants a worldwide, royalty-free, perpetual and irrevocable right and license to use, execute, perform, compile, display, copy, create derivative works of, distribute and sublicense the FreeType Project (in both source and object code forms) and derivative works thereof for any purpose; and to authorize others to exercise some or all of the rights granted herein, subject to the following conditions:

- o Redistribution of source code must retain this license file ('FTL.TXT') unaltered; any additions, deletions or changes to the original files must be clearly indicated in accompanying documentation. The copyright notices of the unaltered, original files must be preserved in all copies of source

files.

- o Redistribution in binary form must provide a disclaimer that states that the software is based in part of the work of the FreeType Team, in the distribution documentation. We also encourage you to put an URL to the FreeType web page in your documentation, though this isn't mandatory.

These conditions apply to any software derived from or based on the FreeType Project, not just the unmodified files. If you use our work, you must acknowledge us. However, no fee need be paid to us.

3. Advertising

Neither the FreeType authors and contributors nor you shall use the name of the other for commercial, advertising, or promotional purposes without specific prior written permission.

We suggest, but do not require, that you use one or more of the following phrases to refer to this software in your documentation or advertising materials: `FreeType Project', `FreeType Engine', `FreeType library', or `FreeType Distribution'.

As you have not signed this license, you are not required to accept it. However, as the FreeType Project is copyrighted material, only this license, or another one contracted with the authors, grants you the right to use, distribute, and modify it. Therefore, by using, distributing, or modifying the FreeType Project, you indicate that you understand and accept all the terms of this license.

4. Contacts

There are two mailing lists related to FreeType:

- o freetype@freetype.org

Discusses general use and applications of FreeType, as well as future and wanted additions to the library and distribution. If you are looking for support, start in this list if you haven't found anything to help you in the documentation.

o `devel@freetype.org`

Discusses bugs, as well as engine internals, design issues, specific licenses, porting, etc.

o `http://www.freetype.org`

Holds the current FreeType web page, which will allow you to download our latest development version and read online documentation.

You can also contact us individually at:

David Turner	<code><david.turner@freetype.org></code>
Robert Wilhelm	<code><robert.wilhelm@freetype.org></code>
Werner Lemberg	<code><werner.lemberg@freetype.org></code>

THE INDEPENDENT JPEG GROUP'S JPEG SOFTWARE

In plain English:

1. We don't promise that this software works. (But if you find any bugs, please let us know!)
2. You can use this software for whatever you want. You don't have to pay us.
3. You may not pretend that you wrote this software. If you use it in a program, you must acknowledge somewhere in your documentation that you've used the IJG code.

In legalese:

The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

This software is copyright (C) 1991-1998, Thomas G. Lane.
All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

- (1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice

unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.

(2) If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".

(3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

These conditions apply to any software derived from or based on the IJG code, not just to the unmodified library. If you use our work, you ought to acknowledge us.

Permission is NOT granted for the use of any IJG author's name or company name in advertising or publicity relating to this software or products derived from it. This software may be referred to only as "the Independent JPEG Group's software".

We specifically permit and encourage the use of this software as the basis of commercial products, provided that all warranty or liability claims are assumed by the product vendor.

ansi2knr.c is included in this distribution by permission of L. Peter Deutsch, sole proprietor of its copyright holder, Aladdin Enterprises of Menlo Park, CA. ansi2knr.c is NOT covered by the above copyright and conditions, but instead by the usual distribution terms of the Free Software Foundation; principally, that you must include source code if you redistribute it. (See the file ansi2knr.c for full details.) However, since ansi2knr.c is not needed as part of any program generated from the IJG code, this does not limit you more than the foregoing paragraphs do.

The Unix configuration script "configure" was produced with GNU Autoconf. It is copyright by the Free Software Foundation but is freely distributable. The same holds for its supporting scripts (config.guess, config.sub, ltconfig, ltmain.sh). Another support script, install-sh, is copyright by M.I.T. but is also freely distributable.

It appears that the arithmetic coding option of the JPEG spec is covered by patents owned by IBM, AT&T, and Mitsubishi. Hence arithmetic coding cannot legally be used without obtaining one or more licenses. For this reason, support for arithmetic coding has been removed from the free JPEG software. (Since arithmetic coding provides only a marginal gain over the unpatented Huffman mode, it is unlikely that very many implementations will support it.) So far as we are aware, there are no patent restrictions on the remaining code.

The IJG distribution formerly included code to read and write GIF files. To avoid entanglement with the Unisys LZW patent, GIF reading support has been removed altogether, and the GIF writer has been simplified to produce "uncompressed GIFs". This technique does not use the LZW algorithm; the resulting GIF files are larger than usual, but are readable by all standard GIF decoders.

We are required to state that

"The Graphics Interchange Format(c) is the Copyright property of CompuServe Incorporated. GIF(sm) is a Service Mark property of CompuServe Incorporated."

NET-SNMP

Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001. An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, Inc. to the project since 2003.

Code has been contributed to this project by many people over the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,

PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;
OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF
ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd.
All rights reserved.

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in the
documentation and/or other materials provided with the distribution.
- * The name of Cambridge Broadband Ltd. may not be used to endorse or
promote products derived from this software without specific prior
written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY
EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR
BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE
OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN
IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara,
California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) ----

Copyright (c) 2003-2005, Sparta, Inc
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the

documentation and/or other materials provided with the distribution.

- * Neither the name of Sparta, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 6: Cisco/BUPTNIC copyright notice (BSD) ----

Copyright (c) 2004, Cisco, Inc and Information Network
Center of Beijing University of Posts and Telecommunications.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,

PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;
OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR
OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF
ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) ----

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003

oss@fabasoft.com

Author: Bernhard Penz <bernhard.penz@fabasoft.com>

Redistribution and use in source and binary forms, with or without
modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice,
this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright
notice, this list of conditions and the following disclaimer in the
documentation and/or other materials provided with the distribution.
- * The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries,
brand or product names may not be used to endorse or promote products
derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS'' AND ANY
EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE
LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR
CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF
SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR
BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE
OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN
IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE PHP LICENSE, VERSION 3.01

Copyright (c) 1999 - 2006 The PHP Group. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, is permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name "PHP" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact group@php.net.
4. Products derived from this software may not be called "PHP", nor may "PHP" appear in their name, without prior written permission from group@php.net. You may indicate that your software works in conjunction with PHP by saying "Foo for PHP" instead of calling it "PHP Foo" or "phpfoo"
5. The PHP Group may publish revised and/or new versions of the license from time to time. Each version will be given a distinguishing version number.
Once covered code has been published under a particular version of the license, you may always continue to use it under the terms of that version. You may also choose to use such covered code under the terms of any subsequent version of the license published by the PHP Group. No one other than the PHP Group has the right to modify the terms applicable to covered code created under this License.
6. Redistributions of any form whatsoever must retain the following acknowledgment:
"This product includes PHP software, freely available from
<<http://www.php.net/software/>>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A

PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PHP DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at group@php.net.

For more information on the PHP Group and the PHP project, please see [<http://www.php.net>](http://www.php.net).

PHP includes the Zend Engine, freely available at [<http://www.zend.com>](http://www.zend.com).

INFO-ZIP

This is version 2007-Mar-4 of the Info-ZIP license.

The definitive version of this document should be available at <ftp://ftp.info-zip.org/pub/infozip/license.html> indefinitely and a copy at <http://www.info-zip.org/pub/infozip/license.html>.

Copyright (c) 1990-2007 Info-ZIP. All rights reserved.

