

HP Network Node Manager i Software

For the Windows[®], HP-UX, Linux, and Solaris operating systems

Software Version: NNMi 9.20

Deployment Reference

Document Release Date: May 2012
Software Release Date: May 2012



Legal Notices

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

© Copyright 2008–2012 Hewlett-Packard Development Company, L.P.

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

HP-UX Release 10.20 and later and HP-UX Release 11.00 and later (in both 32 and 64-bit configurations) on all HP 9000 computers are Open Group UNIX 95 branded products.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

Oracle and Java are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Oracle Technology — Notice of Restricted Rights

Programs delivered subject to the DOD FAR Supplement are 'commercial computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the licensing restrictions set forth in the applicable Oracle license agreement. Otherwise, programs delivered subject to the Federal Acquisition Regulations are 'restricted computer software' and use, duplication, and disclosure of the programs, including documentation, shall be subject to the restrictions in FAR 52.227-19, Commercial Computer Software-Restricted Rights (June 1987). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

For the full Oracle license text, see the license-agreements directory on the NNMi product DVD.

Acknowledgements

This product includes software developed by the Apache Software Foundation.
(<http://www.apache.org>)

This product includes software developed by the Indiana University Extreme! Lab.
(<http://www.extreme.indiana.edu>)

Available Product Documentation

In addition to this guide, the following documentation is available for NNMI:

- *HP Network Node Manager i Software Documentation List*—Available on the HP manuals web site. Use this file to track additions to and revisions within the NNMI documentation set for this version of NNMI. Click a link to access a document on the HP manuals web site.
- *HP Network Node Manager i Software Interactive Installation Guide*—This is an interactive document, and is available on the NNMI 9.20 product media. See the `nnmi_interactive_installation_en_README.txt` file, located on the product media, for more information.
- *HP Network Node Manager i Software Upgrade Reference*—Available on the HP manuals web site. This document provides information to help you upgrade from earlier versions of NNM and NNMI.
- *HP Network Node Manager i Software Release Notes*—Available on the product media and the NNMI management server.
- *HP Network Node Manager i Software System and Device Support Matrix*—Available on the product media and the NNMI management server.
- *HP Network Node Manager iSPI Network Engineering Toolset Planning and Installation Guide*—Available on the NNM iSPI NET diagnostics server product media.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

Or click the **New users - please register** link on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HP sales representative for details.

Support

Visit the HP Software Support Online web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport user ID, go to:

<http://h20229.www2.hp.com/passport-registration.html>

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

About This Guide	19
What Is in This Guide?	19
Path Conventions Used in This Document	20
Revision History	21
For More Information about NNMi	22
Preparation	25
Hardware and Software Requirements	27
Supported Hardware and Software	27
Checking for Required Patches	28
System Configuration (UNIX)	29
Installing NNMi and the NNM iSPIs	29
NNMi Coexistence with HP Performance Insight	29
NNMi Coexistence with HP Operations Agent	30
NNMi 9.1x and NNM iSPI Performance for Metrics Version Requirements	30
Configuration	31
General Concepts for Configuration	33
Task Flow Model	33
Best Practice: Save the Existing Configuration	34
Best Practice: Use the Author Attribute	34
User Interface Model	34
Ordering	35
Node Groups and Interface Groups	35
Group Overlap	36
Node Group Membership	37
Hierarchies/Containment	37
Device Filters	37
Additional Filters	38
Additional Nodes	38
Node Group Status	38
Interface Groups	39
Node/Interface/Address Hierarchy	39
Stop Everything and Start Over Again	40
NNMi Communications	43
Concepts for Communications	44
Levels of Communication Configuration	44

Network Latency and Timeouts	44
SNMP Access Control	45
SNMP Version Preferences	46
Management Address Preferences	47
Polling Protocols	47
Communication Configuration and the nnmsnmp*.ovpl Commands	48
Plan Communications	48
Default Communication Settings	48
Communication Configuration Regions	49
Specific Node Configurations	50
Retry and Timeout Values	50
Active Protocols	50
Multiple Community Strings or Authentication Profiles	51
SNMPv1 and SNMPv2 Community Strings	51
SNMPv3 Authentication Profiles	51
Configure Communications	52
Configuring SNMP Proxy Settings	52
Evaluate Communications	54
Are All Nodes Configured for SNMP?	54
Is SNMP Access Currently Available for a Device?	54
Is the Management IP Address Correct?	54
Is NNMi Using the Correct Communications Settings?	55
Do the State Poller Settings Agree with the Communication Settings?	55
Tune Communications	55
NNMi Discovery	57
Concepts for Discovery	57
NNMi Derives Attributes through Device Profiles	59
Plan Discovery	59
Select Your Primary Discovery Approach	59
List-Based Discovery	60
Rule-Based Discovery	60
Auto-Discovery Rules	61
Auto-Discovery Rule Ordering	61
Exclude Devices from Discovery	61
Ping Sweep	62
Discovery Hints from SNMP Traps	62
Discovery Seeds for Auto-Discovery Rules	62
Best Practices for Auto-Discovery Rules	63
Examples	63
Node Name Resolution	63
Subnet Connection Rules	64
Discovery Seeds	64
Rediscovery Interval	65
Do Not Discover Objects	65
Discover Interface Ranges	66
Monitor Virtual IP Addresses with NNMi	66
Configure Discovery	67

Tips for Configuring Auto-Discovery Rules	67
Tips for Configuring Seeds	67
Evaluate Discovery	68
Follow the Progress of Initial Discovery	68
Were All Seeds Discovered?	69
Do All Nodes Have a Valid Device Profile?	69
Were All Nodes Discovered Properly?	69
Auto-Discovery Rules	70
IP Address Ranges	71
System Object ID Ranges	71
Are All Connections and VLANs Correct?	71
Evaluate Layer 2 Connectivity	71
NNMi Discovery and Duplicate MAC Addresses	72
Rediscover a Device	72
Tune Discovery	72
Discovery Log File	73
Unnumbered Interfaces	73
Enabling the Unnumbered Interface Feature	73
Disabling the Unnumbered Interface Feature	75
Controlling Deletion of Unresponsive Objects	75
NNMi State Polling	77
Concepts for State Polling	77
Plan State Polling	78
Polling Checklist	78
What Can NNMi Monitor?	79
Interfaces to Unmonitored Nodes	80
Stop Monitoring	81
Planning Groups	81
Interface Groups	82
Node Groups	82
Planning Polling Intervals	83
Deciding What Data to Collect	84
Configure State Polling	84
Configure Interface Groups and Node Groups	84
Configure Interface Monitoring	85
Configure Node Monitoring	85
Verify Default Settings	86
Evaluate State Polling	86
Verify the Configuration for Network Monitoring	86
Is the interface or node a member of the right group?	86
Which settings are being applied?	87
Which data is being collected?	87
Evaluate the Performance of Status Polling	87
Is the State Poller keeping up?	87
Tune State Polling	89

NNMi Incidents	91
Concepts for Incidents	91
Incident Lifecycle	92
Trap and Incident Forwarding	93
Comparison: Forwarding Third-Party SNMP Traps to Another Application	94
MIBs	95
Custom Incident Attributes	96
CIAs Added to Closed Management Event Incidents	97
Incident Reduction	98
Incident Suppression, Enrichment, and Dampening	98
Lifecycle Transition Actions	99
Plan Incidents	100
Which Device Traps Should NNMi Process?	100
Which Incidents Should NNMi Display?	100
How Should NNMi Respond to Incidents?	100
Should NNMi Receive Traps from an NNM Management Station?	100
Should NNMi Forward Traps to Another Event Receiver?	100
Configure Incidents	101
Configuring Incident Suppression, Enrichment, and Dampening	101
Configuring Lifecycle Transition Actions	101
Configuring Trap Logs	102
Configuring Incident Logging	102
Configuring Trap Server Properties	102
Batch Load Incident Configurations	104
Generating an Incident Configuration File with <code>nnmincidentcfgdump.ovpl</code>	104
Loading Incident Configurations with <code>nnmincidentcfgload.ovpl</code>	104
Evaluate Incidents	105
Tune Incidents	105
Enabling and Configuring Incidents for Undefined Traps	106
NNMi Console	107
A Practical Example of Using Node Groups	107
Create Node Groups	108
Step 1: Create the My Network Node Group	108
Step 2: Create the USA Node Group	108
Step 3: Create the Colorado Node Group Using Filters	109
Step 4: View the Node Group Members to Check the Node Group Filter Results	109
Step 5: Set Up the Node Group Hierarchy for the My Network Node Group	110
Step 6: Establish the Node Group Hierarchy for the USA Node Group	110
Configure the Node Group Maps	110
Step 1: Create the Node Group Maps	110
Step 2: View the Node Group Maps	111
Step 3: Configure Node Group Status	111
Step 4: Configure Node Group Map Ordering	111
Step 5: Add a Background Image to a Node Group Map	112
Remove Node Groups	112
Step 1: Navigate to the Node Group	113
Step 2: Delete the Node Group	113

- Reducing the Maximum Number of Nodes Displayed in a Network Overview Map 113
- Reducing the Number of Displayed Nodes on a Node Group Map 114
- Configuring Gauges in the Analysis Pane 114
 - Limiting the Number of Gauges Displayed. 115
 - Setting the Refresh Rate for Gauges in the Analysis Pane 115
 - Eliminating Gauges from the Display 115
 - Controlling the Order of Displayed Node Gauges. 115
 - Controlling the Order of Displayed Interface Gauges 115
 - Controlling the Order of Displayed Custom Poller Gauges 116
 - Understanding how Gauge Properties are Applied 116
 - Determining the Names of Gauges 116
 - Troubleshooting Gauge Problems 117
 - Too Many Gauges. 117
 - Missing Gauges 117
- Customizing Device Profile Icons 117

Advanced Configuration 119

Licensing NNMi 121

- Preparing to Install a Permanent License Key 121
 - Checking the License Type and the Number of Managed Nodes 121
- Obtaining and Installing a Permanent License Key 122
 - Using Autopass and your HP Order Number (not possible behind a firewall) 122
 - From the Command Line 122
- Obtaining Additional License Keys 122

Working with Certificates for NNMi. 125

- Putting it All Together 126
- Generating a Certificate Authority Certificate 127
- Configuring Application Failover to use Self-Signed Certificates. 130
- Configuring Application Failover to use a Certificate Authority 132
- Configuring High Availability to use Self-Signed or Certificate Authority Certificates 134
 - Configuring High Availability to use Self-Signed Certificates. 134
 - Configuring High Availability for a New Certificate 134
- Configuring the Global Network Management Feature to use Self-Signed Certificates. 135
- Configuring the Global Network Management Feature to use a Certificate Authority 136
- Configuring Global Network Management with Application Failover to use Self-Signed Certificates . . 137
- Configuring an SSL Connection to the Directory Service 138

Using Single Sign-On with NNMi 141

- SSO Access for NNMi 142
- Enabling SSO for a Single Domain. 143
- Enabling SSO for NNMi Management Servers Located in Different Domains 143
- SSO Access for NNMi and the NNM iSPIs. 144
- Disabling SSO 146
- SSO Security Notes 146

Configuring the Telnet and SSH Protocols for Use by NNMi 149

- Disable the Telnet or SSH Menu Item 149

Configure a Telnet or SSH Client for the Browser on Windows	150
Windows Operating System-Provided Telnet Client	152
Third-Party Telnet Client (Standard Windows)	153
Third-Party Telnet Client (Windows on Windows)	154
Third-Party SSH Client (Standard Windows and Windows on Windows)	155
Configure Firefox to use Telnet or SSH on Linux	157
Telnet on Linux	157
Secure Shell on Linux	158
Example Files for Changing the Windows Registry	158
Example nnmtnet.reg	159
Example nnmputtytelnet.reg	159
Example nnmtnet32on64.reg	159
Example nnmssh.reg	159
Integrating NNMi with a Directory Service through LDAP	161
NNMi User Access Information and Configuration Options	161
Option 1: All NNMi User Information in the NNMi Database	162
Option 2: Some NNMi User Information in the NNMi Database and Some NNMi User Information in the Directory Service	163
Option 3: All NNMi User Information in the Directory Service	164
Configuring NNMi to Access a Directory Service	165
Changing the Directory Service Access Configuration to Support the NNMi Security Model	173
Directory Service Queries	175
Directory Service Access	175
Directory Service Content	175
Information Owned by the Directory Service Administrator	179
User Identification	180
Configuring NNMi User Access from the Directory Service (Detailed Approach)	181
User Group Identification	183
Configuring User Group Retrieval from the Directory Service (Detailed Approach)	183
Directory Service Configuration for Storing NNMi User Groups	185
Troubleshooting the Directory Service Integration	186
ldap.properties Configuration File Reference	187
Examples	192
Managing Overlapping IP Addresses in NAT Environments	193
What is NAT?	193
What are the Benefits of NAT?	193
What Types of NAT are Supported?	194
How is NAT Implemented in NNMi?	194
Static NAT Considerations	194
Hardware and Software Requirements and Static NAT	195
Communication and Static NAT	196
Administering ICMP Polling of the Management Address in a Static NAT Environment	196
Enabling ICMP Polling of the Management Address in a NAT Environment	196
How this Changes NNMi	196
Discovery and Static NAT	197

- Traps and Static NAT 197
 - SNMPv2c Traps 198
 - SNMPv1 Traps 199
- Subnets and Static NAT 201
- Global Network Management and Static NAT 201
- Dynamic NAT and PAT Considerations 201
 - Hardware and Software Requirements and Dynamic NAT and PAT 203
 - Discovery and Dynamic NAT and PAT 203
 - Subnets and Dynamic NAT and PAT 203
 - Global Network Management and Dynamic NAT and PAT 204
- Overlapping IP Address Mapping 204
 - Private IP Address Ranges 204
- NNMi Security and Multi-Tenancy 207**
 - Effects of Limiting Object Access 208
 - The NNMi Security Model 209
 - Security Groups 210
 - Example Security Group Structure 211
 - The NNMi Tenant Model 214
 - Tenants 214
 - Example Tenant Structure 215
 - NNMi Security and Multi-Tenancy Configuration 217
 - Configuration Tools 218
 - Configuring Tenants 219
 - Configuring Security Groups 221
 - Verifying the Configuration 222
 - Exporting the NNMi Security and Multi-Tenancy Configuration 224
 - NNMi Security, Multi-Tenancy, and Global Network Management 225
 - Initial GNM Configuration 226
 - GNM Maintenance 228
 - Including Select Interfaces in NPS Reports 228
- Global Network Management 229**
 - Global Network Management Benefits 229
 - Is Global Network Management a Good Tool for Managing my Network? 230
 - Do I Need Continuous Multi-Site Network Monitoring? 230
 - Can my Critical Devices be Visible? 230
 - Licensing Considerations 231
 - Practical Global Network Management Examples 231
 - Review the Requirements 232
 - Regional Manager and Global Manager Connections 233

Initial Preparation	234
Port Availability: Configuring the Firewall	234
Configuring Self-Signed Certificates	234
Configuring Global Network Management for Application Failover	234
NNMi Management Server Sizing Considerations	235
Synchronizing System Clocks	235
Using the Application Failover Feature with Self-Signed Certificates in Global Network Management	235
Using Self-Signed Certificates in Global Network Management	235
Using a Certificate Authority in Global Network Management	235
List the Critical Equipment you Want to Monitor	236
Review the Global and Regional Managers' Management Domains	236
Review NNMi Help Topics	237
SSO and the Actions Menu	237
Configuring Single Sign-On for Global Network Management	237
Configuring Forwarding Filters on the Regional Managers	239
Configuring a Forwarding Filter to Limit Forwarded Nodes	239
Connecting a Global Manager with a Regional Manager	246
Determining the Connection States from global1 to regional1 and regional2	250
Reviewing global1 Inventory	251
Disconnecting Communication between global1 and regional1	253
Additional Information	255
Discovery and Data Synchronization	255
Status Poll or Configuration Poll a Device	258
Determining Device Status and NNMi Incident Generation using a Global Manager	259
Configuring Application Failover for Global Network Management	260
Configuring Application Failover on a Global Manager	260
Troubleshooting Tips for Global Network Management	262
Troubleshooting Information in the NNMi Help	262
Clock Synchronization	262
Global Network Management System Information	262
Synchronize Regional Manager Discovery from a Global Manager	263
Remedying a Destroyed Database on global1	263
Upgrading Global and Regional Managers from NNMi 9.0x/9.1x to NNMi 9.20	264
NNMi Versions Supported by Global Network Management	264
Global Network Management Upgrade Steps	264
Global Network Management and NNM iSPIs or Third-Party Integrations	265
HP Network Node Manager iSPI Performance for Metrics Software	265
Global Network Management and Address Translation Protocols	265
Configuring NNMi Advanced for IPv6	267
Feature Description	267
Prerequisites	269
Licensing	269
Supported Configuration	270
Management Server	270
Supported SNMP MIBs for IPv6	271
Installing NNMi	271

Activating IPv6 Features	272
Deactivating IPv6 Features	274
IPv6 Monitoring Following Deactivation	274
IPv6 Inventory Following Deactivation	274
Known Issues When Cleaning Up IPv6 Inventory	276
Running NNMi in a Solaris Zones Environment	277
Installing NNMi in a Solaris Zone	277
Trap Forwarding in a Solaris Zone	277
Running NNMi Application Failover in a Solaris Zones Environment	278
Running NNMi under HA in a Solaris Zones Environment	278
Resilience	281
Configuring NNMi for Application Failover	283
Application Failover Overview	284
Application Failover Basic Setup	284
Configuring NNMi for Application Failover	286
Configuring your Cluster with the NNMi Cluster Setup Wizard (Embedded Database Users only)	287
Setting Cluster Communications (Optional)	288
Using the Application Failover Feature	288
Application Failover Behavior Using the Embedded Database	289
Application Failover Behavior Using an Oracle Database	291
Application Failover Scenarios	292
Additional ovstart and ovstop Options	292
Application Failover Incidents	293
Returning to the Original Configuration Following a Failover	293
NNM iSPIs and Application Failover	294
NNM iSPI Installation Information	294
Integrated Applications	295
Disabling Application Failover	296
Administrative Tasks and Application Failover	298
Application Failover and Upgrading to NNMi 9.20	298
Embedded Database	298
Oracle Database	301
Application Failover and NNMi Patches	303
Applying Patches for Application Failover (Shut Down Both Active and Standby)	303
Applying Patches for Application Failover (Keep One Active NNMi Management Server)	305
Application Failover and Restarting the NNMi Management Servers	307
Application Failover Control after a Communication Failure	307
Application Failover and Recovery from a Previous Database Backup (Embedded Database Only)	307
Network Latency/Bandwidth Considerations	309
Application Failover and the NNMi Embedded Database	309
Network Traffic in and Application Failover Environment	310
An Application Failover Traffic Test	311
Configuring NNMi in a High Availability Cluster	313
HA Concepts	314

HA Terms	315
NNMi HA Cluster Scenarios	316
Manpages	319
Verifying the Prerequisites to Configuring NNMi for HA	319
Configuring High Availability	321
Configure NNMi Certificates for HA	321
Configure NNMi for HA	321
NNMi HA Configuration Information	323
Configuring NNMi on the Primary Cluster Node	325
Configuring NNMi on the Secondary Cluster Nodes	327
Configure NNM iSPIs for HA	328
NNM iSPI Performance for Metrics, NNM iSPI Performance for QA, and NNM iSPI Performance for Traffic	328
NNM iSPI for MPLS, NNM iSPI for IP Multicast, and NNM iSPI for IP Telephony	329
NNM iSPI Network Engineering Toolset Software and NNMi Running under HA	329
Configure NNMi for HA in an Oracle Environment	329
NNMi Dependency on Oracle	330
Configuring NNMi for HA in an Oracle Environment	330
Shared NNMi Data	331
Data on the NNMi Shared Disk	331
Replication of Configuration Files	332
Disabling Data Replication	332
Prepare the Shared Disk Manually	332
Configuring a SAN or a Physically Connected Disk	333
Setting the HA Variables in the ov.conf File	333
Moving the Shared Disk into the NNMi HA Resource Group	334
A Note about Shared Disk Configuration on Windows Server	334
Licensing NNMi in an HA Cluster	334
Maintaining the HA Configuration	336
Maintenance Mode	336
Putting an HA Resource Group into Maintenance Mode	336
Removing an HA Resource Group from Maintenance Mode	336
Maintaining NNMi in an HA Cluster	337
Starting and Stopping NNMi	337
Changing NNMi Hostnames and IP Addresses in a Cluster Environment	337
Stopping NNMi Without Causing Failover	339
Restarting NNMi after Maintenance	340
Maintaining Add-on NNM iSPIs in an NNMi HA Cluster	340
Unconfiguring NNMi from an HA Cluster	340
Running NNMi Outside HA with the Existing Database	343
Patching NNMi under HA	344
Upgrading NNMi under HA from NNMi 9.0x/9.1x to NNMi 9.20	345
Upgrade NNMi with the Embedded Database on all Supported Operating Systems	345
Upgrade NNMi with Oracle on all Supported Operating Systems	349
Troubleshooting the HA Configuration	350
Common Configuration Mistakes	350
Configuration Issues with RHCS 6	351
HA Resource Testing	351

General HA Troubleshooting	352
Error: Wrong Number of Arguments	352
Resource Hosting Subsystem Process Stops Unexpectedly (Windows Server 2008 R2)	353
Product Startup Times Out (Solaris)	354
Product Startup Times Out (Windows MSCS 2008)	354
Log Files on the Active Cluster Node Are Not Updating	354
Cannot Start the NNMi HA Resource Group on a Particular Cluster Node	354
NNMi-Specific HA Troubleshooting	356
Re-Enable NNMi for HA after All Cluster Nodes are Unconfigured	356
NNMi Does Not Start Correctly Under HA	356
Changes to NNMi Data are Not Seen after Failover	357
nmsdbmgr Does Not Start after HA Configuration	357
pmd Does Not Start after HA Configuration	358
NNMi Runs Correctly on Only One HA Cluster Node (Windows)	358
Disk Failover Does Not Occur	359
Shared Disk is Not Accessible (Windows)	359
Shared Disk Does Not Contain Current Data	359
Shared Disk Files Are Not Found by the Secondary Node after Failover	359
NNM iSPI-Specific HA Troubleshooting	360
HA Configuration Reference	361
NNMi HA Configuration Files	361
NNMi-Provided HA Configuration Scripts	361
NNMi HA Configuration Log Files	363
NNMi Northbound Interface	365
NNMi Northbound Interface	366
Value	366
Supported Versions	366
Terminology	366
Documentation	367
Enabling the NNMi Northbound Interface	367
Using the NNMi Northbound Interface	368
Incident Forwarding	368
Incident Lifecycle State Change Notifications	369
Incident Correlation Notifications	369
Incident Deletion Notifications	370
Event Forwarding Filter	370
Changing the NNMi Northbound Interface	371
Disabling the NNMi Northbound Interface	371
Troubleshooting the NNMi Northbound Interface	372
Application Failover and the NNMi Northbound Interface	373
Local Northbound Application	373
Remote Northbound Application	374
NNMi Northbound Interface Destination Form Reference	374
Northbound Application Connection Parameters	375
NNMi Northbound Interface Integration Content	376
NNMi Northbound Interface Destination Status Information	378
MIB Information used by the NNMi Northbound Interface	379

Maintaining NNMi**381**

NNMi Backup and Restore Tools	383
Backup and Restore Commands	384
Backing up NNMi Data	384
Backup Type	384
Backup Scope	385
Restoring NNMi Data	388
Same System Restore	389
Different System Restore	389
Backup and Restore Strategies	389
Back up All Data Periodically	389
Back up Data Before Changing the Configuration	390
Back up Data Before Upgrading NNMi or the Operating System	390
Restore File System Files Only	391
Backing up and Restoring the Embedded Database Only	391
Using Backup and Restore Tools in an HA Environment	391
Backup	391
Restore	392
Maintaining NNMi	393
Administering Access Control Lists for NNMi Folders	394
Administering a Custom Poller Collection Export	395
Changing the Custom Poller Collections Export Directory	395
Changing the Maximum Amount of Disk Space for Custom Poller Collections Export	395
Changing the Custom Poller Metric Accumulation Interval	396
Administering Incident Actions	397
Setting the Number of Simultaneous Actions	397
Setting the Number of Threads for Jython Actions	397
Setting the Action Server Name Parameter	398
Changing the Action Server Queue Size	399
Incident Actions Log	399
Blocking Trap Storms using the hosted-on-trapstorm.conf File	400
Blocking Incidents using the trapFilter.conf File	400
Configuring Character Set Encoding Settings for NNMi	401
Configuring NNMi to Require Encryption for Remote Access	401
Configuring the Auto-Trim Oldest SNMP Trap Incidents Feature	402
Enabling the Auto-Trim Oldest SNMP Trap Incidents Feature (No Incident Archive)	402
Enabling the Auto-Trim Oldest SNMP Trap Incidents Feature (Incident Archive Enabled)	403
Reducing the Number of Stored SNMP Trap Incidents	404
Monitoring the Auto-Trim Oldest SNMP Trap Incidents Feature	404
Disabling the Auto-Trim Oldest SNMP Trap Incidents Feature	404
Modifying NNMi Gauge Titles to Show SNMP MIB Variable Names	405
Modifying NNMi Normalization Properties	405
Changing Normalization Properties Following an Initial Discovery	406
Modifying Simultaneous SNMP Requests	407
Modifying the Embedded Database Port	407
Modifying MIB Browser Parameters	408

NNMi Self Monitoring	409
Suppressing the Use of Discovery Protocols for Specific Nodes	409
Suppressing the Use of Discovery Protocol Collections	410
Suppressing the Use of VLAN-indexing for Large Switches	411
Suppressing the Use of VLAN-indexing	411
Configuring Node Component Status	412
Propagating Node Component Status to a Node.	412
Configuring a Node Component's Status so that it does not Propagate to the Node	413
Overriding Node Component Status Values	414
NNMi Logging	415
NNMi Log Files	415
Changing Logging File Properties	416
Sign-in and Sign-out Logging.	416
NNMi Security	417
Configuring SSL Communications for Web Access and RMI Communications	417
Allowing Non-Root UNIX Users to Start and Stop NNMi.	417
Providing a Password for Embedded Database Tools	417
Additional Information	419
Manually Configuring NNMi for Application Failover	421
NNMi Environment Variables	425
Environment Variables Used in This Document	425
Other Available Environment Variables.	425
NNMi 9.20 and Well-Known Ports	429
NNMi 9.20 iSPI Well-Known Ports	435
Suggested Configuration Changes	445
Problems and Solutions	445
Glossary	451
We appreciate your feedback!	459

About This Guide



(1) First installation or test bed

Follow steps in NNMi Installation Guide



(2) Production deployment and migration from previous versions

Read NNMi Deployment Reference (this book)



This chapter contains the following topics:

- [What Is in This Guide?](#)
- [Path Conventions Used in This Document](#)
- [Revision History](#)
- [For More Information about NNMi](#)

What Is in This Guide?

This guide contains a collection of information and best practices for deploying HP Network Node Manager i Software, including NNMi and NNMi Advanced. This guide is for an expert system administrator, network engineer, or HP support engineer with experience deploying and managing networks in large installations.

This guide assumes that you have already installed NNMi in a limited (test) environment, and that you are familiar with start-up configuration tasks, such as using the Quick Start Configuration wizard to configure community strings, set up discovery for a limited range of network nodes, and create an initial administrator account. To learn more about these tasks, see the *HP Network Node Manager i Software Interactive Installation Guide* (see [Available Product Documentation](#) on page 3).

HP updates this guide between product releases, as new information becomes available. For information about retrieving an updated version of this document, see [Available Product Documentation](#) on page 3.

Path Conventions Used in This Document

For commands located in the NNMi bin directory, this document does not include the command path. The NNMi bin directory is located as follows:

- *Windows Server 2008*: <drive>\Program Files (x86)\HP\HP BTO Software\bin
- *UNIX*[®]: /opt/OV/bin

This document primarily uses the following two NNMi environment variables to reference file and directory locations. This list shows the default values. Actual values depend on the selections that you made during NNMi installation.

- *Windows Server 2008*:
 - %NnmInstallDir%: <drive>\Program Files (x86)\HP\HP BTO Software
 - %NnmDataDir%: <drive>\ProgramData\HP\HP BTO Software



On Windows systems, the NNMi installation process creates these system environment variables, so they are always available to all users.

- *UNIX*:
 - \$NnmInstallDir: /opt/OV
 - \$NnmDataDir: /var/opt/OV



On UNIX systems, you must manually create these environment variables if you want to use them.

Additionally, this document references some of the NNMi environment variables that you can source as part of your user log-on configuration on the NNMi management server. These variables are of the form `NNM_*`. For information about this extended list of NNMi environment variables, see [Other Available Environment Variables](#) on page 425.

Revision History

The following table lists the major changes for each new release of this document.

Document Release Date	Description of Major Changes
March 2011 (9.10)	Entirely updated. <ul style="list-style-type: none">• Fourth English edition.• Third Japanese edition.
May 2012 (9.20)	<ul style="list-style-type: none">• Extracted integration information to form separate documents.• Added Managing Overlapping IP Addresses in NAT Environments.• Moved Upgrading from NNMi 9.0x or 9.1x information to the <i>NNMi Upgrade Reference</i>.• Added NNMi Security.• Added Configuring your Cluster with the NNMi Cluster Setup Wizard (Embedded Database Users only).• Updated NNMi 9.20 and Well-Known Ports.• Added NNMi 9.20 iSPI Well-Known Ports.

For More Information about NNMi

To obtain a complete set of information about the NNMi product, use this guide along with other NNMi documentation. The table below shows all NNMi documents to date, including both guides and white papers.



All information below can be downloaded from <http://h20230.www2.hp.com/selfsolve/manuals>. See [Available Product Documentation](#) on page 3 for more information.

What do you want to do?	Where to find more information
View a list of available documentation for this version of NNMi.	Download the <i>NNMi Documentation List</i> . Use this file to track additions to and revisions within the NNMi documentation set for this version of NNMi. Click a link to access a document on the HP manuals web site.
Install NNMi or NNMi Advanced (first time).	Download the <i>HP Network Node Manager i Software Interactive Installation Guide</i> . This guide contains basic steps to install and un-install the product, plus how to do an initial configuration using the NNMi Quick Start Configuration Wizard. <ul style="list-style-type: none"> • <i>HP Network Node Manager i Software Installation Guide for the Windows Operating System</i> • <i>HP Network Node Manager i Software Installation Guide for the HP-UX Operating System</i> • <i>HP Network Node Manager i Software Installation Guide for the Linux Operating System</i> • <i>HP Network Node Manager i Software Installation Guide for the Solaris Operating System</i>
Plan for network deployment, including links to system requirements.	See Preparation on page 25 of this guide.
Configure NNMi for a production environment.	See Configuration on page 31 of this guide.
Configure NNMi behind the scenes.	See Advanced Configuration on page 119 of this guide.
Maintain the NNMi configuration.	See Maintaining NNMi on page 381 of this guide.
Upgrade to NNMi from previous versions of Network Node Manager i Software.	See <i>HP Network Node Manager i Software Upgrade Reference</i> , available on the HP manuals web site.
Reference NNMi environment variables, ports, and messages.	See Additional Information on page 419 of this guide.
Obtain more information about a specific topic.	Download by example documents and white papers.

What do you want to do?	Where to find more information
Print the NNMi help.	Download PDFs of the help content.
Install the HP NNM iSPI NET (NNM iSPI NET) diagnostics server and learn about NNM iSPI NET functionality.	Download the <i>HP NNM iSPI Network Engineering Toolset Planning and Installation Guide</i> from the Network Node Manager SPI for NET product category for the Windows operating system.
Obtain documentation about the NNMi Developer Toolkit (SDK).	See Licensing NNMi to review information related to the SDK, obtaining and installing an SDK license, and viewing SDK documentation and samples.

Preparation

This section contains the following chapter:

- [Hardware and Software Requirements](#)

Hardware and Software Requirements

This chapter contains the following topics:

- Supported Hardware and Software
- Checking for Required Patches
- System Configuration (UNIX)
- Installing NNMi and the NNM iSPIs
- NNMi Coexistence with HP Performance Insight
- NNMi Coexistence with HP Operations Agent
- NNMi 9.1x and NNM iSPI Performance for Metrics Version Requirements

Supported Hardware and Software

Before installing NNMi, read the information about NNMi hardware and software requirements described in [Table 1](#).



For current versions of all documents listed here, go to:

<http://h20230.www2.hp.com/selfsolve/manuals>

Table 1 Software and Hardware Pre-Installation Checklist

Complete (y/n)	Document to Read
	<p><i>HP Network Node Manager i Software Interactive Installation Guide</i></p> <ul style="list-style-type: none"> • Filename = <i>nnmi_interactive_installation_en.zip</i> or <i>nnmi_interactive_installation_en.jar</i> • Instructions Filename: <i>nnmi_interactive_installation_en_README.txt</i> • Windows Media = DVD main drive (root) • UNIX Media = Root directory
	<p><i>NNMi Release Notes</i></p> <ul style="list-style-type: none"> • Filename = <i>releasenotes_en.html</i> • Windows Media = DVD main drive (root) • UNIX Media = Root directory • NNMi console = Help > NNMi Documentation Library > Release Notes
	<p><i>NNMi System and Device Support Matrix</i></p> <ul style="list-style-type: none"> • Filename = <i>supportmatrix_en.html</i> • Windows Media = DVD main drive (root) • UNIX Media = Root directory • NNMi console = Linked from the release notes

- HP updates the *NNMi System and Device Support Matrix* as new information becomes available. Before you deploy NNMi, check for the most recent NNMi support matrix for your version of the software at:

http://www.hp.com/go/hpsoftwaresupport/support_matrices

(You must have an HP Passport ID to access this web site.)

- If you plan to install NNM Smart Plug-ins (NNM iSPIs), include the system requirements for those products as you plan the NNMi deployment.

Checking for Required Patches

NNMi ships an embedded Java virtual machine and JDK version 1.6. Java requires specific operating system patches to function properly.

If you plan to install NNMi on servers running supported operating systems other than HP-UX, consult the release notes for those operating systems.

System Configuration (UNIX)

If you cannot display NNMi manpages on the NNMi management server, verify that the `MANPATH` variable contains the `/opt/OV/man` location. If it does not, add the `/opt/OV/man` location to the `MANPATH` variable.

- ▶ NNMi uses a configuration file located in the `/etc/opt/OV` directory. Do not delete this directory.

Installing NNMi and the NNM iSPIs

If you plan to use any of the HP NNM iSPIs along with NNMi, you must install NNMi before installing any of the HP NNM iSPIs.

NNMi Coexistence with HP Performance Insight

If you plan to install NNMi on the same server as HP Performance Insight, follow this procedure to avoid problems with the installation sequence and port conflicts:

- 1 Install HP Performance Insight first.

- ▶ Do not install NNMi until after you complete [step 1](#) and [step 2](#).

- 2 Stop all HP Performance Insight processes.
- 3 Install NNMi. See the *HP Network Node Manager i Software Interactive Installation Guide* for specific instructions.
- 4 Stop all NNMi processes:

```
ovstop -c
```
- 5 Modify the `nms-local.properties` file to resolve any port conflicts. You can find this file in the following directory:
 - *Windows*: `%NNM_CONF%\nnm\props`
 - *UNIX*: `$NNM_CONF/nnm/props`
- 6 Start HP Performance Insight processes.

- 7 Start all NNMi processes:

```
ovstart -c
```



When NNMi is installed on the same server as HP Performance Insight, uninstalling NNMi causes an exception when running the HP PI MIB Browser. To prevent this exception, complete the following steps:

- 1 Uninstall NNMi.

- 2 Recreate the snmpmib MIB database:

```
a mkdir -p /var/opt/OV/shared/nnm/conf/
```

```
b /opt/OV/sbin/nnmloadmib -load /usr/OVPI/mibs/GENMIB2IF.mib
```

- 3 Use the `nnmloadmib.ovpl` command to load additional MIBs.

NNMi Coexistence with HP Operations Agent

If you plan to install an HP Operations agent on the NNMi management server (for communicating with HP Operations Manager (HPOM)), install NNMi before installing the HP Operations agent.

NNMi 9.1x and NNM iSPI Performance for Metrics Version Requirements


NNMi 9.1x and NNM iSPI Performance for Metrics must have equivalent versions:

- NNM iSPI Performance for Metrics version 9.10 is only supported with NNMi 9.10.
- NNM iSPI Performance for Metrics version 9.11 is only supported with NNMi 9.1x patch 1 (9.11).


Configuration

This section contains the following chapters:

- [General Concepts for Configuration](#)
- [NNMi Communications](#)
- [NNMi Discovery](#)
- [NNMi State Polling](#)
- [NNMi Incidents](#)
- [NNMi Console](#)



General Concepts for Configuration



Read this chapter for an introduction to concepts that are explained in more detail later in this guide. This chapter also contains some best practices that apply to all HP Network Node Manager i Software (NNMi) configuration areas.

This chapter contains the following topics:

- [Task Flow Model](#)
- [Best Practice: Save the Existing Configuration](#)
- [Best Practice: Use the Author Attribute](#)
- [User Interface Model](#)
- [Ordering](#)
- [Node Groups and Interface Groups](#)
- [Node/Interface/Address Hierarchy](#)
- [Stop Everything and Start Over Again](#)

Task Flow Model

The chapters in the configuration section of this guide support the following task flow:

- 1 **Concepts**—Gain a general understanding of the configuration area. The information in this guide supplements the information in the NNMi help.
- 2 **Plan**—Decide how you want to approach the configuration. This is a good time to begin or update your company's network management documentation.
- 3 **Configure**—Use a combination of the NNMi console, configuration files, and command line interface to enter the configuration into NNMi. See the NNMi help for specific procedures.



Writing, amending, or changing configurations in the embedded database using command line interfaces (such as PSQL commands) or external utilities is not supported. Attempting to do so may cause irreparable damage to the database.

- 4 **Evaluate**—In the NNMi console, examine the results of your configuration. Adjust the configuration as necessary to achieve the desired results.
- 5 **Tune**—Optional. Adjust the configuration to improve NNMi performance.

Best Practice: Save the Existing Configuration

It is a good idea to save a copy of the existing configuration before you make any major configuration changes. If you do not like the results of your configuration changes, it is easy to revert to your saved configuration.

Use the `nmmconfigexport.ovpl` command to save the current configuration. To recover a saved configuration, use the `nmmconfigimport.ovpl` command.

For information about how to use these commands, see the appropriate reference pages, or the UNIX manpages.



The `nmmconfigexport.ovpl` command does not retain SNMPv3 credentials. For more information, see the `nmmconfigexport.ovpl` reference page, or the UNIX manpage.

See also the *HP Network Node Manager i Software Step-by-Step Guide to Using NNMi Import and Export Tools White Paper*.

Best Practice: Use the Author Attribute

Many NNMi configuration forms include the **Author** attribute.

As you create or modify the configurations on these forms, set the **Author** attribute to a value that identifies your organization. When you export the NNMi configuration, you can specify an author value to pull only those items that your organization has customized.

When you upgrade NNMi, the installer does not overwrite any configurations whose author value is not HP.

User Interface Model

Some NNMi console forms use a transactional approach to updating the database. The changes that you make in the NNMi console forms do not take effect until you save and close the forms all of the way back to the NNMi console. If you close a form that contains unsaved changes (on that form or on a contained form), NNMi warns you about the unsaved changes and gives you a chance to cancel the close.



The **Discovery Seed** form is one exception to the transactional approach. This form is provided on the **Discovery Configuration** form as a convenience, but it is disconnected from the rest of discovery configuration. For this reason, you must save and close the **Discovery Configuration** form to implement your auto-discovery *rules* before you configure any discovery seeds for those rules.

Ordering

Some NNMi console configuration forms include the **Ordering** attribute, which sets the priority for applying the configurations. For one configuration area, NNMi evaluates each item against the configurations from the smallest (lowest) ordering number to the next lowest ordering number, and so on, until NNMi finds a match. At that point, NNMi uses the information from the matching configuration and ceases to look for any more matches. (The communication configuration is an exception. NNMi continues to search for information at other levels to complete the communication settings.)

The **Ordering** attribute plays an important role in NNMi configuration. If you see unexpected discovery or status results, check the ordering of the configurations for that area.

Ordering applies within the local context. The Menus and Menu Items tables contain multiple objects with the same ordering number because of the local context idea.

Ordering numbers are also used in the following places, but with different meanings:

- Ordering on the **Menu** and **Menu Item** forms sets the order of items in the local context of the associated menu.
- Topology maps ordering on the **Node Group Map Settings** form sets the order of items in the **Topology Maps** workspace.

For specific information about how the **Ordering** attribute affects a given configuration area, see the NNMi help for that area.

Best practice For each configuration area, apply low ordering numbers to the most restrictive configurations, and apply high ordering numbers to the least restrictive configurations.

Best practice For each configuration area, all ordering numbers must be unique. During initial configuration use ordering numbers with a standard interval to provide flexibility for future modifications to the configuration. For example, give the first three configurations the ordering numbers 100, 200, and 300.