For the purposes of this copyright and license, "Info-ZIP" is defined as the following set of individuals:

Mark Adler, John Bush, Karl Davis, Harald Denker, Jean-Michel Dubois, Jean-loup Gailly, Hunter Goatley, Ed Gordon, Ian Gorman, Chris Herborth, Dirk Haase, Greg Hartwig, Robert Heath, Jonathan Hudson, Paul Kienitz, David Kirschbaum, Johnny Lee, Onno van der Linden, Igor Mandrichenko, Steve P. Miller, Sergio Monesi, Keith Owens, George Petrov, Greg Roelofs, Kai Uwe Rommel, Steve Salisbury, Dave Smith, Steven M. Schweda, Christian Spieler, Cosmin Truta, Antoine Verheijen, Paul von Behren, Rich Wales, Mike White.

This software is provided "as is," without warranty of any kind, express or implied. In no event shall Info-ZIP or its contributors be held liable for any direct, indirect, incidental, special or consequential damages arising out of the use of or inability to use this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the above disclaimer and the following restrictions:

1. Redistributions of source code (in whole or in part) must retain the above copyright notice, definition, disclaimer, and this list of conditions.
2. Redistributions in binary form (compiled executables and libraries) must reproduce the above copyright notice, definition, disclaimer, and this list of conditions in documentation and/or other materials provided with the distribution. The sole exception to this condition is redistribution of a standard UnZipSFX binary (including SFXWiz) as part of a self-extracting archive; that is permitted without inclusion of this license, as long as the normal SFX banner has not been removed from the binary or disabled.
3. Altered versions--including, but not limited to, ports to new operating systems, existing ports with new graphical interfaces, versions with modified or added functionality, and dynamic, shared, or static library versions not from Info-ZIP--must be plainly marked as such and must not be misrepresented as being the original source or, if binaries, compiled from the original source. Such altered versions also must not be misrepresented as being Info-ZIP releases--including, but not limited to, labeling of the altered versions with the names "Info-ZIP" (or any variation thereof, including, but not limited to, different capitalizations), "Pocket UnZip," "WiZ" or "MacZip" without the explicit permission of Info-ZIP. Such altered versions are further prohibited from misrepresentative use of the Zip-Bugs or Info-ZIP e-mail addresses or the Info-ZIP URL(s), such as to imply Info-ZIP will provide support for the altered versions.
4. Info-ZIP retains the right to use the names "Info-ZIP," "Zip," "UnZip," "UnZipSFX," "WiZ," "Pocket UnZip," "Pocket Zip," and "MacZip" for its own source and binary releases.

MIT LICENSE

Copyright (c) 2010 Paul T. McGuire

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

EHCACHE

The open source Ehcache project is licensed under the Apache 2.0 License. The text of the license is available below:

```
/**
 * Copyright 2003-2010 Terracotta, Inc.
 *
 * Licensed under the Apache License, Version 2.0 (the "License");
 * you may not use this file except in compliance with the License.
 * You may obtain a copy of the License at
 *
 *     http://www.apache.org/licenses/LICENSE-2.0
 *
 * Unless required by applicable law or agreed to in writing, software
 * distributed under the License is distributed on an "AS IS" BASIS,
 * WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
 * See the License for the specific language governing permissions and
 * limitations under the License.
```



Appendix H Threat Protection Rules

This document contains information about the system and auto threat protection rules for the Infoblox Advanced DNS Protection solution. All rules are grouped by rule categories. System and auto rule are automatically updated during rule updates.

You can create custom rules using rule templates. For information about Advanced DNS Protection and custom rule templates, refer to [Infoblox Advanced DNS Protection](#) on page 1213.

Information in this document is grouped by the following rule categories:

- [DNS Cache Poisoning](#) on page 1372
- [DNS Message Type](#) on page 1373
- [General DDoS](#) on page 1378
- [Reconnaissance](#) on page 1379
- [DNS Malware](#) on page 1379
- [DNS Protocol Anomalies](#) on page 1380
- [TCP/UDP Flood](#) on page 1381
- [DNS Tunneling](#) on page 1382
- [DNS Amplification and Reflection](#) on page 1383
- [NTP](#) on page 1384
- [BGP](#) on page 1385
- [OSPF](#) on page 1386
- [ICMP](#) on page 1387
- [Default Pass/Drop](#) on page 1392

DNS CACHE POISONING

DNS cache poisoning involves inserting a false address record for an Internet domain into a DNS query. If the DNS server accepts the record, subsequent requests for the address of the domain are answered with the address of a server controlled by the attacker. For as long as the false entry is cached, incoming web requests and emails will go to the attacker's address. New cache poisoning attacks, such as the "birthday paradox," use brute force, flooding DNS responses and queries at the same time, hoping to get a match on one of the responses and poison the cache. The following table lists auto rules that Advanced DNS Protection uses to mitigate DNS cache poisoning on your advanced appliance.

You can override the default values for the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.1 DNS Cache Poisoning Rules

Rule ID	Rule Type	Rule Name	Description
100000100	Auto	EARLY PASS UDP response traffic	This rule passes UDP DNS response packets (from upstream DNS servers or external DNS primaries) if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
100000200	AUTO	EARLY PASS TCP response traffic	This rule passes TCP DNS responses initiated by the appliance.
100000300	Auto	PASS ACK packets from NIOS initiated connections	This rule passes DNS ACK packets from NIOS initiated connections if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).

DNS MESSAGE TYPE

The following table lists the system and auto rules that are used to mitigate DNS message type attacks on your advanced appliance.

You can override the default values for the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

All rules for DNS record types are system rules. By default, they are configured as **Pass** rules. You can override this and change the rule action to **Drop**. Note that when you do that, the appliance drops all DNS packets that contain the requested record type.

Note: You cannot disable Auto rules.

Table H.2 DNS Message Type Rules

Rule ID	Rule Type	Rule Name	Usage
100100100	Auto	EARLY PASS UDP notify messages	This rule passes UDP DNS NOTIFY messages if the packet rate is less than the specified Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
100100200	Auto	EARLY PASS TCP notify messages	This rule passes TCP DNS NOTIFY messages if the packet rate is less than the specified Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130100100	Auto	RATELIMIT PASS UDP DNS AXFR zone transfer requests	This rule passes UDP DNS full zone transfer requests if the packet rate is less than the specified Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130100200	Auto	RATELIMIT PASS TCP DNS AXFR zone transfer requests	This rule passes TCP DNS full zone transfer requests if the packet rate is less than the specified Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130100300	Auto	RATELIMIT PASS UDP DNS IXFR zone Transfer requests	This rule passes UDP DNS incremental zone transfer requests if the packet rate is less than the specified Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).

Rule ID	Rule Type	Rule Name	Usage
130100400	Auto	RATELIMIT PASS TCP DNS IXFR zone Transfer requests	This rule passes TCP DNS incremental zone transfer requests if the packet rate is less than the specified Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130200100	Auto	DROP UDP DNS AXFR zone transfer requests	This rule drops any DNS UDP full zone transfer requests when zone transfer is disabled. You can configure only the Events per second parameter.
130200200	Auto	DROP TCP DNS AXFR zone transfer requests	This rule drops any DNS TCP full zone transfer requests when zone transfer is disabled. You can configure only the Events per second parameter.
130200300	Auto	DROP UDP DNS IXFR zone Transfer requests	This rule drops any DNS UDP incremental zone transfer requests when zone transfer is disabled. You can configure only the Events per second parameter.
130200400	Auto	DROP TCP DNS IXFR zone Transfer requests	This rule drops any DNS TCP incremental zone transfer requests when zone transfer is disabled. You can configure only the Events per second parameter.
130500100	System	DNS A record	You can configure this rule to pass or drop UDP packets when A record is requested in the DNS query. The default Action = Pass .
130500200	System	DNS AAAA record	You can configure this rule to pass or drop UDP packets when AAAA record is requested in the DNS query. The default Action = Pass .
130500300	System	DNS CNAME record	You can configure this rule to pass or drop UDP packets when CNAME record is requested in the DNS query. The default Action = Pass .
130500400	System	DNS DS record	You can configure this rule to pass or drop UDP packets when DS record is requested in the DNS query. The default Action = Pass .
130500500	System	DNS PTR record	You can configure this rule to pass or drop UDP packets when PTR record is requested in the DNS query. The default Action = Pass .
130500600	System	DNS NS record	You can configure this rule to pass or drop UDP packets when NS record is requested in the DNS query. The default Action = Pass .
130500700	System	DNS NSEC record	You can configure this rule to pass or drop UDP packets when NSEC record is requested in the DNS query. The default Action = Pass .
130500800	System	DNS NSEC3 record	You can configure this rule to pass or drop UDP packets when NSEC3 record is requested in the DNS query. The default Action = Pass .
130500900	System	DNS NSEC3PARAM record	You can configure this rule to pass or drop UDP packets when NSEC3PARAM record is requested in the DNS query. The default Action = Pass .