Node Groups and Interface Groups

In NNMi, the primary filtering technique is to group nodes or interfaces, and then applying settings to a group or filtering visualizations by group. Node groups can be used for any or all of the following purposes:

- Monitoring settings
- Incident payload filtering
- Table filtering
- Customizing map views
- Filtering the nodes passed from a regional manager to the global manager for the global network management feature

Interface groups can be used for either or both of the following purposes:

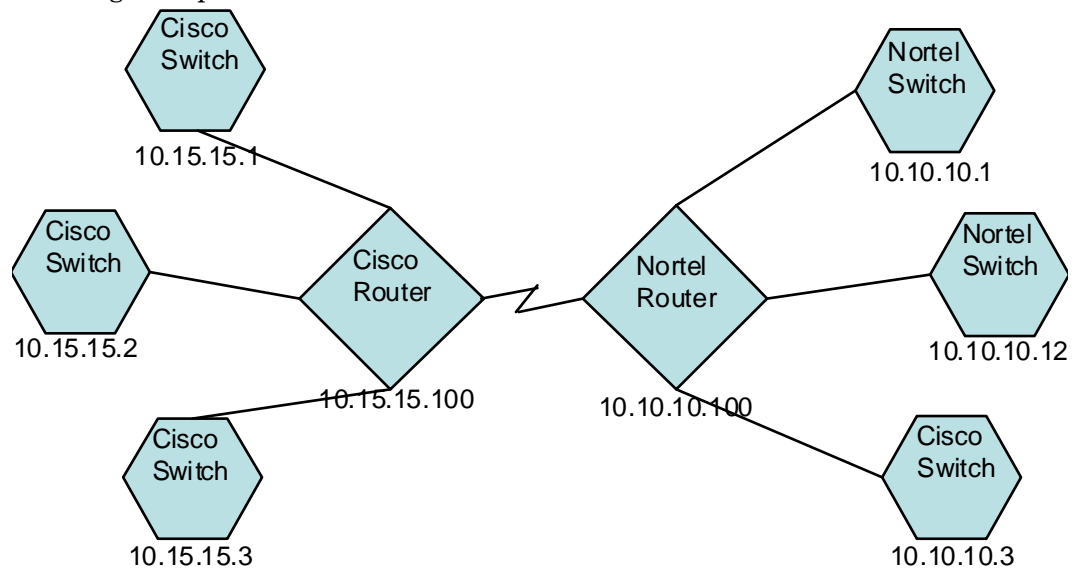
- Excluding interfaces from discovery

- Monitoring settings
- Incident payload filtering
- Table filtering

You can create a hierarchy of node groups based on any filterable attributes to control map view drill-down, monitoring or both settings inheritance.

Group Overlap

Regardless of the intended uses for group definitions, the first step is to define which nodes or interfaces are members of a group. Because you can create groups for different purposes, each object can be included in multiple groups. Consider the following example:



- For monitoring purposes, you might want to set a polling interval of 3 minutes for all switches, regardless of vendor or location. You can do this with a device category filter.
- For maintenance purposes, you might want to group all Cisco switches so that you can place them OUT OF SERVICE together for IOS upgrades. You can do this with a vendor filter.
- For visualization, you might want to group all devices on the 10.10.*.* site into a container with propagated status. You can do this with an IP address filter.

The Cisco switch with IP address 10.10.10.3 would qualify for all three groups.

You want to find the balance between having a usable rich set of groups available for configuration and viewing, and overloading the list with superfluous entries that will never be used.

Node Group Membership

NNMi determines node group membership by comparing each discovered node to each of the configured node groups.



- All nodes specified on the **Additional Nodes** tab are members of the node group.
- Rarely use the **Additional Nodes** tab to add nodes to a node group, as it consumes excessive resources on the NNMi management server.
- All nodes that are members of at least one node group specified on the **Child Node Groups** tab are members of the node group.
- Any node that matches one or more entries (if any exist) on the **Device Filters** tab *and* the filter specified on the **Additional Filters** tab is a member of the node group.

Hierarchies/Containment

You can create simple, reusable, atomic groups and combine them hierarchically for monitoring or visualization. Using hierarchical containers for nodes greatly enhances map views by providing cues about the location or type of object at fault. NNMi gives you complete control of the definition of the groups and their drill-down order.

You can create simple, reusable atomic groups first, and then specify them as child groups as you build up. Alternatively, you can specify your largest parent group first and create child groups as you go.

For example, a network might contain Cisco switches, Cisco routers, Nortel switches, and Nortel routers. You can create parent groups for Cisco devices and for all switches. Because the hierarchy is specified when you create the parent and designate its children, each child group, such as Cisco switches, can have multiple parents.

Hierarchies work well for the following situations:

- Types of nodes with similar monitoring needs
- Geographical locations of nodes
- Types of nodes to be taken OUT OF SERVICE together
- Groups of nodes by operator job responsibility

When you use groups in map views and table views, you see a (configurable) propagated status for the group.



Keep in mind that as you use group definitions to specify monitoring configuration, hierarchy does *not* imply ordering for settings. The settings with the lowest ordering number apply to a node. By carefully incrementing ordering numbers, you can emulate inheritance concepts for settings.

The configuration interface automatically prevents circular hierarchy definitions.

Device Filters

During discovery, NNMi collects direct information through SNMP queries and derives other information from that through device profiles. (For more information, see [NNMi Derives Attributes through Device Profiles](#) on page 59.) By gathering the system object ID, NNMi can index through the correct device profile to derive the following information:

- Vendor

- Device category
- Device family within the category

These derived values, in addition to the device profile itself, are available for use as filters.

For example, you can group all objects from a specific vendor, regardless of device type and family. Or you can group all devices of a type such as router, across vendors.

Additional Filters

With the additional filters editor, you can create custom logic to match fields including:

- hostname (Hostname)
- mgmtIPAddress (Management Address)
- hostedIPAddress (Address)
- sysName (System Name)
- sysLocation (System Location)
- sysContact (System Contact)
- capability (Unique Key of the Capability)
- customAttrName (Custom Attribute Name)
- customAttrValue (Custom Attribute Value)

Filters can include the AND, OR, NOT, EXISTS, NOT EXISTS, and grouping (parentheses) operations. For more information, see *Specify Node Group Additional Filters* in the NNMi help.

Capabilities are primarily intended for other programs that integrate with NNMi. For example, router redundancy and component health add capabilities (fields) to the NNMi database. You can view these capabilities by examining the node details from a device that has already been discovered.

Custom attributes can be added by iSPIs, or you can create your own custom attributes. If you have not purchased the Web Services SDK, you must place values in the field for each node manually. For example, an asset number or serial number might be an attribute that is not a capability.

Additional Nodes

It is better to use **Additional Filters** to qualify nodes for node groups. If the network contains critical devices that are too difficult to qualify using filters, add them to a group by individual hostname. Only add nodes to a node group by individual hostnames as a last resort.



Rarely use the **Additional Nodes** tab to add nodes to a node group, as it consumes excessive resources on the NNMi management server.

Node Group Status

When configured to do so, NNMi determines the status of a node group using one of the following algorithms:

- Set the node group status to match the most severe status of any node in the node group. To use this approach, select the **Propagate Most Severe Status** check box on the **Status Configuration** form.
- Set the node group status using the thresholds set for each target status. For example, the default threshold for the target status of Minor is 20%. NNMi sets the status of the node group to Minor when 20% (or more) of the nodes in the node group have Minor status. To use this approach, clear the **Propagate Most Severe Status** check box on the **Status Configuration** form. You can change the percentage thresholds for the target thresholds on the **Node Group Status Settings** tab of this form.

Because status calculations for large node groups can be resource-intensive, node group status calculation is off by default for new installations of NNMi. (Upgrades from NNMi 8.x retain the prior status calculation settings.) You can enable status calculation with the **Calculate Status** check box on the **Node Group** form for each node group.

Interface Groups

Interface groups filter interfaces within nodes by IFTType or by other attributes, such as ifAlias, ifDescr, ifName, ifIndex, IP address, and so forth. Interface groups carry no hierarchy or containment, although you can further qualify membership based on the node group for the node hosting the interface.

Interface groups can be filtered on custom capabilities and attributes similarly to node groups.

Qualifications for interface groups are AND'd together within and across tabs.

Node/Interface/Address Hierarchy

NNMi assigns monitoring settings in the following manner:

- 1 **Interface Settings**—NNMi monitors each of the node's interfaces and IP addresses based on the first matching **Interface Settings** definition. The first match is the **Interface Settings** definition with the lowest ordering number.
- 2 **Node Settings**—NNMi monitors each node and each previously unmatched interface or IP address based on the first matching **Node Settings** definition. The first match is the **Node Settings** definition with the lowest ordering number.
 - ▶ Child node groups are included in the ordering hierarchy. If the parent node group has a lower ordering number (for example, parent=10, child=20), then the monitoring configuration specified for the parent node group also applies to the nodes in the child node group. To override a parent node group monitoring configuration, set the ordering number for the child node group to a number that is lower than the parent (for example, parent=20, child=10).
- 3 **Default Settings**—If no match is found for a node, interface, or IP address in [step 1](#) or [step 2](#), NNMi applies the default monitoring configuration settings.

Stop Everything and Start Over Again

If you want to completely restart discovery and redo all of the NNMi configuration, or if the NNMi database has become corrupted, you can reset the NNMi configuration and database. This process deletes *all* of the NNMi configuration, topology, and incidents.

For information about the commands identified in this procedure, see the appropriate reference pages, or the UNIX manpages.

Follow these steps:

- 1 Stop the NNMi services:

```
ovstop -c
```

- 2 Optional. Because this procedure deletes the database, you might want to back up the existing database before proceeding:

```
nnmbackup.ovpl -type offline -target <backup_directory>
```

- 3 Optional. If you want to keep any of the current NNMi configuration, use the `nnmconfigexport.ovpl` command to output the NNMi configuration to an XML file.



The `nnmconfigexport.ovpl` command does not retain SNMPv3 credentials. See the *nnmconfigexport.ovpl* reference page, or the UNIX manpage, for more information.

- 4 Optional. Use the `nnmtrimincidents.ovpl` command to archive the NNMi incidents. Incidents are archived in the CSV format, as described in the `nnmtrimincidents.ovpl` reference page or UNIX manpage.
- 5 Drop and recreate the NNMi database.
 - For the embedded database, run the following command:

```
nnmresetembdb.ovpl -nostart
```

- For an Oracle database, ask the Oracle database administrator to drop and recreate the NNMi database. Maintain the database instance name.

- 6 If you have installed iSPIs or stand-alone products that integrate with NNMi, reset those products to remove the old topology identifiers. For specific procedures, see the product documentation.
- 7 Start the NNMi services:

```
ovstart -c
```

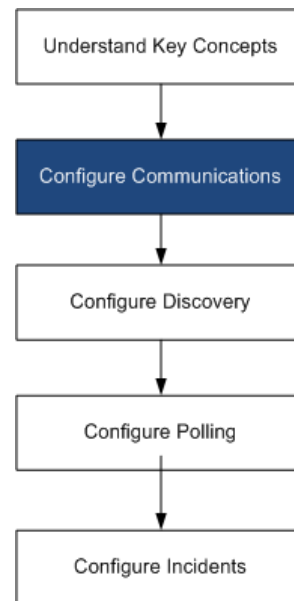
NNMi now has only the default configurations as if you had just installed the product on a new system.

- 8 Start configuring NNMi. Do one of the following:
 - Use the Quick Start Configuration Wizard.
 - Enter information into the **Configuration** workspace in the NNMi console.
 - Use the `nnmconfigimport.ovpl` command to import some or all of the NNMi configuration that you saved in [step 3](#).



If you are using the `nnmconfigimport.ovpl` command to import large amounts of configurations (such as 9,500 node groups or 10,000 incident configurations), consider using the `-timeout` option to adjust the import transaction timeout from its default value of 60 minutes (3600 seconds) to something longer. See the `nnmconfigimport.ovpl` reference page, or the UNIX manpage, for more information.

NNMi Communications



HP Network Node Manager i Software (NNMi) uses both Simple Network Management Protocol (SNMP) and Internet Control Message Protocol (ICMP ping) to discover devices and to monitor device status and health. To establish viable communication in your environment, you configure NNMi with the access credentials and appropriate timeout and retry values for different devices and areas of your network. You can disable a protocol in some areas of your network to reduce traffic or to respect firewalls.

The communication values that you configure form the foundation of NNMi discovery and state polling. NNMi applies the appropriate values for each device when making queries for discovery or polling. Thus, if you configure NNMi to disallow SNMP communication within some region of your network, neither NNMi discovery nor NNMi state polling can send SNMP requests to that region.

This chapter contains the following topics:

- [Concepts for Communications](#)
- [Plan Communications](#)
- [Configure Communications](#)
- [Evaluate Communications](#)
- [Tune Communications](#)

Concepts for Communications

NNMi uses SNMP and ICMP primarily in a request-response manner. Responses to ICMP ping requests verify address responsiveness. Responses to SNMP requests for specific MIB objects provide more comprehensive information about a node.

The following concepts apply to NNMi communications configuration:

- [Levels of Communication Configuration](#)
- [Network Latency and Timeouts](#)
- [SNMP Access Control](#)
- [SNMP Version Preferences](#)
- [Management Address Preferences](#)
- [Polling Protocols](#)
- [Communication Configuration and the `nnmsnmp*.ovpl` Commands](#)

Levels of Communication Configuration

NNMi communication configuration provides the following levels:

- Specific nodes
- Regions
- Global defaults

At each level you can configure access credentials, timeout and retry values, ICMP and SNMP protocol enablement, and SNMP access settings. If you leave settings blank at one level, NNMi applies the next level of defaults.

When communicating with a given node, NNMi applies the configuration settings as follows:

- 1 If the node matches a **specific node** configuration, NNMi uses any communication values in that configuration.
- 2 If any settings are not yet defined, NNMi determines whether the node belongs to any **regions**. Because regions might overlap, NNMi uses the matching region with the lowest ordering number. NNMi uses the values specified for that region to fill in the blanks left from the applicable specific node setting (if any). The settings for additional regions are not considered.
- 3 If any settings are still not yet defined, NNMi uses the **global default** settings to fill in the remaining blanks.

The values used for ICMP and SNMP communication with a particular device might be built up cumulatively until all required settings are determined.

Network Latency and Timeouts

Normal network latency influences the amount of time the NNMi management server must wait to get answers to ICMP and SNMP queries. Different areas of a network customarily have different turnaround times. For example, the local network where the NNMi management server resides could provide nearly instantaneous response, while responses from a device in a remote geographical region accessed through a

dial-up wide area link would typically take much longer. In addition, heavily-loaded devices might be too busy to respond to ICMP or SNMP queries immediately. When deciding which timeout and retry settings to configure, consider these latency concerns.

You can configure specific timeout and retry settings for both network regions and specific devices. The settings you choose determine how long NNMi waits for an answer and how many times NNMi requests data before abandoning the request when no answer is received.

For each request retry, NNMi adds the configured timeout value to the previous timeout value. Thus, the pause gets longer between each retry. For example, when NNMi is configured to use timeout of 5 seconds and three retries, NNMi waits 5 seconds for a response to the first request, 10 seconds for a response to the second request, and 15 seconds for a response to the third request before giving up until the next polling cycle.

SNMP Access Control

Communication with SNMP agents on managed devices requires access control credentials:

- SNMPv1 and SNMPv2c

A community string in each NNMi request must match a community string configured in the responding SNMP agent. All communication passes through the network in clear text (no encryption).

- SNMPv3

Communication with the SNMP agent complies with the user-based security model (USM). Each SNMP agent has a list of configured user names and their associated authentication requirements (the authentication profile). Formatting of all communication is controlled through configuration settings. NNMi SNMP requests must specify a valid user and follow the authentication and privacy controls configured for that user.

- Authentication protocol uses hash-based message authentication code (HMAC) using your choice of either the message-digest algorithm 5 (MD5) or the secure hash algorithm (SHA).
- Privacy protocol uses no encryption or the data encryption standard - cipher block chaining (DES-CBC) symmetric encryption protocol.



DES-CBC is considered a weak cipher. Therefore, if you are using DES-CBC, HP recommends that you choose a stronger cipher. To change your cipher selection:

- 1 In the NNMi console, click the **Configuration** workspace.
- 2 Expand the **Incidents** folder.
- 3 Expand the **Trap Server** folder.
- 4 Click **Trap Forwarding Configuration**.
- 5 In the **Privacy Protocol** list, select a strong cipher.



Avoid using DES-CBC when configuring SNMPv3 communication on the nodes that NNMi manages.

NNMi supports the specification of multiple SNMP access control credentials for a region of your network (defined through IP address filters or hostname filters). NNMi attempts communication with a device in that region by trying all configured values at a given SNMP security level in parallel. You can specify the minimum SNMP security level that NNMi uses in that region. NNMi uses the first value returned by each node (response from the device's SNMP agent) for discovery and monitoring purposes.

SNMP Version Preferences

The SNMP protocol itself has evolved over the years from version 1 to version 2(c) and now version 3, with increasing security capabilities (among others). NNMi can handle any or a mix of all versions in your network environment.

The first SNMP response NNMi receives for a particular node determines the communication credentials and SNMP version used by NNMi for communication with that node.



The SNMP version selection for a node plays a role in NNMi accepting traps from that node:

- If the source node or source object of the incoming trap has been discovered by NNMi using SNMPv3, NNMi accepts incoming SNMPv1, SNMPv2c, and SNMPv3 traps.
- If the source node or source object of the incoming trap has been discovered by NNMi using SNMPv1 or SNMPv2c, NNMi discards incoming SNMPv3 traps.

You specify the minimum level of SNMP version and security settings that are acceptable in each area of your network. The options for the SNMP Minimum Security Level field are as follows:

- **Community Only (SNMPv1 only)**—NNMi attempts to communicate using SNMPv1 with the configured values for community strings, timeouts, and retries. NNMi does not try any SNMPv2c or any SNMPv3 settings.
- **Community Only (SNMPv1 or v2c)**—NNMi attempts to communicate using SNMPv2c with the configured values for community strings, timeouts, and retries. If there is no response to any community string using SNMPv2c, NNMi attempts to communicate using SNMPv1 with the configured values for community strings, timeouts, and retries. NNMi does not try any SNMPv3 settings.
- **Community**—NNMi attempts to communicate using SNMPv2c with the configured values for community strings, timeouts, and retries. If there is no response to any community string using SNMPv2c, NNMi attempts to communicate using SNMPv1 with the configured values for community strings, timeouts, and retries. If none work, NNMi tries SNMPv3.
- **No Authentication, No Privacy**—For users with no authentication and no privacy, NNMi attempts to communicate using SNMPv3 with the configured values for timeouts and retries. If none work, NNMi tries users with authentication and no privacy followed by users with authentication and privacy, if necessary.
- **Authentication, No Privacy**—For users with authentication and no privacy, NNMi attempts to communicate using SNMPv3 with the configured values for timeouts and retries. If none work, NNMi tries users with authentication and privacy.

- **Authentication, Privacy**—For users with authentication and privacy, NNMi attempts to communicate using SNMPv3 with the configured values for timeouts and retries.

Management Address Preferences

A node's **management address** is the address NNMi uses to communicate with the node's SNMP agent. You can specify the management address for a node (in the specific node settings), or you can let NNMi choose an address from the IP addresses associated with the node. You can fine-tune this behavior in the discovery configuration settings by excluding certain addresses from discovery. For information about how NNMi determines the management address, see *Node Form* in the NNMi help.

NNMi discovers and monitors devices on an ongoing basis. *After the first NNMi discovery cycle*, the **Enable SNMP Address Rediscovery** field controls NNMi behavior when previously discovered SNMP agents quit responding (for example, when you reconfigure the device's SNMP agent).

- If the **Enable SNMP Address Rediscovery** check box is selected, NNMi retries any configured values in search of one that works.
- If the **Enable SNMP Address Rediscovery** check box is cleared, NNMi reports the device as "Down" and does not attempt to find another communication configuration setting for that device.



The **Enable SNMP Address Rediscovery** check box is available at all levels of communication configuration.



The **Discover Any SNMP Device** and **Non-SNMP Devices** auto-discovery rule configuration fields influence the way NNMi uses SNMP. For more information, see *Configure Basic Settings for the Auto-Discovery Rule* in the NNMi help.

Polling Protocols

You can prevent NNMi from using SNMP or ICMP in portions of your network (for example, when firewalls in your infrastructure prohibit ICMP or SNMP traffic).

Disabling ICMP traffic to the devices in an area of the network has the following results in NNMi:

- The optional auto-discovery rule ping sweep feature cannot locate additional nodes in that region of your network. All nodes must either be seeded or available through answers to MIB object requests, such as neighbor's ARP cache, Cisco Discovery Protocol (CDP), or Extreme Discovery Protocol (EDP). Wide area network devices might be missed unless you seed every one of them.
- The State Poller cannot monitor devices that are not configured to respond to SNMP requests. (However, if the device responds to SNMP, State Poller does not use ICMP.)
- Operators cannot use **Actions > Ping** to check device reachability during troubleshooting.

Disabling SNMP traffic to the devices in an area of the network has the following results in NNMi:

- Discovery cannot gather any information about the devices except that they exist. All devices receive the `No SNMP` device profile.
- Discovery cannot find additional neighboring devices through queries. All devices must be directly seeded.
- Discovery cannot gather connectivity information from the devices, so they appear unconnected on NNMi maps.
- For devices with the `No SNMP` device profile, the State Poller respects the defaults of monitoring that device using only ICMP (ping).
- The State Poller cannot gather component health or performance data from the devices.
- The Causal Engine cannot contact the devices to perform neighbor analysis and locate the root cause of incidents.

Communication Configuration and the `nnmsnmp*.ovpl` Commands

The `nnmsnmp*.ovpl` commands look up the values for unspecified device communication settings in the NNMi database. This approach requires that the `ovjboss` process be running. If `ovjboss` is not running, the `nnmsnmp*.ovpl` commands behave as follows:

- For SNMPv1 and SNMPv2c agents, the commands use default values for any unspecified communication settings.
- For SNMPv3 agents, if you specify a user and password the commands use default values for any unspecified communication settings. If you do not specify a user and password, the commands fail.

Plan Communications

Make decisions in the following areas:

- [Default Communication Settings](#)
- [Communication Configuration Regions](#)
- [Specific Node Configurations](#)
- [Retry and Timeout Values](#)
- [Active Protocols](#)
- [Multiple Community Strings or Authentication Profiles](#)

Default Communication Settings


Because NNMi uses default values to complete any configuration settings that were not specified for the applicable region or specific node, set defaults to be reasonable for the majority of your network.

- Are there commonly-used community strings that NNMi should try?
- What default timeout and retry values are reasonable in your network?

Communication Configuration Regions

Regions represent areas of the network where similar communication settings make sense. For example, the local network around the NNMi management server usually returns responses very quickly. Areas of your network that are multiple hops away typically take longer to respond.

You do not need to configure each subnet or area of your network. You can combine areas into one region based on similar lag times. Consider the following network map:



For timeout and retry purposes, you might want to configure the following regions:

- Region A for Net 1
- Region B to include Net 10, Net 20, and Net 30
- Region C for the more distant outlying networks

You would decide how best to group Net 170, depending on whether traffic management configuration is set to prefer the one-hop or two-hop path from the NNMi management server.

Regions are also used to group devices with similar access credentials. If all routers in your network use the same community string (or a small set of possible community strings) and you can identify the routers with a naming convention (for example, `rtrnnn.yourdomain.com`), you can configure a region containing all routers so that they are handled similarly. If you cannot use a wildcard to group the devices, you can configure each as a specific node.

Plan your region configurations so that you can apply the same timeout and retry value and access credential configurations to all nodes in a region.

Region definitions can overlap, and a device might qualify for multiple regions. NNMi applies the settings from the region with the lowest ordering number (and no other matching regions).

Specific Node Configurations

For any device with unique communication configuration requirements, use the specific node settings to specify the communication settings for that node. Example uses of specific node settings include the following:

- A node that might not respond well to SNMPv2c/SNMPv3 GetBulk requests
- A node whose name does not match the name pattern of other similar nodes



You can enable or disable SNMP communication for a specific device. See *Specific Node Settings Form* in the NNMi help.

Retry and Timeout Values

Configuring longer timeouts and more retries can result in more responses from devices that are busy or distant. This higher response rate eliminates false down messages. However, it also lengthens the time to determine that actual down devices require attention. Finding the balance for each area of your network is important and might require a period of testing and adjusting values in your environment.

To get an idea of current lag time for each hop, do the following:

- *Windows*: Run a `tracert` to a device in each network area.
- *UNIX*: Run a `traceroute` to a device in each network area.

Active Protocols

You have two opportunities to control the type of traffic NNMi generates when communicating with devices in your network: communication and monitoring configuration settings. Use the communication settings when firewalls in your infrastructure prohibit ICMP or SNMP traffic. Use monitoring settings to fine tune protocol usage when you do not need a particular subset of data about devices. If either communication or monitoring settings disable a protocol for a device, NNMi does not generate that type of traffic to the device.



Disabling SNMP communication significantly compromises the NNMi status and health monitoring of your network.

Note whether each region or specific device should receive ICMP traffic.

You do not need to explicitly disable SNMP communication with devices for which you do not supply access credentials. By default, NNMi assigns those devices to the `No SNMP` device profile and monitors them using ICMP only.

Multiple Community Strings or Authentication Profiles

Plan the community strings and authentication profiles to be tried for each area of your network. For the default and region settings, you can configure multiple community strings and authentication profiles to be tried in parallel.



While trying probable community strings, NNMi queries might cause devices to generate authentication failures. Inform your operations department that authentication failures might safely be ignored while NNMi completes its initial discovery. Alternatively, you can minimize the number of authentication failures by configuring your regions (and the associated community strings and authentication protocols to try) as tightly as possible.

If your environment uses SNMPv1 or v2c *and* SNMPv3, determine the minimum acceptable security level for each region.

SNMPv1 and SNMPv2 Community Strings

For regions where SNMPv1 or v2c access is acceptable, gather the community strings in use within the region and any unique community strings required by specific devices.

SNMPv3 Authentication Profiles

For regions containing SNMPv3-accessible devices, determine the minimum acceptable default authentication profiles, the authentication profiles appropriate for each region, and the unique authentication credentials in use on specific devices (if any). Also determine the authentication and privacy protocols in use within your network.

For SNMPv3 communication, NNMi supports the following authentication protocols:

- HMAC-MD5-96
- HMAC-SHA-1

For SNMPv3 communication, NNMi supports the following privacy protocols:

- DES-CBC
- TripleDES
- AES-128
- AES-192
- AES-256

You can specify one (or no) authentication protocol and one (or no) privacy protocol for each specific node or region setting.



Use of the TripleDES, AES-192, or AES-256 privacy protocols requires the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files library, which is installed automatically as part of the NNMi installation process. If you accidentally delete the library, you can restore it by following the procedure in [Suggested Configuration Changes](#) on page 445.

Configure Communications

After reading the information in this section, see *Configuring Communication Protocol* in the NNMi help for specific procedures.

- It is a good idea to save a copy of the existing configuration before you make any major configuration changes. For more information, see [Best Practice: Save the Existing Configuration](#) on page 34.

Configure the following areas of communication:

- Default settings
- Region definitions and their settings
- Specific node settings

For specific nodes, you can enter node settings through the NNMi console or through a configuration file.

- **Save and Close** all **Communication Configuration** forms all of the way back to the NNMi console to implement your changes.

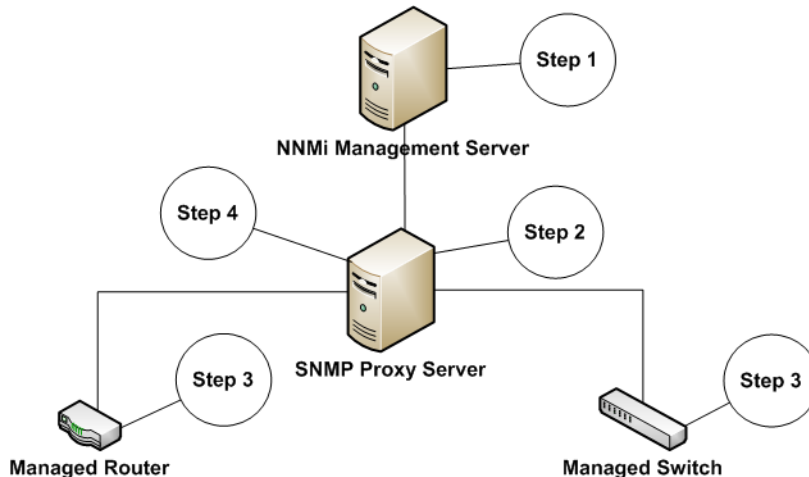
Best practice

Double-check the ordering numbers for the defined regions. If a node qualifies for membership in multiple regions, NNMi applies the settings from the region with the lowest ordering number to that node.

Configuring SNMP Proxy Settings

Some networks use an SNMP proxy agent to communicate with network devices. [Figure 1](#) shows the SNMP communication steps NNMi uses if you configure an **SNMP Proxy Address** and an **SNMP Proxy Port** using **Configuration > Communication Configuration** from the NNMi console. NNMi supports SNMP proxy servers that support using the `SecurityPackAgentAddressOid` OID (.1.3.6.1.4.1.99.12.45.1.1).

Figure 1 Using Proxy Servers



- 1 The NNMi management server sends an SNMP request to an SNMP proxy address and SNMP proxy port to obtain information from the managed router and the managed switch. The NNMi management server encodes the remote address

and port for the managed router and switch in a special proxy varbind, SecurityPackAgentAddressOid (.1.3.6.1.4.1.99.12.45.1.1), and adds this varbind to the SNMP request.

- 2 The SNMP proxy server reads the special proxy varbind, determines where to send the SNMP request, then sends SNMP requests to the managed router and switch to obtain the information requested by the NNMi management server.
- 3 The managed switch and router respond to the SNMP proxy server (using the SNMP Proxy Address and SNMP Proxy Port) with the requested information.
- 4 The SNMP proxy server responds to the NNMi management server (using the configured SNMP port).

When configured to use a proxy server, NNMi uses the following OIDs to handle SNMP responses:

- SecurityPackAgentAddressOid .1.3.6.1.4.1.99.12.45.1.1 (From SNMP Research NetDiscover SECURITY-PACK-MIB)
- SecurityPackNotificationAddressOid .1.3.6.1.4.1.99.12.45.2.1 (From SNMP Research NetDiscover SECURITY-PACK-MIB)
- ProxyOid .1.3.6.1.4.1.11.2.17.5.1.0 (HP)
- TrapForwardingAddressTypeOid .1.3.6.1.4.1.11.2.17.2.19.1.1.2.0 (HP)
- TrapForwardingAddressOid .1.3.6.1.4.1.11.2.17.2.19.1.1.3.0 (HP)
- Rfc3584TrapAddressOid .1.3.6.1.6.3.18.1.3.0 (RFC 3584)
- Rfc3584TrapCommunityOid .1.3.6.1.6.3.18.1.4.0 (RFC 3584)

When using NNMi with an SNMP proxy server, ask the proxy vendor if they support the OIDs in this list.

Evaluate Communications

This section lists ways to evaluate the progress and success of the communications settings. Most of these tasks can be completed only after discovery has completed.

Consider the following:

- [Are All Nodes Configured for SNMP?](#)
- [Is SNMP Access Currently Available for a Device?](#)
- [Is the Management IP Address Correct?](#)
- [Is NNMi Using the Correct Communications Settings?](#)
- [Do the State Poller Settings Agree with the Communication Settings?](#)

Are All Nodes Configured for SNMP?

- 1 Open the **Nodes** inventory view.
- 2 Filter the **Device Profile** column to contain the string `NO SNMP`.
 - For each of the devices that you want to manage, configure communication settings for the specific node. Alternatively, you can expand a region to include the node and update the access credentials.
 - If the communication settings are correct, verify that the SNMP agent on the device is running and properly configured (including ACLs).

Is SNMP Access Currently Available for a Device?

- 1 Select the node in an inventory view.
- 2 Select **Actions > Status Poll** or **Actions > Configuration Poll**.

If the results show any SNMP values, communication is operational.

You can also test communication from the command line with the `nnmsnmpwalk.ovpl` command. For more information, see the `nnmsnmpwalk.ovpl` reference page, or the UNIX manpage.

Is the Management IP Address Correct?

To determine which management address NNMi has selected for a device, follow these steps:

- 1 Select the node in an inventory view.
- 2 Select **Actions > Communication Settings**.
- 3 In the **Communication Configuration** window, verify that the management address of the SNMP agent listed in the Active SNMP Agent Settings list is correct.

Is NNMi Using the Correct Communications Settings?

Missing or incorrect SNMP community strings can result in incomplete discovery or can negatively affect the discovery performance.

To verify the communication settings configured for a device, use the `nmmcommconf.ovpl` command or follow these steps:

- 1 Select the node in an inventory view.
- 2 Select **Actions > Communication Settings**.
- 3 In the **Communication Configuration** window, verify that the values listed in the SNMP configuration settings table are the settings you want NNMi to use for this node.

If the communication settings are not correct, use the source information in the SNMP configuration settings table as a starting point for fixing the problem. You might need to change the configuration or the ordering number of a region or specific node.

Do the State Poller Settings Agree with the Communication Settings?

Even if the communication settings permit protocol traffic to an area of your network, that type of traffic might be disabled in the monitoring settings. To determine whether the settings are being overridden:

- 1 Select the node in an inventory view.
- 2 Select **Actions > Monitoring Settings**.

If either the Monitoring Settings or the Communication Settings disable a type of traffic to the device, that traffic will not be sent from NNMi.

Tune Communications

Reduce authentication failures

If NNMi is generating too many authentication traps during discovery, configure smaller regions or specific nodes with smaller groups of access credentials for NNMi to try.

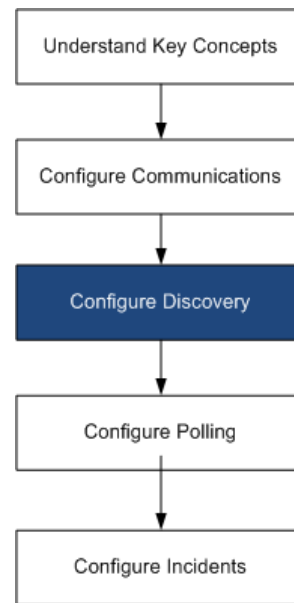
Tune timeouts and retries

When NNMi attempts to contact a device using SNMP during discovery, the communication configuration determines whether NNMi can gather the necessary device information. When the communication configuration does not include the correct SNMP community strings, or if NNMi is discovering non-SNMP devices, NNMi uses the configured settings for SNMP timeouts and retries. In this case, large timeout values or a high number of retries can negatively affect the overall performance of discovery. If your network contains devices that you know respond slowly to SNMP/ICMP requests, consider using the **Regions** or **Specific Node Settings** tabs on the **Communication Configuration** form to fine tune the timeout and retry values for just these devices.

Reduce default community strings

Having a large number of default community strings can negatively affect discovery performance. Instead of entering many default community strings, fine tune the community string configuration for particular areas of your network by using the **Regions** or **Specific Node Settings** tabs on the **Communication Configuration** form.

NNMi Discovery



One of the most important network management tasks is keeping your view of the network topology current. HP Network Node Manager i Software (NNMi) discovery populates the topology inventory with information about the nodes in your network. NNMi maintains this topology information through ongoing spiral discovery, which ensures that root cause analysis and the troubleshooting tools provide accurate information regarding incidents.

This chapter provides information to help you configure NNMi discovery. For an introduction to how discovery works and for detailed information about how to configure discovery, see *Discovering Your Network* in the NNMi help.

This chapter contains the following topics:

- [Concepts for Discovery](#)
- [Plan Discovery](#)
- [Configure Discovery](#)
- [Evaluate Discovery](#)
- [Tune Discovery](#)

Concepts for Discovery

The NNMi default behavior of discovering only routers and switches enables you to focus your network management on the critical or most important devices. In other words, target the backbone of the network first. Generally, you should avoid managing end nodes (for example, personal computers or printers) unless the end node is identified as a critical resource. For example, database and application servers might be considered critical resources.

NNMi provides several ways to control what devices to discover and include in the NNMi topology. Your discovery configuration can be very simple, quite complex, or anywhere in between, depending on how your network is organized and what you want to manage with NNMi.

▶ NNMi does not perform any default discovery. You must configure discovery before any devices appear in the NNMi topology.

Each discovered node (physical or virtually hosted) counts toward the license limit, regardless of whether NNMi is actively managing that node. The capacity of your NNMi license might influence your approach to discovery.

▶ For information about configuring Discovery to discover a large number of nodes, see the NNMi help.

Status monitoring considerations might also influence your choices. By default, the State Poller only monitors interfaces connected to devices NNMi has discovered. You can override this default for some areas of your network, and you can discover the devices beyond the edge of your responsibility. (For information about the State Poller, see [NNMi State Polling](#) on page 77.)

NNMi provides two primary discovery configuration models:

- **List-based discovery**—Explicitly tell NNMi exactly which devices should be added to the database and monitored through a list of seeds.
- **Rule-based discovery**—Tell NNMi which areas of your network and device types should be added to the database, give NNMi a starting address in each area, and then let NNMi discover the defined devices.

You can use any combination of list-based and rule-based discovery to configure what NNMi should discover. Initial discovery adds these devices to the NNMi topology, and then spiral discovery routinely rediscovers the network to ensure that the topology remains current.

▶ NNMi uses tenancy to support networks with overlapping address domains that may exist within static Network Address Translation (NAT), dynamic Network Address Translation (NAT), or Port Address Translation (PAT) areas of your network management domain. If you have such networks, put the overlapping address domains into different tenants (this is done using seeded discovery). See the NNMi help for more information.

▶ If you plan to configure multi-tenancy, configure tenants before initiating network discovery.

NNMi Derives Attributes through Device Profiles

As NNMi discovers devices, it uses SNMP to gather some attributes directly. One of the key attributes is the MIB II system object ID (`sysObjectID`). From the system object ID, NNMi derives additional attributes, such as vendor, device category, and device family.

During discovery, NNMi collects the MIB II system capabilities and stores them in the topology portion of the database. System capabilities are visible on the **Node** form. However, these capabilities are not used by any other portion of NNMi (specifically, monitoring configuration). NNMi uses the device category (from the device profile for the system object ID) to match devices into node groups. In node view tables, the **Device Category** column identifies the device category for each node.

NNMi ships with thousands of device profiles for system object IDs that were available at the time of release. You can configure custom device profiles for the unique devices in your environment to map these devices to category, vendor, and so forth.

Plan Discovery

Make decisions in the following areas:

- [Select Your Primary Discovery Approach](#)
- [Auto-Discovery Rules](#)
- [Node Name Resolution](#)
- [Subnet Connection Rules](#)
- [Discovery Seeds](#)
- [Rediscovery Interval](#)
- [Do Not Discover Objects](#)
- [Discover Interface Ranges](#)
- [Monitor Virtual IP Addresses with NNMi](#)

Select Your Primary Discovery Approach

Decide whether to do entirely list-based discovery, entirely rule-based discovery, or a combination of both approaches.

List-Based Discovery

With list-based discovery, you explicitly specify (as a discovery seed) each node that NNMi should discover.



NNMi uses tenancy to support networks with overlapping address domains that may exist within static Network Address Translation (NAT), dynamic Network Address Translation (NAT), or Port Address Translation (PAT) areas of your network management domain. If you have such networks, put the overlapping address domains into different tenants (this is done using seeded discovery). See the NNMi help for more information.



If you plan to configure multi-tenancy, list-based discovery is the recommended discovery approach.

Benefits of using only list-based discovery include:

- Provides very tight control over what NNMi manages.
- Supports the specification of a non-default tenant at discovery time.
- Simplest configuration.
- Good for fairly static networks.
- A good way to start using NNMi. You can add auto-discovery rules over time.

Disadvantages of using only list-based discovery include:

- NNMi does not discover new nodes as they are added to the network.
- You must provide the complete list of nodes to be discovered.

Rule-Based Discovery

With rule-based discovery, you create one or more auto-discovery rules to define the areas of the network that NNMi should discover and include in the NNMi topology. For each rule, you must provide one or more discovery seeds (by explicitly naming seeds or by enabling ping sweep), and then NNMi discovers the network automatically.

Benefits of using rule-based discovery include:

- Good for large networks. NNMi can discover a large number of devices based on minimal configuration input.
- Good for networks that change frequently. New devices that are added to the network are discovered without administrator intervention (assuming that each device is covered by an auto-discovery rule).
- Ensures that any new device added to your network is discovered to comply with service level agreements for managing new devices in a timely manner or security guidelines to flag unauthorized new devices.

Disadvantages of using rule-based discovery include:

- It is easier to run into license limitations.
- Depending on the structure of your network, tuning auto-discovery rules can be complex.
- If auto-discovery rules are very broad and NNMi discovers many more devices than you want to manage, you might want to delete the unneeded devices from NNMi topology. Node deletion can be time consuming.
- All non-seeded nodes receive the default tenant at discovery. If you want to use NNMi multi-tenancy, you must update the tenant assignment after discovery.

Rule-based discovery
only

Auto-Discovery Rules

Auto-Discovery Rule Ordering

The value of an auto-discovery rule's **Ordering** attribute affects discovery ranges in the following ways:

- IP address ranges

If a device falls within two auto-discovery rules, the settings in the auto-discovery rule with the lowest ordering number applies. For example, if an auto-discovery rule excludes a set of IP addresses, then no other auto-discovery rules with higher ordering numbers process those nodes and the nodes within that range of addresses are not discovered unless they are listed as discovery seeds.
- System object ID ranges
 - If no IP address range is included in an auto-discovery rule, then the system object ID settings apply to all auto-discovery rules with higher ordering numbers.
 - If an IP address range is included in an auto-discovery rule, the system object ID range applies only within the auto-discovery rule.

Exclude Devices from Discovery

- To prevent discovery of certain object types, create an auto-discovery rule with a low ordering number that ignores the system object IDs that you do not want discovered. Do not include an IP address range in this rule. By giving this auto-discovery rule a low ordering number, the discovery process quickly passes by the objects that match this rule.

- The **Ignored by Rule** setting for an IP address range or a system object ID range affects that auto-discovery rule only. The devices included in an ignored range are available to be included in another auto-discovery rule.



Some networks use routing protocols such as Hot Standby Router Protocol (HSRP) and Virtual Router Redundancy Protocol (VRRP) to provide router redundancy. When routers are configured in an router redundancy group (RRG), as they are when using HSRP, the routers configured in the RRG share a protected IP address (one active and one standby). NNMi does not support the discovery and management of multiple RRGs configured with the same protected IP address. Each RRG must have a unique protected IP address.

Ping Sweep

You can use ping sweep to locate devices within the IP address ranges of the configured auto-discovery rules. For initial discovery, you might want to enable ping sweep for all rules. Doing so provides enough information to NNMi discovery that you do not need to configure discovery seeds.



Ping sweep works for subnets of 16 bits or smaller, for example, 10.10.*.*.

Ping sweeps are especially useful for discovering devices across a WAN that you do not control, such as an ISP network.



Firewalls often view ping sweeps as attacks on the network, in which case, a firewall might block all traffic from a device that emits ping sweeps.

Best practice

Enable ping sweep for small discovery ranges only.

Discovery Hints from SNMP Traps

As of NNMi 9.01, NNMi processes the source IP address of received SNMP traps as hints to auto-discovery rules. This function is especially useful for discovering devices across a WAN.