Rule ID	Rule Type	Rule Name	Usage
130501000	System	DNS MX record	You can configure this rule to pass or drop UDP packets when MX record is requested in the DNS query. The default Action = Pass.
130501100	System	DNS SRV record	You can configure this rule to pass or drop UDP packets when SRV record is requested in the DNS query. The default Action = Pass.
130501200	System	DNS TXT record	You can configure this rule to pass or drop UDP packets when TXT record is requested in the DNS query. The default Action = Pass.
130501300	System	DNS DNAME record	You can configure this rule to pass or drop UDP packets when DNAME record is requested in the DNS query. The default Action = Pass.
130501400	System	DNS RRSIG record	You can configure this rule to pass or drop UDP packets when RRSIG record is requested in the DNS query. The default Action = Pass.
130501500	System	DNS NAPTR record	You can configure this rule to pass or drop UDP packets when NAPTR record is requested in the DNS query. The default Action = Pass.
130501600	System	DNS DNSKEY record	You can configure this rule to pass or drop UDP packets when DNSKEY record is requested in the DNS query. The default Action = Pass.
130501700	System	DNS SPF record	You can configure this rule to pass or drop UDP packets when SPF record is requested in the DNS query. The default Action = Pass.
130501800	System	DNS DHCID record	You can configure this rule to pass or drop UDP packets when DHCID record is requested in the DNS query. The default Action = Pass.
130501900	System	DNS SOA record	You can configure this rule to pass or drop UDP packets when SOA record is requested in the DNS query. The default Action = Pass.
130502000	System	DNS SIG record	You can configure this rule to pass or drop UDP packets when SIG record is requested in the DNS query. The default Action = Pass.
130502100	System	DNS LOC record	You can configure this rule to pass or drop UDP packets when LOC record is requested in the DNS query. The default Action = Pass.
130502200	System	DNS SSHFP record	You can configure this rule to pass or drop UDP packets when SSHFP record is requested in the DNS query. The default Action = Pass.
130502300	System	DNS IPSECKEY record	You can configure this rule to pass or drop UDP packets when IPSECKEY record is requested in the DNS query. The default Action = Pass.
130502400	System	DNS TKEY record	You can configure this rule to pass or drop UDP packets when TKEY record is requested in the DNS query. The default Action = Pass.

Rule ID	Rule Type	Rule Name	Usage
130502500	System	DNS TSIG record	You can configure this rule to pass or drop UDP packets when TSIG record is requested in the DNS query. The default Action = Pass .
130502600	System	DNS TA record	You can configure this rule to pass or drop UDP packets when TA record is requested in the DNS query. The default Action = Pass .
130502700	System	DNS DLV record	You can configure this rule to pass or drop UDP packets when DLV record is requested in the DNS query. The default Action = Pass .
130502800	System	DNS ANY record	You can configure this rule to pass or drop UDP packets when ANY record is requested in the DNS query. The default Action = Pass .
130502900	System	DNS A record TCP	You can configure this rule to pass or drop TCP packets when A record is requested in the DNS query. The default Action = Pass .
130503000	System	DNS AAAA record TCP	You can configure this rule to pass or drop TCP packets when AAAA record is requested in the DNS query. The default Action = Pass .
130503100	System	DNS CNAME record TCP	You can configure this rule to pass or drop TCP packets when CNAME record is requested in the DNS query. The default Action = Pass .
130503200	System	DNS DS record TCP	You can configure this rule to pass or drop TCP packets when DS record is requested in the DNS query. The default Action = Pass .
130503300	System	DNS PTR record TCP	You can configure this rule to pass or drop TCP packets when PTR record is requested in the DNS query. The default Action = Pass .
130503400	System	DNS NS record TCP	You can configure this rule to pass or drop TCP packets when NS record is requested in the DNS query. The default Action = Pass .
130503500	System	DNS NSEC record TCP	You can configure this rule to pass or drop TCP packets when NSEC record is requested in the DNS query. The default Action = Pass .
130503600	System	DNS NSEC3 record TCP	You can configure this rule to pass or drop TCP packets when NSEC3 record is requested in the DNS query. The default Action = Pass .
130503700	System	DNS NSEC3PARAM record TCP	You can configure this rule to pass or drop TCP packets when NSEC3PARAM record is requested in the DNS query. The default Action = Pass .
130503800	System	DNS MX record TCP	You can configure this rule to pass or drop TCP packets when MX record is requested in the DNS query. The default Action = Pass .
130503900	System	DNS SRV record TCP	You can configure this rule to pass or drop TCP packets when SRV record is requested in the DNS query. The default Action = Pass .

Rule ID	Rule Type	Rule Name	Usage
130504000	System	DNS TXT record TCP	You can configure this rule to pass or drop TCP packets when TXT record is requested in the DNS query. The default Action = Pass.
130504100	System	DNS DNAME record TCP	You can configure this rule to pass or drop TCP packets when DNAME record is requested in the DNS query. The default Action = Pass.
130504200	System	DNS RRSIG record TCP	You can configure this rule to pass or drop TCP packets when RRSIG record is requested in the DNS query. The default Action = Pass.
130504300	System	DNS NAPTR record TCP	You can configure this rule to pass or drop TCP packets when NAPTR record is requested in the DNS query. The default Action = Pass.
130504400	System	DNS DNSKEY record TCP	You can configure this rule to pass or drop TCP packets when DNSKEY record is requested in the DNS query. The default Action = Pass.
130504500	System	DNS SPF record TCP	You can configure this rule to pass or drop TCP packets when SPF record is requested in the DNS query. The default Action = Pass.
130504600	System	DNS DHCID record TCP	You can configure this rule to pass or drop TCP packets when DHCID record is requested in the DNS query. The default Action = Pass.
130504700	System	DNS SOA record TCP	You can configure this rule to pass or drop TCP packets when SOA record is requested in the DNS query. The default Action = Pass.
130504800	System	DNS SIG record TCP	You can configure this rule to pass or drop TCP packets when SIG record is requested in the DNS query. The default Action = Pass.
130504900	System	DNS ROC record TCP	You can configure this rule to pass or drop TCP packets when ROC record is requested in the DNS query. The default Action = Pass.
130505000	System	DNS SSHFP record TCP	You can configure this rule to pass or drop TCP packets when SSHFP record is requested in the DNS query. The default Action = Pass.
130505100	System	DNS IPSECKEY record TCP	You can configure this rule to pass or drop TCP packets when IPSECKEY record is requested in the DNS query. The default Action = Pass.
130505200	System	DNS TKEY record TCP	You can configure this rule to pass or drop TCP packets when TKEY record is requested in the DNS query. The default Action = Pass.
130505300	System	DNS TSIG record TCP	You can configure this rule to pass or drop TCP packets when TSIG record is requested in the DNS query. The default Action = Pass.
130505400	System	DNS TA record TCP	You can configure this rule to pass or drop TCP packets when TA record is requested in the DNS query. The default Action = Pass.