Discovery Seeds for Auto-Discovery Rules

Provide at least one discovery seed per auto-discovery rule. The options for providing the seeds are as follows:

- Enter seeds on the **Discovery Seed** form by clicking **Seeds** under **Discovery** in the **Configuration** workspace.
- Use the `nnmloadseeds.ovpl` command to load information from a seed file.
- Enable ping sweep for the rule, at least for initial discovery.
- Configure a device to send SNMP traps to the NNMi management server.

Best Practices for Auto-Discovery Rules


- Because NNMi automatically manages all discovered devices, use IP address ranges that closely match the areas of the network that you want to manage.
 - You can use multiple IP address ranges within an auto-discovery rule to restrict discovery.
 - You can add a large IP address range to an auto-discovery rule and then exclude some IP addresses from discovery within that rule.
- The system object ID range specification is a prefix, not an absolute value. For example, the range 1.3.6.1.4.1.11 is the same as 1.3.6.1.4.1.11.*.

Examples

Discovery Rule Overlap

Figure 2 shows two discovery ranges that overlap. The circle on the left represents an IP address range or a system object ID range to be ignored by NNMi discovery. The circle on the right represents an IP address range or a system object ID range to be discovered and included in the NNMi topology. The overlapping region might be included or ignored by discovery, depending on the ordering of these auto-discovery rules.

Figure 2 Overlapping Discovery Ranges



Limit Device Type Discovery

To discover all HP devices in your network that are not printers, create one auto-discovery rule with a range to include the HP enterprise system object ID (1.3.6.1.4.1.11). In this auto-discovery rule, create a second range to ignore the system object IDs of HP printers (1.3.6.1.4.1.11.2.3 9). Leave the IP address range unset.

Node Name Resolution

By default, NNMi attempts to identify a node in the following order:

- 1 Short DNS name
- 2 Short sysName
- 3 IP Address



If you change a node's hostname, there is a delay before NNMi data reflects the name change, because NNMi caches DNS names to enhance performance.

The following scenarios describe situations in which you might want to change the default order for node name resolution:

- If your organization is dependent on others to update the DNS configuration, you might set a policy of defining the sysName for each new device as it is added to the network. In this case, set select sysName as the first choice for node name resolution so that NNMi can discover the new device as soon as it is deployed in the network. (Maintain the sysName over the life of the device.)
- If your organization does not set or maintain the sysName for managed devices, select sysName as the third option for node name resolution.

Best practice

If you use the full or short DNS name as the primary naming convention, confirm that you have forward and reverse DNS resolution from the NNMi management server to all managed devices.



When the full DNS name is the naming convention, labels on the topology maps can be long.

Best practice

NNMi selects the lowest loopback address as the management address for Cisco devices, so put DNS resolution on the lowest loopback address for each Cisco device. (NNMi 8.0x selects the highest loopback address as the management address.)

Subnet Connection Rules

List-based discovery
only

For list-based discovery, NNMi uses the subnet connection rules to detect connections that span a WAN. NNMi evaluates the subnet membership of the device it has discovered on each end of a probable connection (by examining their IP addresses and subnet prefixes) and looks at subnet connection rules for a match.

Rule-based discovery
only

When auto-discovery rules are enabled and NNMi finds a device configured with a subnet prefix between /28 and /31:

- 1 NNMi checks for an applicable subnet connection rule.
- 2 If a match is found, NNMi uses each valid address in the subnet as a hint and attempts a discovery on that address.

Best practice

Use the default connection rules. Only modify them if you have a problem.

Discovery Seeds

List the devices to use as discovery seeds.

Best practice

One of the NNMi rules for selecting the preferred management IP address specifies using the first discovered IP address as the management address. You can influence NNMi by configuring the preferred IP address as the seed address.

Best practice

For Cisco devices, use a loopback address as the discovery seed because loopback addresses are more reliably reachable than other addresses on a device. Ensure that DNS is correctly configured to resolve the device hostname to the loopback address.

List-based discovery
only

For list-based discovery, list all devices that you want NNMi to manage. You might be able to export this list from asset management software or from some other tool.

Because NNMi does not automatically add any devices to this list, ensure that the list includes every device for which you have responsibility or which influences your monitoring and status calculations.

Rule-based discovery
only

Discovery seeds are optional for rule-based discovery:

- If ping sweep is enabled for an auto-discovery rule, you do not need to specify a seed for that rule.
- For each auto-discovery rule with ping sweep disabled, identify at least one seed per rule. If a rule includes multiple IP address regions, you might need a seed in each routable region because routers do not keep ARP entries across WAN links.

Best practice

For the most complete rule-based discovery, use routers, not switches, as discovery seeds because routers generally have much larger ARP caches than do switches. A core router connected to a network that you want to discover is an excellent choice for a discovery seed.

Rediscovery Interval

NNMi rechecks the configuration information from each device in the database according to the configured rediscovery interval. In addition, NNMi collects the ARP cache from each router covered by an auto-discovery rule and looks for new nodes on the network.

Any change in the communication-related configuration of a device, such as interface renumbering, automatically triggers NNMi to update its data for that device and its neighbors.

The following changes do not trigger an automatic rediscovery; devices are updated only at the configured rediscovery interval:

- Changes within a node (for example, firmware upgrade or system contact).
- New nodes added to the network.

Select the rediscovery interval to match the level of change in the network. For a highly-dynamic network, you might want to use the minimum interval of 24 hours. For more stable networks, you can safely extend that period.

Do Not Discover Objects

In NNMi, there are three ways that you can configure NNMi to disregard certain objects:

- On the **Communication Configuration** form, you can turn off ICMP communication, SNMP communication, or both at different levels: globally, for communication regions, or for specific hostnames or IP addresses. For information about the impacts of disabling one or both of these protocols, see [Polling Protocols](#) on page 47.
- On the **Discovery Configuration** form, you can set up an auto-discovery rule that instructs NNMi to never gather hints from certain IP addresses or SNMP system object IDs. Nodes matching the criteria still appear on the map and in the database, but spiral discovery does not extend to the neighboring devices beyond those IP addresses or object types.
- On the **Discovery Configuration** form, you can set up an auto-discovery rule that instructs NNMi to exclude specific IP address ranges, IP addresses, or both from the database. Spiral discovery does not display those addresses on any node's list of addresses or use those addresses when establishing connections between devices, so NNMi never monitors the health of those addresses.

- On the **Excluded IP Addresses** tab of the **Discovery Configuration** form, you can exclude a range of IP addresses from being discovered by configuring an excluded IP addresses filter.



If you exclude an IP address range, any duplicates of addresses in static Network Address Translation (NAT), dynamic Network Address Translation (NAT), or Port Address Translation (PAT) areas of your network management domain are also excluded.

NNMi uses tenancy to support networks with overlapping address domains. If you have such networks, put the overlapping address domains into different tenants (this is done using seeded discovery). See the NNMi help for more information.

- On the **Excluded Interfaces** tab of the **Discovery Configuration** form, you can exclude a certain type of interface from the discovery process by selecting an Interface Group. See the NNMi help for more information.

Discover Interface Ranges

NNMi allows you to specify a range of interfaces to be discovered by defining a filter. This is particularly helpful when you have large nodes where you only want to discover a subset of the interfaces. When you specify a range of interfaces to be discovered, NNMi does not ask for information about interfaces outside that range; whereas, using the excluded interface option filters interfaces after retrieving the information from the device. Therefore, range-based discovery can improve discovery performance for large devices, especially when you do not want to manage all the interfaces on such devices.

The included interface ranges filter, defined on the **Included Interface Ranges** tab of the **Discovery Configuration** form, uses the System Object ID prefix and the ifIndex values to define the interface range. See the NNMi help for more information.

Monitor Virtual IP Addresses with NNMi

NNMi discovers and monitors devices such as clustered servers that share a virtual IP address. After a cluster fails over to a new active node, NNMi associates the virtual IP address with the new active node. This association is not immediate, as some time might pass between failover and NNMi discovering the change.

There are several actions you can take to configure NNMi for your specific situation:

If you want NNMi to monitor a virtual IP address, *use only one of the following options*:

- **Option 1:** For this option, NNMi manages N+1 non-SNMP devices, where N represents the number of members in the cluster discovered with a non-virtual IP address. NNMi discovers the additional (+1) non-SNMP node, and it is configured with the virtual IP address.

Do nothing to stop NNMi from discovering a virtual IP address. Using this approach, NNMi discovers the virtual IP address and the physical IP addresses associated with the Network Interface (NIC) cards on devices configured to use this virtual IP address. NNMi discovers and monitors each device as a separate non-SNMP node.

- Option 2: Configure NNMi to use a device's physical IP address as the Preferred Management Address of a clustered server. For instructions on how to do this, see the *Specific Node Settings Form (Communication Settings)* topic in the NNMi help.



NNMi might not immediately recognize the transfer of a virtual IP address from one active node to a new active node. NNMi might show the status of a virtual IP address using a node other than the current active node in the cluster.

If you do not want NNMi to monitor a virtual IP address, do the following using the NNMi console:

- 1 Click **Discovery Configuration** in the **Configuration** workspace.
- 2 Click the **Excluded IP Addresses** tab.
- 3 Add the virtual IP address or range of addresses to the list of addresses to be excluded from discovery.
- 4 Save your work.

Configure Discovery

This section lists configuration tips and provides some configuration examples. After reading the information in this section, see *Configure Discovery* in the NNMi help for specific procedures.



Because NNMi launches discovery from seeds as soon as you **Save and Close** the **Discovery Seed** form, ensure that you do the following before you configure seeds:

- Complete all communication configuration.
- Complete all auto-discovery rules (if any).
- Configure subnet connection rules.
- Configure name resolution preferences.
- **Save and Close** all of the way back to the console.



It is a good idea to save a copy of the existing configuration before you make any major configuration changes. For more information, see [Best Practice: Save the Existing Configuration](#) on page 34.

Tips for Configuring Auto-Discovery Rules

- As you define a new auto-discovery rule, check each setting carefully. For a new rule, auto-discovery is enabled by default, IP address ranges are included by default, and system object ID ranges are *ignored* by default.

Tips for Configuring Seeds

- If you already have a file that lists the nodes to be discovered, format this information as a seed file and use the `nnmloadseeds.ovpl` command to import the node list into NNMi.

- In the seed file, specify IP addresses as a way of influencing the IP address that NNMi chooses as the management address. (If you use hostnames, DNS provides the IP address for each node.)
- Good formats for the entries in the seed file are shown here:

```
IP_address1 # node name
```

```
IP_address2, <tenant_UUID_or_tenant_name> # node name
```

These formats are easy for both NNMi and human readers.

- For maintenance purposes, it is better to use only one seed file. Add nodes as needed and then rerun the `nmlloadseeds.ovpl` command. NNMi discovers the new nodes but does not re-evaluate the existing nodes.



If the seed file cannot be loaded, try making the file readable by `nmsproc` (644 permissions).

- Removing a node from the seed file does not remove it from the NNMi topology. Delete the node directly in the NNMi console.
- Deleting a node from a map or inventory view does not delete the seed.
- If you want NNMi to rediscover a node, delete that node from a map or inventory view *and* from the **Seeds** form in the **Discovery** area of the **Configuration** workspace in the NNMi console, and the re-enter the node in the NNMi console, or run the `nmlloadseeds.ovpl` command.
- Completely configure a discovery rule *before* you specify a seed for that rule. That is, click **Save and Close** on the **Discovery Configuration** form. (The **Discovery Seed** form is a separate form that is not part of the **Discovery Configuration** form in the database model. As a result, when you save the information on the **Discovery Seed** form, NNMi updates the seed configuration immediately.)

Rule-based discovery
only

Evaluate Discovery

This section lists ways to evaluate the progress and success of discovery.

Follow the Progress of Initial Discovery

NNMi discovery is dynamic and ongoing; it is never complete, so you will never see a “discovery completed” message. The process of initial discovery and connection takes some time. The following items suggest ways to gauge the progress of initial discovery:

- On the **Database** tab of the **System Information** window, watch for the node count to reach the expected level and stabilize. This window does not refresh automatically. During initial discovery, open the **System Information** window several times.
- Under **Discovery** in the **Configuration** workspace, look at the **Seeds** page. Refresh this page until all seeds show the `Node created` results, which indicates that the device has been added to the topology database. This result does *not* indicate that NNMi has gathered all information from the device and processed its connectivity.

- Open the **Node** form for representative nodes. When the **Discovery State** field (located on the **General** tab) transitions to `Discovery Completed`, NNMi has gathered the node's basic characteristics as well as the node's ARP cache and discovery protocol neighbors, if applicable. This state does *not* indicate that NNMi has completed connectivity analysis for the device.
- In the **Nodes** inventory view, scan to see that key devices are present from different areas of your network.
- Open the **Layer 2 Neighbor View** for representative nodes to determine whether connectivity analysis has completed for that area.
- Review the **Layer 2 Connections** and **VLANs** inventory views to gauge the progress of layer 2 processing.

Were All Seeds Discovered?

- 1 From the **Configuration** workspace, under **Discovery**, click **Seeds**.
- 2 On the **Seeds** page, sort the list of nodes by the **Discovery Seed Results** column. For any node in an error state, consider the following:
 - Failed discovery due to an unreachable node or unresolved DNS name or IP address—For these types of failures, verify network connectivity to the node and check for accurate DNS name resolution. To work around DNS issues, use the IP address to seed the node or include the hostname in a `hostnlookup.conf` file. For problems due to IP addresses that should not be resolved to hostnames, include the IP addresses in a `ipnlookup.conf` file. See the `hostnlookup.conf` and `ipnlookup.conf` reference pages, or the UNIX manpages, for more information.
 - License node count exceeded—This scenario occurs when the number of devices already discovered reached your license limit. You can either delete some discovered nodes or purchase additional node pack licenses.
 - Node discovered but no SNMP response—SNMP communication problems can occur for seeded devices as well as devices that are discovered through auto-discovery. For more information, see [Evaluate Communications](#) on page 54.

Do All Nodes Have a Valid Device Profile?

- 1 Open the **Nodes** inventory view.
- 2 Filter the **Device Profile** column to contain the string `No Device Profile`.
- 3 If a node is discovered but has no device profile, add a new device profile (from **Configuration > Device Profiles**), and then perform a configuration poll on the node to update its data.

Were All Nodes Discovered Properly?

To avoid discovery problems, NNMi should only manage nodes using a unique IP address that does not appear on any other node in the management domain. For example, if a node suddenly disappears or gets merged with another node in the database, and it is part of a Router Redundancy Group (RRG), there are special requirements. To manage a router that participates in an RRG, you must use a unique

IP address (which is not a protected address) as the management address of the router, and SNMP must be enabled on that address. NNMi will not properly manage a router if it tries to use a protected IP address as the management address.

Examine the data in the **Nodes** inventory view. If any nodes do not have a management address, check the communication settings for those nodes as described in [Are All Nodes Configured for SNMP?](#) on page 54.

If any expected nodes are missing from the **Nodes** inventory view, check the following:

- On each missing node, verify that the discovery protocol (for example, CDP) is correctly configured.
- If a missing node is on a WAN, enable ping sweep for the auto-discovery rule that includes that node.

List-based discovery
only

Auto-Discovery Rules

If you see unexpected discovery results, re-evaluate the auto-discovery rules.

When NNMi discovery finds an address hint, it uses the first matching rule to determine if a node should be created. If no rules are matched, NNMi discovery discards the hint. The ordering number for auto-discovery rules determines the order in which the auto-discovery rule configuration settings are applied.

For each auto-discovery rule, check the following settings:

- **Discover Included Nodes** must be enabled for auto-discovery to occur for the rule.
- Verify that the following settings are correct for the type of nodes you want discovered for the rule:
 - **Discover Any SNMP Device**
 - **Discover Non-SNMP Devices**

Remember that only routers and switches are discovered by default and non-SNMP nodes are *not* discovered. Enabling these settings without considering your environment can result in NNMi discovering more nodes than intended.

IP Address Ranges

The IP address of a discovery hint must match an **Include in Rule** entry in the IP address range list. If there are no included IP address ranges in an auto-discovery rule, then all address hints are considered a match. (For this case, see [Tips for Configuring Auto-Discovery Rules](#) on page 67.) Additionally, the hint must *not* match any entry marked **Ignored by Rule**. If all checks successfully match, this rule's configuration is used for handling the hint.

- If you are not discovering some expected devices, check your configured IP ranges to ensure that the IP addresses for those devices are included in a range and not ignored by a rule with a lower ordering number.
- If you are discovering more devices that you want, modify the include ranges or add ignored ranges for the IP addresses of the devices that you do not want discovered. Also, determine if **Discover Any SNMP Device** is enabled.

System Object ID Ranges

The system object ID (OID) from a discovery hint must match an **Include in Rule** entry in the system object ID ranges list. If there are no included system object ID ranges in an auto-discovery rule, then all object IDs are considered a match. Additionally, the OID must not match any entry marked **Ignored by Rule**. If all checks successfully match, this rule's configuration is used for handling the hint.

- Use the system object ID ranges to either expand auto-discovery to include more than the default routers and switches, or to exclude specific routers and switches.
- Each node must match both the IP address range and the system object ID range specified before it is discovered and added to the topology database.

Are All Connections and VLANs Correct?

NNMi creates Layer 2 connections and VLANs as a separate step after devices are added to the topology. Give NNMi plenty of time for initial discovery before evaluating connections and VLANs.

Evaluate Layer 2 Connectivity

To evaluate Layer 2 connectivity, create a node group for each network area of interest, and then display a topology map for that node group. (In the **Node Groups** inventory, select a node group, and then click **Actions > Node Group Map**.) Look for any nodes that are not connected to the other nodes in this map.

To evaluate VLANs, from the **VLANs** inventory view, open each **VLAN** form, and then examine the list of ports for that VLAN.

NNMi Discovery and Duplicate MAC Addresses

During discovery, NNMi reads the Forwarding Database (FDB) tables from Ethernet switches within a network to help NNMi determine communication paths between network devices. NNMi searches these FDB tables for information about discovered nodes. When an NNMi management server finds FDB references to duplicate Media Access Control (MAC) addresses, it does the following:

- If two or more discovered nodes contain an interface associated with the same Media Access Control (MAC) address within the same Tenant or with one of those nodes in Default Tenant and one in any other Tenant, NNMi disregards the communication paths reported for those duplicate MAC addresses in the FDB. This might result in missing connections on NNMi maps in network areas that include those duplicate MAC addresses.

NNMi Advanced - Global Network Management feature: If two NNMi management servers discover nodes that contain an interface associated with the same Media Access Control (MAC) address, the Global NNMi management server's maps could be missing connections that are visible on the Regional NNMi management server's maps.

- If a single node contains multiple interfaces that have the same MAC address, NNMi gathers all communication path information for those interfaces and displays that information on NNMi maps.

Forwarding Database (FDB) information can cause NNMi to establish wrong L2 Connections in the following cases:

- When the FDB is configured as cache and contains obsolete data.
- In network environments with hardware from a variety of vendors, each generating different and sometimes conflicting FDB data.

Optional: NNMi administrators can configure Discovery to ignore this FDB data for one Node Group.

Rediscover a Device

- 1 Perform a configuration poll of the device.
- 2 Delete the device.

If the device is a seed, delete the seed, and then re-add the seed.

Tune Discovery

For general discovery performance, fine tune the discovery configuration to discover only critical and important devices.

- Filter by IP address range, system object ID, or both.
- Limit discovery of non-SNMP devices and any SNMP devices (devices that are not switches or routers).

To delete one or more nodes from the NNMi database on the command line, use the `nmnodedelete.ovpl` command. This command deletes nodes, but not seed definitions, from the NNMi database.

To delete one or more seed definitions from the NNMi database on the command line, use the `nmseeddelete.ovpl` command.

There are special discovery circumstances that might be remedied by suppressing discovery protocol collections or VLAN-indexing. See [Suppressing the Use of Discovery Protocols for Specific Nodes](#) on page 409 or [Suppressing the Use of VLAN-indexing for Large Switches](#) on page 411 for more information.

Discovery Log File

Look in the `nm.?.?.log` file for messages containing the keyword `Exception` for the classes beginning with the string `com.hp.ov.nms.disco`. For information about log files, see [NNMi Logging](#) on page 415.

Unnumbered Interfaces

Prior to NNMi 9.10 patch 2, NNMi did not discover layer 2 connections for unnumbered interfaces unless you enabled xDP. NNMi 9.20 patch 2, provides an unnumbered interface discovery and monitoring solution that supports devices that use the default MIB-II `ipRoutingTable` and `ipCidrRoutingTable`.

The solution described in this section provides a way for NNMi 9.20 to discover and monitor IPv4 unnumbered interfaces and the associated layer 2 connections.

The solution explained in this section functions as follows in a global network management configuration:

- It works normally on a remote NNMi management server.
- It works on a global NNMi management server only for nodes managed on that server.
- It does not work on a global NNMi management server for nodes managed by a remote NNMi management server.