Rule ID	Rule Type	Rule Name	Usage
130505500	System	DNS DLV record TCP	You can configure this rule to pass or drop TCP packets when DLV record is requested in the DNS query. The default Action = Pass .
130505600	System	DNS ANY record TCP	You can configure this rule to pass or drop TCP packets when ANY record is requested in the DNS query. The default Action = Pass .

GENERAL DDoS

The following table lists the auto rules that are used to mitigate general DDoS attacks on your advanced appliance. You can configure the following rule parameter for all rules in this category:

- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.3 General DDoS Rules

Rule ID	Rule Type	Rule Name	Description
110000100	Auto	EARLY DROP DoS packets with same source and destination IP	This rule drops any IP packets that contain the same source IP and destination IP.
110000200	Auto	EARLY DROP DoS UDP packets with same source and destination IP	This rule drops UDP packets that contain the same source and destination IP.
110000300	Auto	EARLY DROP DoS TCP packets with same source and destination IP	This rule drops TCP packets that contain the same source and destination IP.
130400300	Auto	DROP IPv6 loopback address spoofing	This rule blocks any IP packets that attempt to forge the IPv6 loopback address.
130400400	Auto	DROP IPv6 loopback address spoofing	This rule blocks any IP packets that attempt to forge the IPv6 loopback address.

RECONNAISSANCE

Reconnaissance attacks consist of attempts to get information on the network environment before launching a large DDoS or other types of attacks. Techniques include port scanning and finding versions and authors. These attacks exhibit abnormal behavior patterns that, if identified, can provide early warnings.

The following table lists the auto rules that are used to mitigate reconnaissance attacks on your advanced appliance.

You can configure the following rule parameter for all rules in this category:

- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.4 Reconnaissance Rules

Rule ID	Rule Type	Rule Name	Description
110100100	Auto	EARLY DROP DNS named author attempts	This rule drops UDP packets that contain attempts to find information about authors.
110100200	Auto	EARLY DROP DNS named version attempts	This rule drops UDP packets that contain attempts to find version information.

DNS MALWARE

DNS malware is software used to disrupt your DNS service, gather sensitive information, or gain access to your appliance. It can include downloaders, backdoors, trojan horses, and other malicious software.

The following table lists the auto rules that are used to mitigate DNS malware on your advanced appliance.

You can configure the following rule parameter for all rules in this category:

- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.5 DNS Malware Rules

Rule ID	Rule Type	Rule Name	Description
110100300	Auto	EARLY DROP UDP MALWARE backdoor	This rule drops packets against the backdoor malware BKDR_QUEJOB.EVL, which poses as an installer of FaceBook messenger. This malware may be spread as a malicious attachment in email messages.
130300300	Auto	DROP MALWARE trojan downloader	This rule drops packets against trojan downloaders, which download and install new versions of malicious programs, including Trojans and AdWare, on your appliance. Once downloaded, the programs are launched or included on a list of programs that run automatically when your system starts.

Rule ID	Rule Type	Rule Name	Description
130300400	Auto	DROP MALWARE possible Hiloti	This rule drops packets against trojan Hiloti malicious programs that may download potentially malicious files from a remote server and report system information back to the server. Once downloaded, the programs are launched or included on a list of programs that run automatically when your system starts.

DNS PROTOCOL ANOMALIES

DNS protocol anomalies send malformed DNS packets, including unexpected header and payload values, to the targeted server. This causes the server to stop responding or crash, which results in an infinite loop in server threads. These anomalies sometimes take the form of impersonation attacks.

The following table lists rules that are used to mitigate DNS protocol anomalies on your advanced appliance.

You can configure the following rule parameter for all rules in this category:

- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.6 DNS Protocol Anomalies Rules

Rule ID	Rule Type	Rule Name	Description
110100400	Auto	EARLY DROP UDP DNS question name too long	This rule drops UDP DNS packets when the DNS Question Name is too long.
110100500	Auto	EARLY DROP UDP DNS label too long	This rule drops UDP DNS packets when the DNS Label in the name being queried is too long.
110100600	Auto	EARLY DROP UDP query invalid question count	This rule drops UDP DNS packets when the number of entries in the question section is invalid.
110100700	Auto	EARLY DROP UDP query invalid question class	This rule drops UDP DNS packets when the RR (resource record) class being queried is invalid.
110100800	Auto	EARLY DROP UDP query invalid question string	This rule drops UDP DNS packets that contain invalid question string.
110100900	Auto	EARLY DROP query multiple questions	This rule drops UDP DNS packets when there are multiple questions being queried at one time.
130000700	Auto	EARLY DROP TCP non-DNS query	This rule drops TCP packets that do not contain valid DNS queries.
130000800	Auto	EARLY DROP TCP query multiple questions	This rule drops TCP DNS packets when there are multiple questions being queried at one time.
130100500	Auto	DROP UDP DNS invalid IXFR query with zero or more than one Authority	This rule drops UDP DNS packets when incremental zone transfer requests contain zero or more than one Authority.

Rule ID	Rule Type	Rule Name	Description
130100600	Auto	DROP TCP DNS invalid IXFR query with zero or more than one Authority	This rule drops TCP DNS packets when incremental zone transfer requests contain zero or more than one Authority.
130300100	Auto	DROP UDP invalid DNS query with Authority	This rule drops UDP DNS packets when DNS queries contain invalid Authority.
130300200	Auto	DROP TCP invalid DNS query with Authority	This rule drops TCP DNS packets when DNS queries contain invalid Authority.
130300500	Auto	DROP TCP non-DNS or non-compliant DNS traffic (AD bit set)	This rule drops TCP DNS packets that contain invalid or non-compliant DNS traffic.

TCP/UDP Flood

TCP and UDP flood attacks are volumetric attacks with massive numbers of packets that consume network bandwidth and resources. They exploit TCP and UDP.

The following table lists the system and auto rules that are used to mitigate TCP/UDP floods on your advanced appliance.

Depending on the rules, you can override the default values for all or some of the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.7 TCP/UDP Flood Rules

Rule ID	Rule Type	Rule Name	Description
130000100	System	WARN & BLOCK high rate inbound UDP DNS queries	This rule warns and then blocks large amount of inbound UDP DNS queries if the packet rate is less than the Packets per second value (default = 40). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP.
130000200	System	RATELIMIT PASS high rate inbound UDP DNS queries	This rule passes large amount of inbound UDP DNS queries if the packet rate is less than the Packets per second value (default = 60). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 5).
130000300	System	WARN & BLOCK TCP high rate inbound connection attempts	This rule warns and then blocks large amount of inbound TCP DNS queries if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP.

Rule ID	Rule Type	Rule Name	Description
130000400	Auto	RATELIMIT PASS TCP high rate inbound connection attempts	This rule passes large amount of inbound TCP DNS queries if the packet rate is less than the Packets per second value (default = 60). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130000500	Auto	RATELIMIT PASS UDP high rate inbound large DNS queries (anti-tunneling)	This rule passes large amount of inbound UDP DNS queries if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130000600	Auto	RATELIMIT PASS TCP high rate inbound large DNS queries (anti-tunneling)	This rule passes large amount of inbound TCP DNS queries if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130000600	Auto	RATELIMIT PASS TCP high rate inbound large DNS queries (anti-tunneling)	This rule passes large amount of inbound TCP DNS queries if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).

DNS TUNNELING

DNS tunneling attacks involve tunneling another protocol through DNS port 53 for the purpose of data exfiltration. Outbound and inbound data being communicated is encoded into small chunks and fitted into DNS queries and DNS responses.