Enabling the Unnumbered Interface Feature

- 1 Create a node group that includes the devices containing the unnumbered interfaces. Either create a single node group that contains the device identifiers, or create a parent node group that represents multiple child node groups containing the device identifiers.
- 2 Create the following file:

Windows: `%NNM_DATA%\shared\nm\conf\disco\UnnumberedNodeGroup.conf`

UNIX: `$NNM_DATA/shared/nm/conf/disco/UnnumberedNodeGroup.conf`

- 3 Add a single node group name to this file. Again, this file must contain the name of a single node group containing the device identifiers, or it can be the name of a parent node group that represents multiple child node groups containing the device identifiers.

```
# This is the name of an node group containing devices with
unnumbered interfaces.
Unnumbered Node Group
```

In the example shown above, a node group named Unnumbered Node Group exists in NNMi. Add comment information as a separate line prefaced by a # character.

- 4 *Optional Step:* Create the following file:

Windows: %NNM_DATA%\shared\nnm\conf\disco\UnnumberedSubnets.conf

UNIX: \$NNM_DATA/shared/nnm/conf/disco/UnnumberedSubnets.conf

- 5 *Optional Step:* Add information to this file to show the specific routing address range you need NNMi to discover. You can add multiple lines of IPv4 CIDR subnet entries in random order to this file.

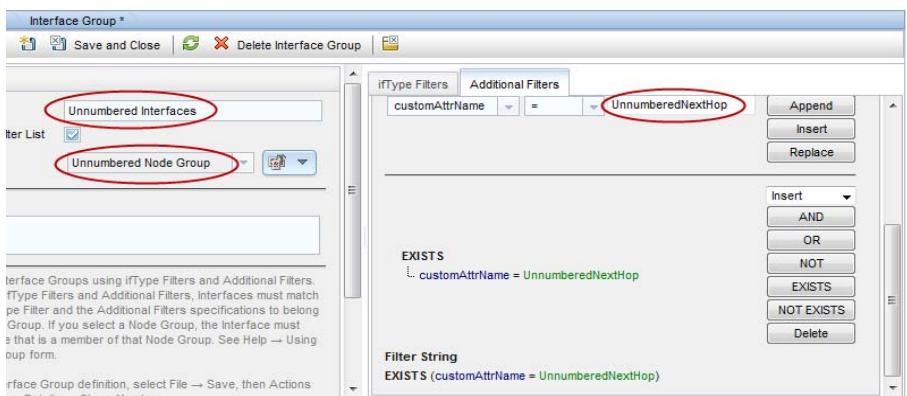


If you do not create and configure this file, NNMi will do a full MIB-II routing table walk against those nodes in the configured node group; by using the UnnumberedSubnets.conf file, NNMi requests MIB data from only those routes falling in the specified subnet destinations. It is a good practice to use this file and reduce the amount of discovery traffic and performance effect on the devices.

Below are some example entries to the UnnumberedSubnets.conf file.

```
10.1.5.0/18 #This entry filters the following routes: 10.1.0-63.
15.2.126.0/16 #This entry filters the following routes: 15.2.*.*
192.168.1.0/24 #This entry filters the following routes:
192.168.1.0-255
```

- 6 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.
- 7 Wait for NNMi to complete the next discovery cycle.
- 8 To find all of the unnumbered interfaces, configure a new interface group to include those interfaces having custom attribute named UnnumberedNextHop.



- 9 To view the layer 2 connections created by this solution, navigate to the **Layer 2 Connections** view; then look for the source from ROUTES.

Sta	Name	Topology Source	Status	Last Modified	Notes
	c2950-j1[Fa0/17],c2950-j2[Fa0/17],hp2626-1[Se0/0.1]	ROUTES	✓	Mar 31, 2011 12:00:02 PM	
	Frame Relay Interfaces-mp1sce02[Se0/0.1],mp	SUBNETCONNECTION	✓	Mar 31, 2011 11:59:02 AM	
	Frame Relay Interfaces-cisco2k1[Se0/0.1],cis	SUBNETCONNECTION	✓	Apr 20, 2011 11:51:21 AM	
	Serial Point to Point-cisco4k1[Hs0],cisco4k2[H	SUBNETCONNECTION	✗	Apr 20, 2011 11:49:59 AM	
	Frame Relay Interfaces-cisco4k1[Se0.1],cisco	SUBNETCONNECTION	✓	Apr 20, 2011 11:54:13 AM	

Disabling the Unnumbered Interface Feature

If you decide to disable the unnumbered interface feature, complete these steps:

- 1 Remove the following file:

Windows: %NNM_DATA%\shared\nnm\conf\disco\UnnumberedNodeGroup.conf

UNIX: \$NNM_DATA/shared/nnm/conf/disco/UnnumberedNodeGroup.conf

- 2 Remove the following file if it exists:

Windows: %NNM_DATA%\shared\nnm\conf\disco\UnnumberedSubnets.conf

UNIX: \$NNM_DATA/shared/nnm/conf/disco/UnnumberedSubnets.conf

- 3 Restart the NNMi management server.

a Run the **ovstop** command on the NNMi management server.

b Run the **ovstart** command on the NNMi management server.

- 4 Wait for NNMi to complete the next discovery cycle.

See the `UnnumberedNodeGroup.conf` and `UnnumberedNodeGroup.conf` reference pages, or the UNIX manpage, for more information.

Controlling Deletion of Unresponsive Objects

You can control the deletion of the following unresponsive objects by specifying the number of days to wait after an object has become unresponsive:


- Unresponsive nodes
- Connections that are down

To control the deletion of unresponsive objects, perform the following steps:

- 1 In the **Configuration** workspace, click **Discovery Configuration**.
- 2 In the **Delete Unresponsive Objects Control** area, enter the numbers of days for the system to wait before deleting the applicable objects. Note that a value of zero (0) indicates that the objects should not be deleted.

After the specified waiting period, the unresponsive objects are deleted from the database.

NNMi State Polling



This chapter provides information to help you expand and fine tune network monitoring by configuring the HP Network Node Manager i Software (NNMi) State Poller service. This chapter supplements the information in the NNMi help. For an introduction to how monitoring works and for detailed information about how to configure monitoring, see *Monitoring Network Health* in the NNMi help.

This chapter contains the following topics:

- [Concepts for State Polling](#)
- [Plan State Polling](#)
- [Configure State Polling](#)
- [Evaluate State Polling](#)
- [Tune State Polling](#)

Concepts for State Polling

This section provides a brief overview of network monitoring, including the order that the State Poller uses to evaluate polling groups. After reading the information in this section, continue to [Plan State Polling](#) on page 78 for more specific information.

As with network discovery, you should focus network monitoring on the critical or most important devices in the network. NNMi can only poll devices in the topology database. You control which network devices NNMi monitors, the type of polling to use, and the interval at which to poll.

You can use the interface and node settings on the **Monitoring Configuration** form to refine the status polling of devices, and to set different polling types and intervals for different classes, types of interfaces, and types of nodes.

You can configure State Poller data collection to be based on an ICMP (ping) response, or to be based on SNMP data. NNMi automatically handles the mapping from the type of data collection you enable to the actual MIB objects internally, significantly simplifying configuration.

As you plan polling configuration, you should carefully consider how to set up interface groups and node groups for the State Poller service. If you are new to the concept of *groups*, see [Node Groups and Interface Groups](#) on page 35, and [Node/Interface/Address Hierarchy](#) on page 39 for overview information.

Order of evaluation

Because an interface or node might qualify for multiple groups, the State Poller applies the configured polling interval and polling type in a well-defined order of evaluation. For each object in the discovered topology:

- 1 If the object is an interface, State Poller looks for a qualifying interface group. Groups are evaluated from the lowest Order Number to the highest. The first matching group is used and evaluation stops.
- 2 If no interface group has captured the object, node groups are evaluated from lowest Order Number to highest. The first matching group is used and evaluation stops. Any contained interface which has not qualified for an interface group on its own characteristics inherits the polling settings from its hosting node.
- 3 For devices that are discovered but not included in any node or interface settings definitions, the global monitoring settings (on the **Default Settings** tab of the **Monitoring Configuration** form) establish the monitoring behavior.

Plan State Polling

This section provides information to plan for State Poller configuration, including a polling configuration checklist; and more detailed information to help you plan for monitoring, decide how to create polling groups, and determine what types of data should be captured during the polling process.

Polling Checklist

You can use the checklist below to plan for State Poller configuration.

- What can NNMi monitor?
- What are the logical groups for monitored items, based on object type, location, relative importance, or other criteria?
- How often should NNMi monitor each grouping?
- What data should be collected to capture information about the monitored item? This might include:
 - ICMP (ping) response
 - SNMP fault data
 - SNMP performance data if you have a license for one or more NNM Performance iSPIs
 - Additional SNMP Component Health data

Example polling configuration

To help you understand the polling configuration process, consider this example. Suppose that your network contains the latest proxy servers from ProximiT. You must ensure that these devices can be reached, but you do not require SNMP monitoring of the proxy servers.

1 What can NNMi monitor?

Because you can only monitor what has been discovered, you configure auto-discovery rules to ensure that NNMi's database contains your ProximiT proxy servers. For more information on configuring discovery, see [NNMi Discovery](#) on page 57.

2 What are the logical groups for monitored items?

It makes sense to group the ProximiT proxy servers together and apply the same monitoring settings to all of them. Because you are not doing interface (SNMP) monitoring for the devices, you do not need any interface groups.

You can also use this node group to filter views, to check the status of the proxy servers as a group, and to put the group OUT OF SERVICE to update firmware.

3 How often should NNMi monitor each group?

For your service level agreements, a five minute polling interval for the proxy servers is sufficient.

4 What data should be collected?

Here's where the monitoring configuration differs from other groups. For our ProximiT proxy server example, you enable ICMP fault monitoring and disable SNMP fault and polling monitoring. Without SNMP fault monitoring for the group, Component Health monitoring will not apply.

For more detailed planning information concerning these configuration choices, see the following topics:

- [What Can NNMi Monitor?](#) on page 79
- [Planning Groups](#) on page 81
- [Planning Polling Intervals](#) on page 83
- [Deciding What Data to Collect](#) on page 84

What Can NNMi Monitor?

By default, the NNMi State Poller uses SNMP polls to monitor the following:

- Interfaces that are connected to another known interface on an NNMi-discovered device.
- Router interfaces which host IP addresses.



In most cases, polling only connected interfaces provides sufficiently accurate root-cause analysis. Extending the set of monitored interfaces can impact polling performance.

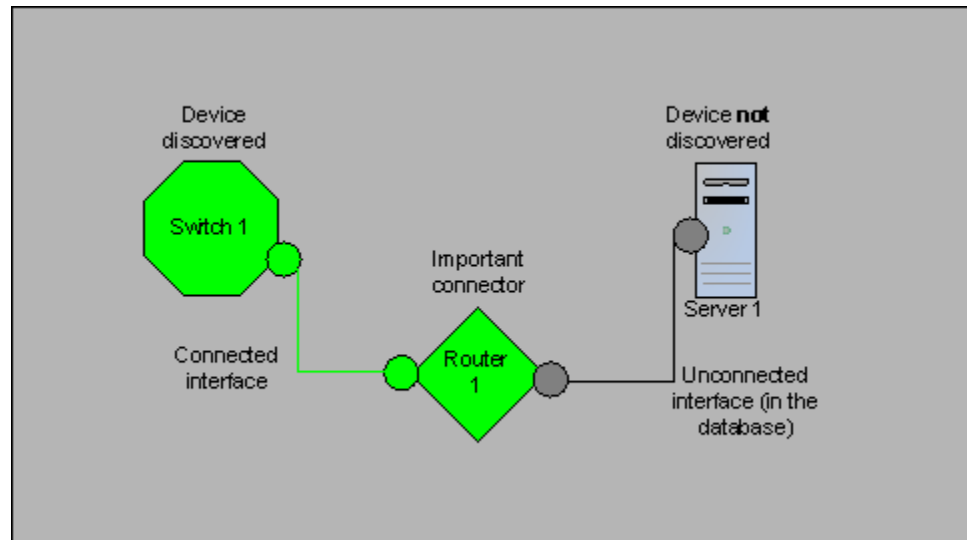
Extend monitoring

You can extend the monitoring to include the following:

- Unconnected interfaces. By default, the only unconnected interfaces that NNMi monitors are those that have IP addresses *and* are included in the **Routers** node group.



NNMi defines an unconnected interface as an interface that is not connected to another device discovered by NNMi, as shown below.



- Interfaces, such as router interfaces, that have an IP address.
- ICMP polling for devices that do not support SNMP. By default, ICMP polling is enabled for the **Non-SNMP Devices** node group.

Interfaces to Unmonitored Nodes

Sometimes you want to know the status of an interface that connects to a device you do not manage directly. For example, you want to know whether the connection to an application or Internet server is up, but you might not be responsible for maintaining that server. If you do not include the server in the discovery rules, NNMi sees the interface that faces the server as unconnected.

There are two ways to monitor the status of an important interface that connects to an unmonitored node.

- Discover the unmonitored node

When you add an unmonitored node to the NNMi topology, NNMi sees the interfaces connecting the node to the rest of the topology as `CONNECTED`. Then NNMi can poll these interfaces according to the monitoring configuration. NNMi discovers the node as `MANAGED`. Unmanage nodes that you do not want NNMi to monitor.



Each discovered node counts toward the license limit, regardless of whether NNMi is actively managing that node.

- Poll the unconnected interface

You can create a node group containing the network devices that provide connectivity for undiscovered nodes. Then enable polling of unconnected interfaces for the node group.

NNMi polls *all* interfaces on the devices in the node group, which can add a lot of traffic for a device with many interfaces.

Stop Monitoring

The NNMi management modes are used to set devices or interfaces to `UNMANAGED` or `OUT OF SERVICE`. `UNMANAGED` is considered to be a permanent situation; you will never care to know the status of the object. `OUT OF SERVICE` is for temporary situations where one or more objects will be offline and down incidents would be superfluous.

Consider the management mode as an overlay across all group settings. Regardless of its group, polling interval, or type, the State Poller does not communicate with an object when its status is set to `UNMANAGED` or `OUT OF SERVICE`.

Best practice

Some of the devices, interfaces, or both you choose to discover and place in the database do not need to be polled. Note those objects which you will permanently set to `UNMANAGED`. You might want to create one or more node groups to enable you to set management modes more easily.

Planning Groups

You must set up node and interface groups before configuring monitoring settings. Therefore, you must consider polling requirements while configuring node and interface groups. Ideally, node and interface groups are configured so that you can monitor important devices frequently, and you can check on non-critical devices less frequently (if at all).

Best practice

Configure one set of node and interface groups for network monitoring. Configure a different set of node groups for network visualization through maps.

These groups are defined through the **Configuration > Node Groups** or **Configuration > Interface Groups** work spaces and are, by default, the same groups that are used to filter incident, node, interface, and address views. To create a separate set of node or interface filters for configuring monitoring settings, open a node or interface group and select the **Add to View Filter List** check box on the **Node Group** or **Interface Group** form. Click **Save and Close**.

You can set polling types and polling intervals at a node group or interface group level on the **Node Settings** and **Interface Settings** tabs of the **Monitoring Configuration** form.

Determine the criteria by which you want to group interfaces, devices, or both by similar polling needs. Here are some factors to consider in your planning:

- Which area of your network contains these devices? Are there timing constraints?
- Do you want to differentiate polling intervals or data gathered by device type? By interface type?
- Does NNMi provide pre-configured groups you can use?

Best practice

You can create group definitions for objects that are likely to go OUT OF SERVICE at the same time, whether by location or some other criteria. For example, you could put all your Cisco routers into OUT OF SERVICE mode while you apply an IOS upgrade.

Interface Groups

Based on your criteria, determine which Interface groups to create. Remember that interface groups are evaluated first (see [Concepts for State Polling](#) on page 77). Interface groups can reference node group membership, so you might end up configuring node groups before interface groups to implement your plan.

Preconfigured interface groups

NNMi has several useful interface groups already configured for you to use. These include:

- All interfaces with an IFType related to ISDN connections
- Interfaces for voice connections
- Interfaces for point-to-point communication
- Software loopback interfaces
- VLAN interfaces
- Interfaces participating in link aggregation protocols

Over time HP might add more default groups to simplify your configuration tasks. You can use existing groups, modify them, or create your own.

Interface groups have two types of qualifiers: node group membership for the hosting node and IFType or other attribute for the interface. You can choose to combine these as follows:

- All interfaces on nodes in a node group are grouped regardless of IFType; do not select any IFTypes or attributes (such as name, alias, description, speed, index, address, or other IFType attributes).
- All interfaces of certain IFTypes or set of attributes are grouped, regardless of the node on which they reside.
- Only interfaces of a certain IFType or attributes that reside on a particular group of nodes are grouped.

Node Groups

After planning interface groups, plan node groups. Not all node groups created for monitoring make sense for filtering views, so you can configure them independently.

Preconfigured node groups

HP provides a default collection of node groups to simplify your configuration tasks. These are based on device categories derived from the system object ID during the Discovery process. The node groups provided by default include:

- Routers
- Networking Infrastructure Devices (such as switches or routers.)
- Microsoft Windows Systems
- Devices for which you do not have the SNMP community string
- Important Nodes. This is used internally by the Causal Engine to provide special handling for devices in the “shadow” of a connector failure. For more information, see *Node Groups As Predefined View Filters* in the NNMi help.

Over time HP might add more default groups to simplify your configuration tasks. You can use existing groups, modify them, or create your own.

You can qualify the definition of related nodes using the following node attributes:

- IP address(es) on the node
- Hostname wildcard convention
- Device Profile derivatives such as category, vendor, and family
- MIB II sysName, sysContact, sysLocation

Best practice

You can create simple, reusable, atomic groups and combine them into hierarchical clusters for monitoring or visualization. Group definitions can overlap, such as “All Routers” and “All systems with IP address ending in .100.” Nodes will probably qualify for multiple groups as well.

Find a balance by creating a rich set of groups for configuration and viewing without overloading the list with superfluous entries that will never be used.

Interaction with Device Profiles

When each device is discovered, NNMi uses its system object ID to index into the list of available Device Profiles. The Device Profile is used to derive additional attributes of the device, such as vendor, product family, and device category.

As you configure node groups, you can use these derived attributes to categorize devices to apply monitoring settings. For example, you might want to poll all switches regardless of vendor throughout your network on a certain polling interval. You can use the derived device category, Switch, as the defining characteristic of your node group. All discovered devices whose system object ID maps to the category, Switches, will receive the configured settings for the node group.

Planning Polling Intervals

For each object group, you select a polling interval that NNMi uses to collect data. The interval can be as short as one minute, or as long as days to best match your Service Level Agreements.

Best practice

Shorter intervals help you become aware of network problems as soon as possible; however, polling too many objects in too short an interval can cause a backlog in the State Poller. Find the best balance between resource utilization and intervals for your environment.



The Causal Engine performs a Status Poll of each node every 24 hours and updates Status, Conclusion, and Incident information as needed. This Status Poll does not affect the timing of the Polling interval configured for the device.

Deciding What Data to Collect

The State Poller service uses polls to gather state information about the monitored devices in your network. Polling can be done using ICMP, SNMP, or both.

ICMP (ping)

ICMP address monitoring uses ping requests to verify the availability of each managed IP address.

SNMP

SNMP monitoring verifies that each monitored SNMP agent is responding to SNMP queries.

- The State Poller is highly optimized to collect configured SNMP information from each monitored object with one query at each interval. When you save configuration changes, the State Poller recalculates the group membership of each object and reapplies the configured interval and set of data to collect.
- SNMP monitoring issues SNMP queries for all monitored interfaces and components, requesting the current values from the MIB II interface table, the HostResources MIB, and vendor-specific MIBs. Some values are used for fault monitoring. If you have the NNM iSPI Performance for Metrics installed, some values are used for performance measurement.

SNMP Component Health data

You might enable or disable Component Health monitoring at the global level. Component Health monitoring for faults follows the fault polling interval settings for the device.

Gathering additional data at each poll does not affect the time to execute the poll. However, additional data stored for each object can increase the memory requirements for State Poller.



Performance monitoring settings are only used with the NNM iSPI Performance for Metrics. Component Health monitoring for performance follows the performance polling interval settings for the device.

Best practice

Batching your monitoring configuration changes is less disruptive to State Poller ongoing operation.

Configure State Polling

This section provides configuration tips and provides some configuration examples. After reading the information in this section, see *Configure Monitoring Behavior* in the NNMi help for specific procedures.



It is a good idea to save a copy of the existing configuration before you make any major configuration changes. For more information, see [Best Practice: Save the Existing Configuration](#) on page 34.

Configure Interface Groups and Node Groups

You create interface groups and node groups in the **Configuration** workspace. For more information, see *Creating Groups of Nodes or Interfaces* in the NNMi help.

Examples For example, to configure a node group for ProximiT proxy servers:

- 1 Open **Configuration > Node Groups** and click **New**.
- 2 Name the group **Proxy Servers** and check **Add to View Filter List**.
- 3 On the **Additional Filters** tab, select the **hostname** attribute, and leave the operator set to **=**.
- 4 For value, enter the wildcard as **prox*.example.com**.
If you had configured a device profile and device category for the ProximiT devices, you could use the **Device Filters** tab to access the **Device Category** selector and base the group on the Proxy Server category you created.
- 5 Click **Save and Close** on the group definition.



You must configure node groups before you can reference them in your interface group configuration.