The following table lists the system rule used to mitigate DNS tunneling on your advanced appliance.

You can override the default values for the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The drop period in seconds.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.
- **Packet size:** DNS packet size. If the DNS packet size exceeds a certain value, the corresponding rule will be triggered. The default is 200.

Table H.8 Anti DNS Tunneling Rules

Rule ID	Rule Type	Rule Name	Description
130000500	System	RATELIMIT PASS UDP high rate inbound large DNS queries (anti tunneling)	This rule passes large amount of inbound UDP DNS queries (attempted for DNS tunneling) if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30). This rule is triggered when the DNS Packet size exceeds a certain value (default = 200).
130000600	Auto	RATELIMIT PASS TCP high rate inbound large DNS queries (anti-tunneling)	This rule passes large amount of inbound TCP DNS queries (attempted for DNS tunneling) if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30). This rule is triggered when the DNS Packet size exceeds a certain value (default = 200).

DNS AMPLIFICATION AND REFLECTION

DNS reflection attacks use a form of IP spoofing, changing the source address in their DNS queries to show the address of their intended target, such as a DNS root server or a top-level domain (TLD) name server operator. DNS reflection and amplification recognizes UDP as an asymmetrical protocol (small requests, large responses) and the existence of open DNS resolvers to the Internet cloud. The result is that small DNS queries reflect large UDP datagram responses to the target address in the original source datagrams. Some recent attacks have used this DDoS technique at a huge scale.

Since DNS runs over UDP and does not require a handshake, it is possible to use the protocol as a means to lock down a host or a network. Designed a specific way, sending a small query to any open DNS resolver can result in a single response containing several kilobytes or more, that are sent to the unwitting spoofed victim. (This type of response typically is sent via TCP, as UDP does not allow for more than 512 bytes in a response datagram. The resulting packet usually exceeds the MTU of the recipient's interfaces, resulting in further packet fragmentation and processing.) Open DNS resolvers may allow for launching DDoS attacks containing hundreds of gigabytes of data. Attackers may also use the EDNS0 DNS protocol extension as a means to enable larger DNS responses. Many network operators, particularly overseas, allow open DNS resolvers to run on their networks, unwittingly allowing attackers to abuse them. Many network operators do provide intelligent rate-limiting to prevent abuse, even while supporting open recursive DNS servers. Hence, issues of this type usually result from mistakes in configuration.

The following table lists the system and auto rules that are used to mitigate DNS amplification and reflection attacks on your advanced appliance.

Depending on the rules, you can override the default values for all or some of the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.9 DNS Amplification and Reflection Rules

Rule ID	Rule Type	Rule Name	Description
130400100	Auto	WARN & DROP DoS DNS possible reflection/ amplification attack attempts	This rule warns and then blocks possible DoS attacks if any source IP sends packets over the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 5).
130400500	System	RATELIMIT PASS UDP DNS root requests with additional RRs	This rule passes large amount of inbound UDP DNS queries if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130400600	System	RATELIMIT PASS UDP DNS root requests	This rule passes large amount of inbound UDP DNS queries if the packet rate is less than the Packets per second value (default = 5). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).

NTP

The following table lists the rule used to mitigate NTP attacks on your advanced appliance when NTP is enabled.

You can override the default values for the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.10 NTP Rules

Rule ID	Rule Type	Rule Name	Description
130600100	Auto	RATELIMIT PASS NTP responses	This rule passes NTP responses when NTP is enabled and if the packet rate is less than the Packets per second value (default = 10). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 15).

BGP

The following table lists the auto rules that are used to mitigate BGP attacks on your advanced appliance when BGP is enabled.

You can override the default values for the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.11 BGP Rules

Rule ID	Rule Type	Rule Name	Description
130700100	AUTO	DROP BGP header length shorter than spec	When BGP is enabled, this rule drops TCP BGP packets that contain message header length that is shorter than the RFC specification.
130700200	AUTO	DROP BGP header length longer than spec	When BGP is enabled, this rule drops TCP BGP packets that contain message header length that is longer than the RFC specification.
130700300	AUTO	DROP BGP spoofed connection reset attempts	When BGP is enabled, this rule drops TCP BGP packets that contain spoofed connection reset.
130700400	AUTO	DROP BGP invalid type 0	When BGP is enabled, this rule drops TCP BGP packets that contain invalid message type 0.
130700500	AUTO	DROP BGP invalid type bigger than 5	When BGP is enabled, this rule drops TCP BGP packets that contain invalid message type greater than 5.
130700550	AUTO	RATELIMIT PASS BGP IPv4 peer TCP connection attempts	This rule passes TCP BGP route advertisement connection attempts from IPv4 peers when BGP is enabled and if the packet rate is less than the Packets per second value (default = 10). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).
130700600	Auto	RATELIMIT PASS BGP allowed with IPv4 peer	This rule passes TCP BGP route advertisement to IPv4 peers when BGP is enabled and if the packet rate is less than the Packets per second value (default = 10). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).

Rule ID	Rule Type	Rule Name	Description
130700650	AUTO	RATELIMIT PASS BGP IPv6 peer TCP connection attempts	This rule passes TCP BGP route advertisement connection attempts from IPv6 peers when BGP is enabled and if the packet rate is less than the Packets per second value (default = 10). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).
130700700	Auto	RATELIMIT PASS BGP allowed with IPv6 peer	This rule passes TCP BGP route advertisement to IPv6 peers when BGP is enabled and if the packet rate is less than the Packets per second value (default = 10). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).
130800100	Auto	DROP BGP unexpected	When BGP is enabled, this rule drops unexpected TCP BGP packets.

OSPF

The following table lists auto rules that are used to mitigate OSPF attacks on your advanced appliance when OSPF is not in use.

Depending on the rules, you can override the default values for one or more of the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.12 OSPF Rules

Rule ID	Rule Type	Rule Name	Description
130900300	Auto	DROP OSPF unexpected	This rule drops unexpected OSPF packets.
130900400	Auto	RATELIMIT PASS OSPF multicast	This rule passes OSPF IPv4 multicast packets if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).
130900500	Auto	RATELIMIT PASS OSPF IPv6 multicast	This rule passes OSPF IPv6 multicast packets if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).

Rule ID	Rule Type	Rule Name	Description
130900600	Auto	RATELIMIT PASS OSPF	This rule passes OSPF packets if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).

ICMP

ICMP attacks use network devices such as routers to send error messages when a requested service is not available or the remote server cannot be reached. Examples of ICMP attacks include ping floods, ping-of-death attacks, and smurf attacks.

The following table lists the system and auto rules that are used to mitigate ICMP attacks on your advanced appliance. Depending on the rules, you can override the default values for one or more of the following rule parameters:

- **Packets per second:** The number of packets per second that the appliance processes before it blocks the traffic.
- **Drop interval:** The time period (in seconds) for which the appliance blocks all traffic beyond the rate limit value.
- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Note: You cannot disable Auto rules.

Table H.13 ICMP Rules

Rule ID		Rule Name	Description
130400200	Auto	DROP ICMP large packets	This rule drops large ICMP large packets.
130900100	Auto	ICMP Ping	This rule passes ICMP ping packets if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130900200	Auto	ICMPv6 Ping	This rule passes ICMPv6 ping packets if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130900200	Auto	DROP OSPF unexpected	This rule drops any unexpected OSPF packets when OSPF is disabled.
130900300	Auto	DROP OSPF unexpected	This rule drops any unexpected OSPF packets when OSPF is disabled.
130900400	Auto	RATELIMIT PASS OSPF multicast	This rule passes OSPF multicast packets if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 60).

Rule ID		Rule Name	Description
130900500	Auto	RATELIMIT PASS OSPF IPv6 multicast	This rule passes OSPF IPv6 multicast packets if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval; default = 60).
130900600	Auto	RATELIMIT PASS OSPF	This rule passes OSPF packets if the packet rate is less than the Packets per second value (default = 50). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130900700	Auto	RATELIMIT PASS ICMPv6 destination unreachable	This rule passes ICMPv6 Destination Unreachable messages if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130900800	Auto	RATELIMIT PASS ICMPv6 packet too big	This rule passes ICMPv6 Packet Too Big messages if the packet rate is less than the Packets per second value (default = 100). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130900900	Auto	RATELIMIT PASS ICMPv6 ping responses	This rule passes ICMPv6 ping responses if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130901000	Auto	RATELIMIT PASS ICMPv6 parameter problem erroneous header	This rule passes ICMPv6 Erroneous Header messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130901100	Auto	RATELIMIT PASS ICMPv6 parameter problem unrecognized next header	This rule passes ICMPv6 Unrecognized Next Header messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901200	Auto	RATELIMIT PASS ICMPv6 parameter problem unrecognized IPv6 option	This rule passes ICMPv6 Unrecognized IPv6 Option messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).

Rule ID		Rule Name	Description
130901300	Auto	RATELIMIT PASS ICMPv6 router solicitation	This rule passes ICMPv6 router solicitation packets if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901400	Auto	RATELIMIT PASS ICMPv6 router advertisement	This rule passes ICMPv6 router advertisement if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901500	Auto	RATELIMIT PASS ICMPv6 neighbor solicitation	This rule passes ICMPv6 neighbor solicitation packets if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901600	Auto	RATELIMIT PASS ICMPv6 neighbor advertisement	This rule passes ICMPv6 neighbor advertisement if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901700	Auto	RATELIMIT PASS ICMPv6 inverse neighbor solicitation	This rule passes ICMPv6 inverse neighbor solicitation messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901800	Auto	RATELIMIT PASS ICMPv6 inverse neighbor advertisement	This rule passes ICMPv6 inverse neighbor advertisement if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130901900	Auto	RATELIMIT PASS ICMPv6 listener query	This rule passes ICMPv6 listener query messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130902000	Auto	RATELIMIT PASS ICMPv6 listener report	This rule passes ICMPv6 listener report messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).

Rule ID		Rule Name	Description
130902100	Auto	RATELIMIT PASS ICMPv6 listener done	This rule passes ICMPv6 listener done messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130902200	Auto	RATELIMIT PASS ICMPv6 listener report v2	This rule passes ICMPv6 listener report v2 messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130902300	Auto	RATELIMIT PASS ICMPV6 multicast router advertisement	This rule passes ICMPv6 multicast router advertisement if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130902400	Auto	RATELIMIT PASS ICMPV6 multicast router solicitation	This rule passes ICMPv6 multicast router solicitation messages if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130902500	Auto	RATELIMIT PASS ICMPV6 multicast router advertisement	This rule passes ICMPv6 multicast router advertisement if the packet rate is less than the Packets per second value (default = 8). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 30).
130902600	Auto	RATELIMIT PASS ICMP ping responses	This rule passes ICMP ping responses if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130902700	Auto	RATELIMIT PASS ICMP router advertisement	This rule passes ICMP router advertisement if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130902800	Auto	RATELIMIT PASS ICMP router solicitation	This rule passes ICMP router solicitation messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).

Rule ID		Rule Name	Description
130902900	Auto	RATELIMIT PASS ICMP time exceeded	This rule passes ICMP time exceeded messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130903000	Auto	RATELIMIT PASS ICMP parameter problems	This rule passes ICMP parameter problems if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130903100	Auto	RATELIMIT PASS ICMPv6 hop limit exceeded or ICMPv4 network unreachable	This rule passes ICMPv6 Hop Limit Exceeded messages or ICMPv4 Network Unreachable messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130903200	Auto	RATELIMIT PASS ICMPv6 fragment reassembly time exceeded or ICMPv4 host unreachable	This rule passes ICMPv6 fragment reassembly time exceeded messages or ICMPv4 host unreachable messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130903300	Auto	RATELIMIT PASS ICMP protocol unreachable	This rule passes ICMP protocol unreachable messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130903400	Auto	RATELIMIT ICMP port unreachable	This rule passes ICMP port unreachable messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).
130903500	Auto	RATELIMIT PASS ICMP fragmentation needed	This rule passes ICMP fragmentation needed messages if the packet rate is less than the Packets per second value (default = 20). If any source IP sends packets over this value, the appliance blocks all such traffic from this source IP for a certain period of time (specified in Drop interval ; default = 10).

DEFAULT PASS/DROP

The following table lists the system rules that are used to pass or drop packets on your advanced appliance. All rules are disabled by default. You can enable them and configure the following, depending on the rules:

- **Events per second:** The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. The default value is 10.

Table H.14 Default Pass/Drop Rules

Rule ID	Rule Type	Rule Name	Description
100000050	System	EARLY PASS TCP with flowbits set	This rule passes TCP traffic that has the flowbits options set and marked OK.
140000100	System	DROP UDP DNS unexpected	This rule drops any unexpected UDP DNS packets.
140000200	System	DROP TCP DNS unexpected	This rule drops any unexpected TCP DNS packets.
140000400	System	PASS TCP established packets	This passes all TCP established packets.
140000500	System	DROP TCP unexpected	This rule drops any unexpected TCP packets.
140000600	System	DROP UDP unexpected	This rule drops any unexpected UDP packets.
140000700	System	DROP ICMP unexpected	This rule drops any unexpected ICMP packets.
140000800	System	DROP unexpected protocol	This rule drops any unexpected protocol packets.

Index

A

- A records
 - adding to shared record groups 683, 685
 - bulk hosts 656
 - DDNS updates 703
 - host records 460
 - sort lists 588
- AAAA records
 - adding 662, 1239, 1240, 1241
 - adding to shared record groups 683, 685
 - sort lists 588
- access control lists
 - for TFTP, HTTP, FTP services 394
- Active Directory
 - authenticating admins 177, 182
 - configuring support for 709
- Adding Widgets to Dashboards 117
- admin groups 154
 - ALL USERS group 154
 - configuring on remote RADIUS server 174
 - defining permissions 160
 - limited-access 154
 - managing 158
 - superusers 154
- admin roles 157
- admins
 - authenticating 169
 - authenticating using RADIUS 173
 - defining admin policy 185
 - defining user profiles 50
 - managing 149
 - notifying 191
 - password length 191
 - using Active Directory to authenticate 177, 182
- anycast 761
- API
 - migrating data 301
- appliance status
 - viewing 1005, 1010
- audit history
 - IP addresses 478
- audit log
 - configuring 1018
 - selecting type 268
 - sending to syslog 1014

B

- backing up
 - guidelines for scheduled tasks 74
- backup files
 - creating and restoring 423
- BGP 767
- blackhole list
 - configuring for DNS 590
- Blacklists 579
- bloxTools environment 431
 - viewing status 1011
- bookmarks 63

BOOTP

- specifying parameters 798

browsers

- limitations 47
- setting time zone 51
- supported 46

bulk host records 656

- admin permissions 202

C

- capacity report 1022

- capture traffic 1021

certificates

- generating a certificate signing request 54
- generating a certificate signing request for a captive portal 926
- generating self-signed 53
- generating self-signed for captive portal 926
- importing 54
- importing for Captive Portal 927