Configure Interface Monitoring

State Poller analyzes interface group membership before node groups. For each of the interface groups you created, as well as any of the preexisting ones you want to use, open the **Monitoring Configuration** dialog and the **Interface Settings** tab to create a custom set of instructions for how State Poller should handle that group. Your instructions will include:

- Enabling or disabling fault polling
- Setting the fault polling interval
- Enabling or disabling performance polling if you have the NNM iSPI Performance for Metrics
- Setting the performance polling interval if you have the NNM iSPI Performance for Metrics
- Setting performance management thresholds if you have the NNM iSPI Performance for Metrics
- Selecting whether NNMi should monitor unconnected interfaces (or unconnected interfaces hosting IP addresses) in the group

You can configure different settings for each interface group. Remember that the State Poller evaluates the list in order from the lowest ordering number to the highest ordering number.

Best practice

Double-check your order numbers, keeping in mind that an object that qualifies for multiple groups has settings applied from the group with the lowest order number.

Configure Node Monitoring

If an object does not qualify for any configured interface group, State Poller evaluates the object for membership in node groups. Settings are applied to the first node group match from the lowest ordering number to the highest ordering number.

For each node group, open the **Monitoring Configuration** form, and then, open the **Node Settings** tab. Create a custom set of instructions as to how State Poller should handle that group. Your instructions can include:

- Enabling or disabling fault polling
- Setting the fault polling interval
- Enabling or disabling performance polling if you have the NNM iSPI Performance for Metrics
- Setting the performance polling interval if you have the NNM iSPI Performance for Metrics
- Setting performance management thresholds if you have the NNM iSPI Performance for Metrics
- Selecting whether NNMi should monitor unconnected interfaces (or unconnected interfaces hosting IP addresses) in the group

You might configure different settings for each node group.

Best practice

Double-check the order numbers, keeping in mind that an object that qualifies for multiple groups has settings applied from the group with the lowest order number.

Verify Default Settings

State Poller applies the settings from the **Default Settings** tab for any object that does not match a defined interface setting or node setting. Review the settings on this tab to ensure they match your environment at the default level. For example, you would rarely poll all unconnected interfaces as a default setting.



Be sure you **Save and Close** all **Monitoring Configuration** dialog boxes all the way back to the console for your changes to be implemented.

Evaluate State Polling

This section lists ways to evaluate the progress and success of the monitoring settings.

Verify the Configuration for Network Monitoring

You can determine the settings that NNMi uses for monitoring a given node or interface, and you can initiate a status poll of a node at any time.

Is the interface or node a member of the right group?

You can verify which interfaces or nodes belong to a group by selecting one of the following in the **Configuration** workspace:

- Node Groups
- Interface Groups

Follow the instructions in the help to show the members of the group. Keep in mind that an object can be a member of multiple groups, and that another group might have a lower ordering number.

Alternatively, you can see the full list of groups to which the object belongs by opening the object (interface or node) and clicking the **Node Groups** or **Interface Groups** tab. This list is alphabetical by group name and does not reflect the ordering numbers that determine which settings are applied.

If the object is not a member of a group:

- 1 Retrieve the device profile for the node in the inventory view.
- 2 Review the attribute mapping for the device profile under **Configuration > Device Profiles**.
- 3 Review the attribute requirements for the node group definition.

If you have a mismatch, you can adjust the category derived in the Device Profile to force that type of device to qualify for your node group. You might need to do an **Actions > Configuration Poll** to update the attributes for the node so that it qualifies.

Which settings are being applied?

To check the monitoring configuration in effect for a specific node, interface, or address, select that object in the appropriate inventory view, and select **Actions > Monitoring Settings**. NNMi displays the current monitoring settings.

Examine the values for **Fault Polling Enabled** and **Fault Polling Interval**. If these values are not as expected, look at the value for **Node Group** or **Interface Group** to see which ordered group match applied.

You might need to check **Actions > Communication Settings** for the object to ensure traffic has not been disabled for it.

Which data is being collected?

You can initiate a status poll of a specific device to validate that the expected types of polls (SNMP, ICMP) are being performed for that device. Select a node, and then click **Actions > Status Poll**. NNMi performs a real-time status check of the device. The output shows the types and results of the polls being performed. If the types of polls are not what you expect, check the monitoring settings for the node and the respective global, interface, or node settings of the monitoring configuration.

Evaluate the Performance of Status Polling

Evaluate the performance of status polling in your environment by using the information in the state poller health check to quantify and assess the operation of the state poller service.

Is the State Poller keeping up?

At any time, you can check the current health statistics about the state poller service on the **State Poller** tab of the **System Information** window, as described in [Table 2](#).

Table 2 State Poller Health Information

Information	Description
Status	Overall status of the state poller service
Poll counters	<ul style="list-style-type: none"> • Collections requested in last minute • Collections completed in last minute • Collections in process
Time to execute skips in last minute	<p>The number of regularly scheduled polls that did not complete within the configured polling interval. A non-zero value indicates that the polling engine is not keeping up or that targets are being polled faster than they can respond.</p> <ul style="list-style-type: none"> • What to watch for: If this value continues to increase, there are problems communicating with the target or NNMi is overloaded. • Action to take: Look in the <code>nmn.?.?.log</code> file for messages for the classes beginning with the string <code>com.hp.ov.nms.statepoller</code> to determine the targets for the skipped polls. <ul style="list-style-type: none"> — If the skipped polls are for the same targets, change the configuration to poll these targets at a less frequent rate or to increase the timeout for these targets. — If the skipped polls are for different targets, check the NNMi system performance, especially the available memory for <code>ovjboss</code>.
Stale collections in last minute	<p>A stale collection is a collection that has not received a response from the polling engine for at least 10 minutes. A healthy system should never have any stale collections.</p> <ul style="list-style-type: none"> • What to watch for: If this value increases consistently, there is a problem with the polling engine. • Action to take: Look in the <code>nmn.?.?.log</code> file for messages for the classes beginning with the string <code>com.hp.ov.nms.statepoller</code> to determine the targets for the stale collections. <ul style="list-style-type: none"> — If the stale collections are for a single target, unmanage the target until you can resolve the problem. — If the stale collections are for different targets, check the performance of the NNMi system and the NNMi database. Stop and restart NNMi.
Poller result queue length	<ul style="list-style-type: none"> • What to watch for: This value should be close to 0 most of the time. • Action to take: If this queue size is very large, <code>ovjboss</code> might be running out of memory.
State mapper input queue length	<ul style="list-style-type: none"> • What to watch for: This value should be close to 0 most of the time. • Action to take: If this queue size is very large, then check the performance of the NNMi system and the NNMi database.
State updater queue time length	<ul style="list-style-type: none"> • What to watch for: This value should be close to 0 most of the time. • Action to take: If this queue size is very large, then check the performance of the NNMi system and the NNMi database.

Tune State Polling

The performance of state polling is affected by the following key variables:

- The number of devices/interfaces to be polled
- The type of polling configured
- The frequency of polling each device


These variables are driven by your network management needs. If you are experiencing performance issues with status polling, consider the following configurations:

- Because polling settings for individual nodes are controlled through their membership in node groups and interface groups, make sure that the groups contain nodes or interfaces with similar polling requirements.
- If you are polling unconnected interfaces or interfaces that host IP addresses, check the configurations to make sure you are only polling the interfaces that are necessary. Enable these polls on the **Node Settings** or **Interface Settings** form (not as a global setting on the **Monitoring Configuration** form) to maintain the most specific control and to select the smallest subset of interfaces to poll.
- Remember that polling unconnected interfaces monitors *all* unconnected interfaces. To monitor only those unconnected interfaces that have IP addresses, enable polling of interfaces that host IP addresses.

Regardless of the monitoring configuration, status polling is dependent on network responsiveness and might be impacted by overall system performance. Although status polling with default polling intervals does not introduce much network load, if the performance of the network link between the server and the polled device is poor, status polling performance is poor. You can configure larger timeouts and a smaller number of retries to reduce the network load, but these configuration changes only go so far. Timely polling requires adequate network performance and sufficient system resources (CPU, memory).

Enabling or disabling the Component Health monitoring has no effect on timeliness of polling. It simply gathers additional MIB objects at the schedule time. However, disabling Component Health monitoring might reduce the amount of memory used by the State Poller.

NNMi Incidents



HP Network Node Manager i Software (NNMi) provides a large number of default incidents and correlations that filter incoming SNMP traps to provide a workable number of incidents in the NNMi console. This chapter provides information to help you fine tune network management by configuring the NNMi incidents. This chapter supplements the information in the NNMi help. For an introduction to NNMi incidents and for detailed information about how to configure incidents, see *Configuring Incidents* in the NNMi help.

This chapter contains the following topics:

- [Concepts for Incidents](#)
- [Plan Incidents](#)
- [Configure Incidents](#)
- [Batch Load Incident Configurations](#)
- [Evaluate Incidents](#)
- [Tune Incidents](#)

Concepts for Incidents

NNMi collects network status information from the following sources:

- The NNMi Causal Engine analyzes the health of your network and provides the ongoing health status reading for each device. The Causal Engine also extensively evaluates and determines the root cause of network problems whenever possible.
- SNMP traps from network devices. The NNMi Causal Engine uses this information as symptoms during its analysis.
- NNM 6.x/7.x events forwarded from one or more NNM 6.x/7.x management stations.
- Syslog messages from HP ArcSight Logger integration.

NNMi converts this network status information into incidents that provide useful information for managing the network. NNMi provides many default incident correlations that reduce the number of incidents for network operators to consider. You can customize the default incident correlations and create new incident correlations to match the network management needs of your environment.

The incident configurations in the NNMi console define the incident types that NNMi can create. If no incident configuration matches a received SNMP trap or NNM 6.x/7.x event or syslog message, that information is discarded. If the management mode of the source object is set to NOT MANAGED or OUT OF SERVICE in the NNMi database, or if the device is not monitored for fault polling, NNMi always discards the incoming trap.



`nmtrapconfig.ovpl -dumpBlockList` outputs information about the current incident configuration, including SNMP traps that were not passed into the incident pipeline because of non-existent or disabled incident configurations.

Additionally, NNMi discards SNMP traps from network devices that are not in the NNMi topology. For information about changing this default behavior, see *Handle Unresolved Incoming Traps* in the NNMi help.

For more information, see the following:

- *About the Event Pipeline* in the NNMi help
- *The NNMi Causal Engine and Incidents* in the NNMi help
- *HP Network Node Manager i-series Software Causal Analysis White Paper*, available from <http://h20230.www2.hp.com/selfsolve/manuals>

Incident Lifecycle

Table 3 describes the stages of an incident's lifecycle.

Table 3 NNMi Incident Lifecycle

Lifecycle State	Description	State Set By	Incident Used By
none	The NNMi event pipeline receives input from all sources and creates incidents as needed.	not applicable	<ul style="list-style-type: none"> • NNMi
Dampened	<p>The incident is in a holding place waiting to be correlated with another incident. The purpose of this waiting period is incident reduction in the incident viewers.</p> <p>The dampening interval can vary per incident type. For more information, see Incident Suppression, Enrichment, and Dampening on page 98.</p>	NNMi	<ul style="list-style-type: none"> • NNMi
Registered	<p>The incident is visible in incident views.</p> <p>The incident is forwarded to any configured destinations (northbound or global manager).</p>	<p>NNMi</p> <p>A user can also set this state in an incident view.</p>	<ul style="list-style-type: none"> • Users • Lifecycle transition actions • Integrations that forward incidents

Table 3 NNMi Incident Lifecycle (cont'd)

Lifecycle State	Description	State Set By	Incident Used By
In Progress	The incident has been assigned to someone who is investigating the problem. The network administrator defines the specific meaning of this state.	User	<ul style="list-style-type: none"> • Users • Lifecycle transition actions • Integrations that forward incidents
Completed	Investigation of the problem indicated by the incident is complete, and a solution is in place. The problem that the incident identifies The network administrator defines the specific meaning of this state.	User	<ul style="list-style-type: none"> • Users • Lifecycle transition actions • Integrations that forward incidents
Closed	Indicates that NNMi determined the problem reported by this Incident is no longer a problem. For example, when you remove an interface from a device, all incidents related to the interface are automatically closed.	User or NNMi	<ul style="list-style-type: none"> • Users • Lifecycle transition actions • Integrations that forward incidents

Trap and Incident Forwarding

Table 4 summarizes the ways to forward traps and incidents from the NNMi management server to another destination. The text following the table compares the NNMi SNMP trap forwarding mechanism with the NNMi northbound interface SNMP trap forwarding mechanism.

Table 4 Supported Ways to Forward Traps and NNMi Incidents

	NNMi Trap Forwarding	NNMi Northbound Interface Trap Forwarding	Global Network Management Trap Forwarding
What to forward	<ul style="list-style-type: none"> • SNMP traps from network devices • syslog messages from HP ArcSight Logger 	<ul style="list-style-type: none"> • SNMP traps from network devices • NNMi management events • syslog messages from HP ArcSight Logger 	<ul style="list-style-type: none"> • SNMP traps from network devices • NNM 6.x/7.x events from NNM management stations • syslog messages from HP ArcSight Logger
Forwarding format	SNMPv1, v2c, or v3 traps, as received (SNMPv3 traps can be converted to SNMPv2c traps)	SNMPv2c traps created from NNMi incidents	NNMi incidents
Added information	In most cases, NNMi adds varbinds to identify the original source object. NNMi does not ever modify SNMPv1 traps.	NNMi adds varbinds to identify the original source object.	Any information added to the incident by the regional manager processes is retained in the forwarded incident.

Table 4 Supported Ways to Forward Traps and NNMi Incidents (cont'd)

	NNMi Trap Forwarding	NNMi Northbound Interface Trap Forwarding	Global Network Management Trap Forwarding
Where to configure	Trap Forward Configuration in the Configuration workspace	HPOM, Northbound Interface, or Netcool in the Integration Module Configuration workspace	Forward to Global Managers tab on an SNMP Trap Configuration form, a Remote NNM 6.x/7.x Event Configuration form, or syslog configuration.
Notes		NNMi provides several integrations built on the NNMi northbound interface. See the NNMi Integration Module for Netcool Software, the HP NNMi—HPOM Integration, and the HP NNMi—HPOM Integration Chapters in the <i>NNMi Integration Reference</i> .	Forward the remote incidents that should be visible in the global manager incident views. Forwarded incidents participate in correlations on the global manager.
For more information	<i>Configuring Trap Forwarding</i> in the NNMi help	See the NNMi Northbound Interface chapter in the <i>NNMi Integration Reference</i> .	<ul style="list-style-type: none"> • <i>Configure Forward to Global Manager Settings for an SNMP Trap Incident</i> in the NNMi help • <i>Configure Forward to Global Managers Settings for a Remote 6.x/7.x Event Incident</i> in the NNMi help

Comparison: Forwarding Third-Party SNMP Traps to Another Application

If you want to forward the SNMP traps that NNMi receives from managed devices to another application, you can use either of the following approaches:

- Use the NNMi SNMP trap forwarding mechanism. For information about how to configure NNMi SNMP trap forwarding, see *Configuring Trap Forwarding* in the NNMi help.
- Use the NNMi northbound interface SNMP trap forwarding mechanism. For information about configuring the NNMi northbound interface to forward received SNMP traps, the NNMi Northbound Interface chapter in the *NNMi Integration Reference*.

The approach to trap identification by the receiving application varies with the SNMP trap forwarding mechanism:

- *Windows (all) and UNIX without original trap forwarding*

This description applies to the default and SNMPv3 to SNMPv2c conversion forwarding options.

The NNMi SNMP trap forwarding mechanism on a Windows NNMi management server enriches each SNMP trap before forwarding it to the trap destination. The trap appears to originate from the NNMi management server. (This information also applies to a UNIX NNMi management server for which the original trap forwarding option is not selected on the **Trap Forwarding Destination** form.)

To ensure the correct association between the trap-sending device and the event in the receiving application, the rules for these traps must be customized for the enriched varbinds. Interpret the value from the `originIPAddress` (.1.3.6.1.4.1.11.2.17.2.19.1.1.3) varbind. The `originIPAddress` value is a byte string of generic type `InetAddress`, either `InetAddressIPv4` or `InetAddressIPv6` as determined by the value of `originIPAddressType` (.1.3.6.1.4.1.11.2.17.2.19.1.1.2) varbind. The rule must read the `originIPAddressType` varbind to determine the type of Internet address (`ipv4(1)`, `ipv6(2)`) value in the `originIPAddress` varbind. The rule might also need to convert the `originIPAddress` value to a display string.

For more information about the varbinds that NNMi adds to forwarded traps, see *Trap Varbinds Provided by NNMi* in the NNMi help, RFC 2851, and the following file:

- *Windows*: %NNM_SNMP_MIBS\Vendor\Hewlett-Packard\hp-nnmi.mib
- *UNIX*: \$NNM_SNMP_MIBS/Vendor/Hewlett-Packard/hp-nnmi.mib

- *UNIX with original trap forwarding*

The NNMi SNMP trap forwarding mechanism on a UNIX NNMi management server can forward the traps in the same format as NNMi receives them. Each trap appears as if the managed device sent it directly to the trap destination, so existing trap processing configured in the receiving application should work without modification.

For more information, see the original trap forwarding option in *Trap Forwarding Destination Form* in the NNMi help.

- *NNMi northbound interface (all operating systems)*

The NNMi northbound interface enriches each SNMP trap before forwarding it to the trap destination. The trap appears to originate from the NNMi management server. To ensure the correct association between the trap-sending device and the event in the receiving application, the rules for these traps must be customized for the enriched varbinds. The `IncidentNodeHostname` (1.3.6.1.4.1.11.2.17.19.2.2.21) and `IncidentNodeMgmtAddr` (1.3.6.1.4.1.11.2.17.19.2.2.24) varbinds identify the original source object.

MIBs

NNMi requires that the following management information base (MIB) files be loaded into the NNMi database:

- All MIB variables used in MIB expressions for the Custom Poller feature, line graphs, or both
- Node components that NNMi monitors for health (for example fan or power supply)
- (NNM iSPI Performance for Metrics) All MIB variables used in threshold monitoring

NNMi requires that the following management information base (MIB) files, or the traps defined in the MIB files, be loaded into the NNMi database:

- All SNMP traps that you want to forward to a northbound destination
- (NNM iSPI NET) All MIB variables accessed from Trap Analytics reports



NNMi provides a `README.txt` file that lists those MIBs that are currently not supported. The `README.txt` file is located in the following directory:

- *Windows*: %NnmInstallDir%\misc\nnm\snmp-mibs
- *UNIX*: \$NnmInstallDir/misc/nnm/snmp-mibs

Custom Incident Attributes

NNMi uses custom incident attributes (CIAs) to attach additional information to incidents.

- For an SNMP trap incident, NNMi stores the original trap varbinds as CIAs for the incident.
- For a management event incident, NNMi adds pertinent information (for example, com.hp.ov.nms.apa.symptom) as CIAs for the incident.

You can use incident CIAs to narrow the scope of configurations such as incident lifecycle transition actions, suppression, deduplication, and enrichment. You can also use CIAs to narrow the availability of the menu items on the Actions menu for an incident view or form.

To determine which CIAs NNMi adds for any given incident, open a sample incident from an incident view, and look at the information on the Custom Attributes tab.

CIA's Added to Closed Management Event Incidents

When the NNMi Causal Engine determines that the conditions that caused a management event incident no longer apply, NNMi sets that incident's lifecycle state to CLOSED and adds the CIA's listed in [Table 5](#) to the incident. NNMi console users can see this information in the **Correlation Notes** field of the **Incident** form. Lifecycle transition actions can use the values of the CIA's directly.

Table 5 Custom Incident Attributes for a Closed Incident

Name	Description
cia.reasonClosed	The reason that NNMi cancelled or closed the incident. This reason is also the conclusion name, for example NodeUp or InterfaceUp. If this field is not set, an NNMi console user closed the incident. To determine the NNMi expected values of the cia.reasonClosed CIA, see <i>How NNMi Closes Incidents</i> in the NNMi help.
cia.incidentDurationMs	The duration, in milliseconds, of the outage, as measured by NNMi from when the status goes down and comes back up. This value is the difference of the cia.timeIncidentDetectedMs and cia.timeIncidentResolvedMs CIA's. It is a more accurate measurement than comparing the timestamps of down and up incidents.
cia.timeIncidentDetectedMs	The timestamp, in milliseconds, when the NNMi Causal Engine first detected the problem.
cia.timeIncidentResolvedMs	The timestamp, in milliseconds, when the NNMi Causal Engine detected that the problem has been resolved.

NNMi adds the CIA's listed in [Table 5](#) to most primary and secondary root cause incidents. For example, a NodeDown incident can have InterfaceDown and AddressDown incidents as secondary root causes. When NNMi closes the NodeDown incident, NNMi also closes the secondary incidents and adds the CIA's with values for each incident context to the secondary incidents.