CLI

- defining network settings 277

CNAME records

- adding 669
- host records 460

code pages

- Microsoft Windows 92

columns

- editing 60

- Configuring 1183

- Configuring Concurrent Zone Transfers 562

- Configuring Recursive Deletions of Networks and Zones 269

- CPU temperature status 1008

- CPU usage status 1008

CSV Import

- blacklist rules 580

- CSV import 86

D

- Dashboard 48, ??–131

Data Import Wizard

- migrating data 301

database replication

- bandwidth considerations 231

- database status 1006

DDNS 689–731

- configuring DHCP 695–705

- configuring DNS 705–708

- DDoS Protection Statistics 138

- delegated zones 638

- detailed status 1004

- bloxTools Environment 436

- NTP status icons 313

- viewing 1004, 1116

DHCP

- configuration checklist 780

- defining admin permissions 205

- dynamic DNS updates 691
 - enabling LAN2 port 358
 - general properties 793, 1032
 - inheriting properties 782
 - NAC Foundation module 917
 - starting and stopping service 822
 - viewing status 1011
 - DHCP failover
 - configuring 884–890
 - Dashboard 125
 - DHCP filters 892
 - DHCP lease history
 - defining admin permissions 213
 - enabling 815
 - DHCP leases
 - clearing active 491
 - converting 488
 - scavenging 794
 - DHCP option filters 901
 - DHCP options
 - option 12 695
 - option 15 695
 - option 60 805
 - option 81 701
 - DHCP ranges
 - applying MAC address filters 907
 - configuring 854, 876
 - Dashboard 125
 - defining admin permissions 198, 207, 210
 - permission to create from template 211
 - using templates 826
 - DHCP templates
 - configuring 826–833
 - defining admin permissions 211
 - DHCPv6 options 810
 - DHCPv6 RFC Compliance 1316
 - disk usage status 1005
 - DNAME records
 - adding 671
 - DNS
 - configuration checklist 549
 - defining admin permissions 199
 - enabling LAN2 port 358
 - IPv6 552
 - logging categories 1015
 - root name servers 587
 - specifying DNS resolvers 376
 - starting and stopping service 565
 - updates for a zone 697, 703
 - using MGMT port 364
 - viewing status 1011
 - DNS anycast
 - BGP 767
 - DNS monitoring 1024
 - DNS Query Trend per IP Block 1156, 1162
 - DNS statistics
 - Dashboard 124
 - DNS views 602–614
 - defining admin permissions 200
 - read-only permission 198, 207
 - DNS64 594
 - DNSKEY resource record 735
 - DNSSEC 734
 - downgrading NIOS software 422
 - DS resource record 740
- ## E
- Ethernet ports 346
 - exporting data 89, 91
 - capacity report 1022
 - from smart folders 145
 - lease records 950
 - extensible attributes
 - 322–333
 - in wizards and editors 332
 - smart folders 140
 - UTF-8 encoding 92
- ## F
- fan status 1007
 - file distribution services
 - status 131
 - TFTP, HTTP, FTP 391
 - filters
 - DHCP 892
 - in global search 69
 - search criteria 67
 - using for smart folders 140
 - viewing permissions list 168
 - Finder panel 48, 59
 - fixed addresses
 - configuring 857, 878
 - defining admin permissions 198, 207, 208
 - displaying lease information 814
 - DNS updates 698
 - host records 460
 - permission to create from template 211
 - using templates 828
 - force restore 429
 - forward zones 640
 - forwarders 569
 - FTP
 - backup files 423
 - configuring 393
 - defining admin permissions 214
 - uploading files 399
 - UTF-8 encoding for file names 92
 - viewing status 1011
- ## G
- global search 69
 - exporting results 91
 - Grid
 - configuring 221–271
 - configuring NTP 316
 - configuring syslog 1013
 - configuring upgrade groups 408
 - DNS updates for DHCP 696
 - downgrading 422
 - monitoring services 1011
 - NAT groups 226
 - promoting master candidate 270
 - removing a member 270

- restarting services 386
- restricting access 344
- setting date, time, time zone 312
- Grid Manager
 - overview 48
- grid master 224
 - configuring 236
 - promoting candidate 270
 - scheduled tasks after promoting 74
- grid members
 - adding to grid 245
 - configuring NTP 318
 - configuring syslog 1014
 - defining admin permissions 195
 - distribution and upgrade status 421
 - logging DHCP leases 816
 - monitoring services 1012
 - permissions 201, 215
 - read-only permission 198, 207
 - restarting services 387
 - setting date, time, time zone 312
- grid status
 - distribution and upgrade 421
 - viewing detailed 1004, 1130
 - widget in Dashboard 121
- GSS-TSIG
 - authenticating DDNS updates 710
- GTM Group Status Information 1192
- GTM Synchronization Groups 1186
- Guidelines for the Next Available Network and IP Address 844

H

- hardware and software requirements 46
- hardware status 1010
- Help panel 59
- high availability (HA) pair
 - configuring grid master 238
 - deploying independent 287
 - grid members 247
 - rebooting 380
 - VRRP advertisements 235
- host records
 - adding IP addresses 479
 - admin permissions 202
 - bulk 656
 - configuring 459
 - hostname restrictions 592
- Hostname Compliance Report 594
 - admin permissions 202
- hostnames
 - for DNS updates 698
 - restrictions 592
- HSM signing 750
- HTTP
 - configuring 394
 - defining admin permissions 214
 - redirecting to HTTPS 344
 - restricting access to appliance 344
 - uploading files 399
 - viewing status 1011
- HTTPS
 - defined 52

I

- ICMP
 - network discovery 497
- ICMP echo requests
 - ping 795
- IDNA RFC Compliance 1316
- IF-MAP
 - Infoblox DHCP server 817
- Ignoring DHCP Client Identifiers 797
- importing CSV files 86
- importing zone data 631
- independent appliances 273
- Infoblox-1552, -1552-A
 - power supply status 1007
- Infoblox-1852-A
 - power supply status 1007
- Infoblox-2000, -2000-A
 - power supply status 1007
 - RAID status 1009
- Informational GUI Banner 269
- IP Address Management panel
 - admin permissions 198, 207
- IP addresses
 - converting associated objects 487
 - reclaiming associated objects 491
 - viewing IP Map 474
- IP Map 48, 474
- IPv6
 - adding AAAA records 662, 1239, 1240, 1241
 - adding PTR records 664
 - configuring DNS 552
 - configuring for host records 463
 - configuring forward reverse-mapping zone 642
 - DNS views 610
 - managing networks and addresses 480
 - reverse-mapping zone 619
 - stub reverse-mapping zone 650

J

- joining networks 473

L

- LAN2 ports
 - configuring 355
 - viewing status 1006
- languages
 - supporting multiple 92
- LCD
 - configure network settings 277
 - disabling 344
 - viewing status 1006
- licenses
 - activating 377
 - managing 377
 - removing 379, 382
- lite upgrades 406
- local admins 169
- logging
 - configuring member 815
- logging in
 - creating a banner 49

- login options 49
 - steps 48
- login banner 49
 - UTF-8 encoding 92
- logs
 - audit log 1018
 - DHCP 815
 - DHCP and DNS data 300
 - DHCP logs 815
 - DNS logging 1015
 - Microsoft 965
 - replication 1020
 - syslog 1012
 - traffic capture 1021
- loopback interface
 - captive portal IP address 924

M

- MAC address filters 897–899
 - defining admin permissions 212, 214
 - using in NAC Foundation module 917
- MAC addresses
 - formats 857
- master candidate 270
- member status
 - Dashboard 122, 123
 - viewing detailed 1004
- memory usage status 1006, 1007
- MGMT port
 - static routes 374
 - using 359
 - viewing status 1006
- MIBs
 - SNMP 1048
- Microsoft servers
 - configuring managing members 957
 - supported Windows versions 438, 955
 - synchronizing DHCP data 983
 - synchronizing DNS data 967, 1193
- Microsoft superscopes 995
- Modifying GTM Group Properties 1188
- Monitoring Client Queries 1128
- MTU
 - for VPN tunnels 270
- multi-ping 491
- MX records
 - adding 665
 - adding to shared record groups 683, 685

N

- name server groups 629–630
- NAT groups 226
- NetBIOS
 - network discovery 497, 524
- network discovery
 - 493–506, 517–??
 - Dashboard 128, 129
 - defining permissions 196, 495, 521
 - managing discovered data 511
 - PortIQ appliances 508
 - viewing discovered data 510
- network list 470, 484
- network map 48, 467
- network statistics

- Dashboard 127
- network views
 - configuring 787
 - displayed in bookmarks 63
 - DNS views 604
- networks
 - adding from Net Map 469, 483
 - configuring and managing 551, 843
 - Dashboard 128
 - defining admin permissions 207
 - joining 473
 - permission to create from template 211
 - resizing 472
 - splitting 472
 - using templates 829
- NIOS appliance
 - managing licenses 377
 - rebooting 380
 - resetting 381
 - restricting HTTP access 344
 - shutting down 380
- NIOS GUI
 - setting timeout 344
- NIOS software
 - downgrading 422
 - reverting 423
 - upgrading 411
- NIOS version
 - upgrading and downgrading 229
 - viewing 411
- NS records 662
- NSEC resource record 738
- NSEC3 resource record 738
- NSEC3PARAM resource record 739
- NTP
 - adding an authentication key 317
 - appliance as client 315
 - appliance as NTP server 319
 - authenticating 314
 - independent appliance 289
 - monitoring status 122, 313
 - using for appliances 313
 - viewing status 1008
- NXDOMAIN redirection 574

P

- passwords
 - length for admins 191
 - setting in user profile 50
- permissions
 - applying 166
 - Dashboard widgets 99
 - defining for admin groups 160
 - defining for DHCP 205
 - defining for DNS 199
 - defining for objects 162
 - defining for TFTP, HTTP, and FTP services 214
 - defining global 161, 164
 - for common tasks 193
 - for grid members 195
 - for network discovery 196
 - overview 160
- phone home 1023
- ping
 - configuring settings 795
 - IP addresses 491

- multi-ping 491
- PortIQ appliances 508
- ports
 - configuring VPN 267
 - DNS settings 562
 - Ethernet and service ports 346–354
 - for grid communications 238
 - LAN2 355
 - MGMT 359
 - source for DNS messages 359
 - viewing status 1006
- primary servers
 - configuring 623
 - name server groups 629
 - stealth 624
- printing data 91
 - capacity report 1022
 - in smart folders 145
- PTR records
 - adding 664
 - bulk hosts 656
 - host records 460

Q

- queries
 - controlling 570
 - recursive 571

R

- RADIUS
 - authenticating admins 173
 - authenticating DHCP clients 937
- RAID
 - monitoring status 1009
- rebooting 380
- recursion
 - configuring 571
 - DNS views 605
 - DNSSEC 754
- recycle bin 63, 64
 - enabling 267
 - restoring zone data 636
- relay agent filters 899–901
- remote admins
 - authenticating 171
 - authenticating using Active Directory 177, 182
 - authenticating using RADIUS 173
 - creating admin groups for 185
- remote console access
 - enabling and disabling 343
- Removing a GTM Group 1190
- replication status
 - viewing 1020
- Requirements and Permissions 1181
- reservations
 - configuring 860
 - defining admin permissions 208
 - using templates 828
- resetting NIOS appliances 381
- resizing networks 472
- resource records
 - adding 660
 - adding to shared record groups 683, 685
 - bulk host records 656

- defining permissions 201, 202, 215
- DNSSEC 735
 - specifying TTL 557
- restarting services 386
- restoring
 - backup files 428
 - using Recycle Bin 64
- restoring database
 - guidelines for scheduled tasks 74
- reverting 423
 - guidelines for scheduled tasks 74
- RFC 2317 618
- Riverbed Steelhead appliances
 - vNIOS appliances 223
- roaming hosts 863
- root name servers 587
- root zone 620
- RRset Orders 561
- RRSIG resource record 737
- Rulesets
 - blacklist 580
 - NXDOMAIN 575

S

- scavenging leases 794
- scheduling tasks
 - canceling scheduled restart 388
 - distributions 413
 - upgrades 416
- search
 - global 69
 - setting filters 67
- secondary servers
 - configuring 626
 - forwarding updates 707
 - name server groups 629
 - notifying external 561
- Security Banner Messages 268
- service status 1011
- Setting Priority for Managed Load Balancers 1188
- shared networks
 - defining admin permissions 207
- shared record groups 680
 - admin permissions 203
- shutting down 380
- smart folders 139–145
 - Finder panel 63
- SNMP 1037–??
 - enable threshold crossing event trap 809
 - Infoblox MIBs 1048
- SOA records
 - adding email address 560
 - stub zone 651
 - zone settings 624
- software and hardware requirements 46
- software distribution 412
 - monitoring 421
- Sophos NAC Advanced server 937
- sort lists
 - DNS 588
- sorting data 60
- SPF records 668
- splitting networks 472
- SRV records

- adding 666
- adding to shared record groups 683, 685
- SSH
 - remote console access 343
 - restrict access to MGMT port 364
- SSL
 - certificates 52–54
 - overview 52
- standalone
 - configuring 275
- Startup Wizard
 - independent appliance 279
 - independent HA pair 287
 - single grid master 242
- static routes
 - configuring 372
- stub zones 644
- subzones 620
- superusers
 - configuring admin groups 155
 - defining global smart folders 141
 - empty recycle bin 65
 - name server groups 629
 - scheduling grid restart 386
- support bundles
 - downloading 429
- syslog 1012–1017
 - exporting 91
- system status
 - Dashboard 123
- system temperature status 1008
- system time
 - monitoring 313
 - setting date and time 312
 - using NTP 313

T

- tables
 - customizing 60
 - printing data 91
 - selecting objects 60
 - setting size 50
- TACACS+
 - authenticating admins 179
- tasks
 - monitoring background 86
- TCP
 - network discovery 498, 523
- technical support
 - download support bundle 429
 - enabling and disabling access 343
 - restrict access to MGMT port 364
- testing upgrades 415
- TFTP
 - backup files 423
 - configuring 393
 - defining admin permissions 214

- uploading files 399
- UTF-8 encoding for file names 92
- viewing status 1011
- time zone
 - setting for grid 267
 - setting for grid and members 312
 - setting for independent appliances 289
 - setting in user profiles 51
 - upgrade schedule 416
- timeout, setting 344
- toolbar 59
- tooltips 60
- Top Devices Denied an IP Address 1151
- traffic capture tool 1021
- TTL
 - DNSSEC key rollovers 744
 - setting for resource records 557
- TXT records
 - adding 668
 - adding to shared record groups 683, 685
 - DDNS updates 703

U

- upgrade groups 408
- upgrade test 415
- upgrading 411
 - configuring upgrade groups 408
 - guidelines for scheduled tasks 74
- uploading NIOS software 412
- user class filters 904
- user names
 - setting in user profile 50
- user profiles 50
- UTF-8 encoding
 - multilingual support 92

V

- views
 - See* DNS views
 - See* network views
- virtual appliances 223
- virtual router ID
 - in VRRP advertisements 235
 - independent HA pair 288
- VMware
 - vNIOS appliances 223
- vNIOS appliances 223
- VRRP advertisements 235

W

- widgets
 - Dashboard 99

Z

zones

- applying name server groups 630
- configuring authoritative 616, 1233
- configuring properties 633
- copying records 608
- defining admin permissions 200, 201, 215
- delegation 638
- enabling and disabling 621
- enabling transfers 583
- forward 640
- importing data 284, 631
- locking and unlocking 621
- removing 634
- restoring data 636
- root 620
- stub 644
- subzones 620