NNMi does not add the CIA's listed in [Table 5](#) to the following default management event incident types:

- Incidents that an NNMi console user closes manually
- Incidents that NNMi closes in response to an object being deleted from the NNMi database
- IslandGroupDown incidents
- NnmClusterFailover, NnmClusterLostStandby, NnmClusterStartup, and NnmClusterTransfer incidents
- Incidents in the following families:
 - Correlation
 - License

- NNMi Health
- Trap Analysis

Incident Reduction

NNMi provides the following customizable correlations for reducing the number of incidents that network operators see in the NNMi console:

- Pairwise correlation—One incident cancels another incident.
- Deduplication correlation—When multiple copies of an incident are received within the specified time window, correlate the duplicates under a deduplication incident. The time window restarts for each newly received duplicate incident. In this way, NNMi correlates the duplicate incidents until it has not received any duplicates for the entire duration of the correlation time window.
- Rate correlation—When the specified number of copies on an incident are received within the specified time window, correlate the duplications under a rate incident. NNMi generates the rate incident when the specified number of incidents has been received, regardless of how much time remains in the time window.

Incident Suppression, Enrichment, and Dampening

NNMi provides a rich feature set for getting the most value from incidents. For each incident type, you can specifically define when an incident is of interest with the following incident configuration options:

- Suppression—When an incident matches the suppression configuration, that incident does not appear in the NNMi console incident views. Incident suppression is useful for incidents (for example, SNMPLinkDown traps) that are important for some nodes (for example routers and switches) but not others.
- Enrichment—When an incident matches the enrichment configuration, NNMi changes one or more incident values (for example, severity or message) according to the contents of the incident. Incident enrichment is useful for processing traps (for example, RMONFallingAlarm) that carry the distinguishing information in the trap varbinds (payload).
- Dampening—When an incident matches the dampening configuration, NNMi delays activity for that incident for the duration of the dampen interval. Incident dampening provides time for the NNMi Causal Engine to perform root cause analysis on the incident, which is useful for providing fewer, more meaningful incidents in the NNMi console.

For each incident type NNMi provides the following levels of configuration for suppression, enrichment, and dampening:

- Interface group settings—Specify incident behavior when the source object is a member of an NNMi interface group. You can specify different behavior for each interface group.
- Node group settings—Specify incident behavior when the source object is a member of an NNMi node group. You can specify different behavior for each node group.
- Default settings—Specify default incident behavior.

For each incident configuration area (suppression, enrichment, and dampening), NNMi uses the following procedure to determine the behavior of a specific incident:

- 1 Check the interface group settings:
 - If the source object matches any interface group settings, carry out the behavior defined in the match with the lowest ordering number and stop looking for a match.
 - If the source object does not match any interface group settings, continue with [step 2](#).
- 2 Check the node group settings:
 - If the source object matches any node group settings, carry out the behavior defined in the match with the lowest ordering number and stop looking for a match.
 - If the source object does not match any node group settings, continue with [step 3](#).
- 3 Carry out the behavior defined in the default settings, if any.

Lifecycle Transition Actions

A lifecycle transition action is an administrator-provided command that runs when an incident lifecycle state changes to match the action configuration. An incident action configuration is specific to one lifecycle state for one incident type. The action configuration identifies the command to run when this incident type transitions to the specified lifecycle state. The command can include arguments that pass incident information to the action code.

The action code can be any Jython file, script, or executable that runs correctly on the NNMi management server. The action code can be specific to one incident type, or it can process many incident types. For example, you might create action code that pages a network operator when NNMi creates a `ConnectionDown`, `NodeDown`, or `NodeOrConnectionDown` incident. You would configure three incident actions, one for the `REGISTERED` lifecycle state for each of these incident types.

Similarly, the action code can be specific to one lifecycle state change, or it can respond to several lifecycle state changes. For example, you might create action code that generates a trouble ticket when NNMi creates an `InterfaceDown` incident and closes the trouble ticket when the `InterfaceDown` incident is canceled. You would configure two incident actions for the `InterfaceDown` incident, one for the `REGISTERED` state and one for the `CLOSED` state.

Each action configuration can include a payload filter based on CIAs that limits when the action is run. For additional filtering, you can use incident enrichment to add a CIA to the incident. NNMi determines the value of that attribute from the incident source. For example, if you have added a custom attribute to some nodes, you can add this information to the incident as a CIA and then base the payload filter for an incident action on this attribute value.

Plan Incidents

Make decisions in the following areas:

- Which Device Traps Should NNMi Process?
- Which Incidents Should NNMi Display?
- How Should NNMi Respond to Incidents?
- Should NNMi Receive Traps from an NNM Management Station?
- Should NNMi Forward Traps to Another Event Receiver?

Which Device Traps Should NNMi Process?

Identify the device traps that are of interest in your network, and plan an incident configuration for each trap. NNMi can process traps without the MIB being loaded into NNMi. If the MIB contains TRAP-TYPE or NOTIFICATION-TYPE macros, you can create skeleton incident configurations for the traps defined in the MIB.

Decide whether you want to see traps from devices that are not in the NNMi topology.

Which Incidents Should NNMi Display?

The default set of incidents is a good place to start. You can expand and reduce the incident set over time.

Plan which incidents can be reduced through deduplication, rate configuration, and pairwise correlation.

How Should NNMi Respond to Incidents?

What actions (for example, sending an email message to a network operator) should NNMi take when certain incidents occur? At what lifecycle state should each action run?

Should NNMi Receive Traps from an NNM Management Station?

If your environment includes one or more NNM 6.x/7.x management stations that will continue managing areas of the network in conjunction with NNMi, identify the NNM 6.x/7.x events that will help NNMi operators manage the network. Plan an incident configuration for each NNM 6.x/7.x event that should be available in the NNMi console.

Should NNMi Forward Traps to Another Event Receiver?

If your environment includes a third-party trap consolidator, decide whether to use the NNMi SNMP trap forwarding mechanism with the NNMi northbound interface SNMP trap forwarding mechanism.

If you choose the NNMi northbound interface SNMP trap forwarding mechanism, load the MIBs for all traps that NNMi will forward to the event receiver.

Configure Incidents

This section lists configuration tips and provides some configuration examples. After reading the information in this section, see *Configuring Incidents* in the NNMi help for specific procedures.



It is a good idea to save a copy of the existing configuration before you make any major configuration changes. For more information, see [Best Practice: Save the Existing Configuration](#) on page 34.

- Configure the incident types that you planned. If possible, start with the skeleton incident configurations from the traps defined in the MIB.
- Load any MIBs that are required for trap forwarding.
- Verify that devices are configured to send traps to the NNMi management server.

Configuring Incident Suppression, Enrichment, and Dampening

While configuring incident suppression, enrichment, and dampening, note the following:

- For each interface group, node group, or default setting, you can specify a payload filter that further refines when the configuration is applicable.
- Configure interface group settings on the **Interface Settings** tab of an incident configuration form.
- Configure node group settings on the **Node Settings** tab of an incident configuration form.
- Configure default settings on the **Suppression, Enrichment, and Dampening** tabs of an incident configuration form.

Configuring Lifecycle Transition Actions

While configuring lifecycle transition actions, note the following:

- By default, NNMi runs actions in the following location:
 - *Windows*: %NnmDataDir%\shared\nnm\actions
 - *UNIX*: \$NNM_DATA/shared/nnm/actions

If an action is not in this location, specify the absolute path to the action in the **Command** field of the **Lifecycle Transition Action** form.



Jython files must be placed in the actions directory.

- Each time you make a change to the action configuration, NNMi rereads the actions directory for Jython files and loads them into NNMi.
- Actions are enabled as a group for an incident type.
- For information about the NNMi information that you can pass to an action, see *Valid Parameters for Configuring Incident Actions* in the NNMi help.

Configuring Trap Logs

NNMi provides the ability to log all incoming SNMP traps into a log file (either a text file or a CSV file). Traps are logged to the following location:

- *Windows*: %NnmDataDir%\nnm\log
- *UNIX*: \$NNM_DATA/nnm/log

Trap log files can be configured using the `nmtrapconfig.ovpl` script. The following format choices are available:

- *CSV (default)* – Traps are logged in the CSV format (`trap.csv`).
- *TXT* – Traps are logged in the TXT format (`trap.log`).
- *BOTH* – Traps are logged in both CSV and TXT (2 log files).
- *OFF* – No traps are logged.

For example, to specify that traps get logged into BOTH modes, you would use the following command:

```
nmtrapconfig.ovpl -setProp trapLoggingMode BOTH -persist
```

Note that the `-persist` argument causes all trap server properties to remain in effect even after the trap service is restarted. If you do not use the `-persist` argument, all trap server properties will be in effect only until the service is stopped.

Traps are written to a rolling file. After the log file size reaches the defined maximum limit (as defined using the `nmtrapconfig.ovpl` script), the file is renamed to `trap.<format>.old`. Any existing file is replaced.

See the `nmtrapconfig.ovpl` reference page, or the UNIX manpage, for more information. See also *Configure Trap Logging* in the NNMi help.

Configuring Incident Logging

You can configure incident logging so that incoming incident information is written to the `incident.log` file. This feature is useful when you want to track and archive your incident history.

Configure and enable incident logging by navigating to the **Incident Logging Configuration** tab in the **Incident Configuration** area of the **Configuration** workspace, and configuring the settings. For more information, see the NNMi help.

Configuring Trap Server Properties

You can set trap server properties (`nmtrapserver.properties`) by using the `nmtrapconfig.ovpl` script.



Although an `nmtrapserver.properties` file exists, do not edit this file directly; use the `nmtrapconfig.ovpl` script to modify the file.

The following default values exist for trap server properties:

Table 6 Trap Server Properties and Default Values

Trap Server Property	Default Value
com.hp.ov.nms.trapd.udpPort	162
com.hp.ov.nms.trapd.rmiPort	1097
com.hp.ov.nms.trapd.trapInterface	all interfaces
com.hp.ov.nms.trapd.recvSocketBufSize	2048 kilobytes
com.hp.ov.nms.trapd.pipeline.qSize	50000 traps
com.hp.ov.nms.trapd.connectToWinSNMP	false
com.hp.ov.nms.trapd.blocking	true
com.hp.ov.nms.trapd.blockTrapRate	50 traps/second
com.hp.nms.trapd.unblockTrapRate	50 traps/second
com.hp.ov.nms.trapd.overallBlockTrapRate	150 traps/second
com.hp.nms.trapd.overallUnblockTrapRate	150 traps/second
com.hp.ov.nms.trapd.analysis.minTrapCount	100 traps
com.hp.ov.nms.trapd.analysis.numSources	10 sources
com.hp.ov.nms.trapd.analysis.windowSize	300 seconds (5 minutes)
com.hp.nms.trapd.updateSourcesPeriod	30 seconds
com.hp.nms.trapd.notifySourcesPeriod	300 seconds
com.hp.ov.nms.trapd.hosted.object.trapstorm.enabled	false
com.hp.ov.nms.trapd.hosted.object.trapstorm.threshold	10 traps/second
com.hp.ov.nms.trapd.database.fileSize	100 megabytes
com.hp.ov.nms.trapd.database.fileCount	5 files
com.hp.ov.nms.trapd.database.qSize	300000 traps
com.hp.ov.nms.trapd.discohint.cacheSize	5000 entries
com.hp.ov.nms.trapd.discohint.cacheEntryTimeout	3600 milliseconds

See the *nmmtrapconfig.ovpl* reference page or the UNIX manpage for more information.

Batch Load Incident Configurations

You can use the following two scripts in conjunction with batch loading of incident configurations: `nmincidentcfgdump.ovpl` and `nmincidentcfgload.ovpl`.

Generating an Incident Configuration File with `nmincidentcfgdump.ovpl`

The NNMi `nmincidentcfgdump.ovpl` script provides a way for you to create or update an Incident Configuration to subsequently load into the NNMi database using the `nmincidentcfgload.ovpl` script. The file is generated in a non-xml format.

You can edit the file using the format descriptions provided in the following directory:

Windows: %NnmInstallDir%/examples/nnm/incidentcfg

UNIX: /opt/OV/examples/nnm/incidentcfg

To generate a file of your Incident Configurations, use the following example syntax:

```
nmincidentcfgdump.ovpl -dump <file_name> -u <NNMiadminUsername>
-p <NNMiadminPassword>
```

See the `nmincidentcfgdump.ovpl` reference page or the UNIX manpage for more information.

Loading Incident Configurations with `nmincidentcfgload.ovpl`

The NNMi `nmincidentcfgload.ovpl` script provides a way for you to load Incident Configurations into the NNMi database from a formatted configuration file.



Use the `nmincidentcfgdump.ovpl` script to create a configuration file of existing Incident Configurations in a non-xml format. You can then edit this file if desired before loading them into the NNMi database.

See the following directory for the required format:

Windows: %NnmInstallDir%/examples/nnm/incidentcfg

UNIX: /opt/OV/examples/nnm/incidentcfg

To validate an Incident Configuration file before it is loaded into the NNMi database, use the following example syntax:

```
nmincidentcfgload.ovpl -validate <file_name> -u <NNMiadminUsername>
-p <NNMiadminPassword>
```

To load Incident Configurations, use the following example syntax:

```
nmincidentcfgload.ovpl -load <file_name> -u <NNMiadminUsername>
-p <NNMiadminPassword>
```

Note the following:

- NNMi updates all configurations that have matching names or other matching key identifiers.



NNMi also overwrites the values of any codes associated with these configurations (for example, incident Family).

- NNMi adds all incident configurations with key identifiers that do not exist in the NNMi database.
- NNMi does not change existing incident configurations with key identifiers that do not match any in the exported file.
- NNMi resolves Universally Unique Object Identifiers (UUIDs) if they are not provided in the configuration file.
- If NNMi is unable to resolve a UUID, a UUID is created.

See the `nmincidentcfgload.ovpl` reference page or the UNIX manpage for more information.

Evaluate Incidents

This section lists ways to evaluate the incident configuration.

- Verify that NNMi receives traps from all managed devices in the network.

If NNMi is not receiving traps, verify the configuration of the firewall on the NNMi management server.



Some anti-virus software includes a firewall that is configured separately from the system firewall.

- Verify that the most important traps are converted to incidents.
- Verify that incident actions run at the correct lifecycle state transitions.
- Verify that NNMi is handling incidents as expected.

The **Actions > Incident Configuration Reports** menu contains several options for testing an existing incident against the current configuration of that incident type. Using one of these menu items does not change the incidents currently in the NNMi console.

Tune Incidents

Reduce the number of incidents in the NNMi console incident views. Use any of the following methods:

- Disable the incident configuration for any incident types that are not needed in the NNMi console.
- Set the management mode of the network objects that you do not need to monitor to `NOT MANAGED` or `OUT OF SERVICE`. NNMi discards most incoming traps from these nodes and their interfaces.
- Set NNMi to not monitor some network objects. NNMi discards most incoming traps from the source objects that are not monitored.
- Identify additional criteria for or relationships between incoming incidents. When these criteria or relationships occur, NNMi modifies the flow of incidents by recognizing the criteria or patterns of incoming management events or SNMP traps and nesting related incidents as correlated children.

Enabling and Configuring Incidents for Undefined Traps

NNMi drops undefined traps silently by default. As of NNMi 9.01, NNMi can identify any undefined SNMP traps that might be dropped.



If you have NNM iSPI NET licensed on the NNMi management server, use the Total Traps Received (by OID) report to research the dropped SNMP traps. See *Analyze Trap Information (NNM iSPI NET)* in the NNMi help for more information.

If you do not have NNM iSPI NET licensed on the NNMi management server, and want to see the missing traps as an incident, configure the Undefined SNMP Trap incident as follows:

- 1 Edit the following file:

- *Windows*: %NNM_PROPS%\nms-jboss.properties
- *UNIX*: \$NNM_PROPS/nms-jboss.properties

- 2 Look for look for a section in the file that resembles the following line:

```
#!com.hp.nnm.events.allowUndefinedTraps=false
```

Change this line as follows:

```
com.hp.nnm.events.allowUndefinedTraps=true
```

- 3 *Optional*. Specify the incident severity using the values explained within the `nms-jboss.properties` file. Look for a section in the file that resembles the following line:

```
#!com.hp.nnm.events.undefinedTrapsSeverity=NORMAL
```

Change this line as follows, substituting a defined severity value for *YourSpecifiedSeverity*.

```
com.hp.nnm.events.undefinedTrapsSeverity=YourSpecifiedSeverity
```

- 4 *Optional*. Specify the incident nature using the values explained within the `nms-jboss.properties` file. Look for a section in the file that resembles the following:

```
#!com.hp.nnm.events.undefinedTrapsNature=INFO
```

Change this line as follows, substituting a defined nature value for *YourSpecifiedNature*.

```
com.hp.nnm.events.undefinedTrapsNature=YourSpecifiedNature
```

- 5 Restart NNMi by running the following commands:

- a **ovstop**
- b **ovstart**

- 6 Review the list of undefined traps and create new incident configurations for those traps that you want to control. Enable the new incident if you want NNMi to display it and disable the new incident if you want NNMi to ignore it. See *Configuring SNMP Trap Incidents* in the NNMi help for more information.

NNMi Console

Use the information in this chapter to understand how to use the NNMi console to configure NNMi to function in specific ways.

This chapter contains the following topics:

- [A Practical Example of Using Node Groups](#)
- [Reducing the Maximum Number of Nodes Displayed in a Network Overview Map](#)
- [Reducing the Number of Displayed Nodes on a Node Group Map](#)
- [Configuring Gauges in the Analysis Pane](#)

A Practical Example of Using Node Groups

The following shows a practical example for configuring node groups.

My Network: A top level *container* node group containing other node groups.

USA: An intermediate *container* node group containing other node groups.

Colorado: A node group containing nodes located in Colorado.

Note the following:

- It is a best practice to design your node group map layout ahead of time.
- It is a best practice to configure one set of node and interface groups for network monitoring. Configure a different set of node groups for network visualization through maps.
- In this example, **Colorado** is the only node group that contains nodes.
- NNMi provides more than one way to configure node groups and node group maps. After you become familiar with the steps described in this document, you might find more efficient ways to create subsequent node groups and node group maps.

This document guides you through the following steps for configuring node groups and node group maps, and then deleting node groups:

Create Node Groups

- Step 1: Create the My Network Node Group
- Step 2: Create the USA Node Group
- Step 3: Create the Colorado Node Group Using Filters
- Step 4: View the Node Group Members to Check the Node Group Filter Results
- Step 5: Set Up the Node Group Hierarchy for the My Network Node Group
- Step 6: Establish the Node Group Hierarchy for the USA Node Group



Parent node groups might not contain any nodes. Instead they contain only child node groups in the definition. In this example, the *My Network* and *USA* node groups are parent node groups that contain only child node groups.

Configure the Node Group Maps

- Step 1: Create the Node Group Maps
- Step 2: View the Node Group Maps
- Step 3: Configure Node Group Status
- Step 4: Configure Node Group Map Ordering
- Step 5: Add a Background Image to a Node Group Map

Remove Node Groups

- Step 1: Navigate to the Node Group
- Step 2: Delete the Node Group

Create Node Groups

We begin by creating the Node Groups to include in our Node Group maps.

Step 1: Create the My Network Node Group

To create the **My Network** Node Group:

- 1 Navigate to the **Configuration** workspace.
- 2 Select **Node Groups**.
- 3 Click the **New** icon.
- 4 In the **Name** attribute, enter: **My Network**.
- 5 In the **Notes** attribute, enter: **This is the top level Node Group**.
- 6 Click **Save and Close** to save this configuration.

Step 2: Create the USA Node Group

- 1 Navigate to the **Configuration** workspace.

- 2 Select **Node Groups**.
- 3 Click the **New** icon.
- 4 In the **Name** attribute, enter: **USA**.
- 5 Click **Save and Close** to save this configuration.

Step 3: Create the Colorado Node Group Using Filters

To create the **Colorado** node group, use the Filter Editor to establish a filter to select the nodes.



When possible, use the **Additional Filters** tab rather than specifying a list of nodes using the **Additional Nodes** tab. Using a node group filter enables NNMi to automatically place a node into the correct node group as new nodes are added to the network.

- 1 Navigate to the **Configuration** workspace.
- 2 Select **Node Groups**.
- 3 Click the **New** icon.
- 4 In the **Name** attribute, enter: **Colorado**.
- 5 Select the **Additional Filters** tab.
- 6 Click **OR** to specify that you want NNMi to match a node if the node matches either of the hostname values you enter.
- 7 In the Filter Editor **Attribute** field, select `hostname`.
Selecting `hostname` specifies that NNMi should match hostname values when determining whether a node belongs to this node group.
- 8 In the **Operator** field, select `like`.
Selecting `like` enables you to use wildcard characters in the search.
- 9 In the **Value** field, enter a value that represents the devices you want the node group to contain. For example, `cisco*.ntc.example.com` represents devices named `cisco<replace with this text>.<network_domain>`.
- 10 Click **Append**.
- 11 In the **Attribute** field, select `hostname`.
- 12 In the **Operator** field, select `like`.
- 13 In the **Value** field, enter a wildcard that represents the remaining device names you want to add to the Colorado node group. For this example, use `cisco?*`.
- 14 Click **Append**.
- 15 Click **Save** to save the node group without closing the window.

Step 4: View the Node Group Members to Check the Node Group Filter Results

To test the node group filter, you can view the members of the node group you just created.

Select **Actions->Node Group Details->Show Members** to launch a view containing all of the nodes in the node group.



Examine the node group filter definition results until you are confident the node group filter is correct.

Step 5: Set Up the Node Group Hierarchy for the My Network Node Group

Establish a hierarchy for the node groups, starting with the top level node group, **My Network**.

- 1 Return to the **Node Groups** option in the **Configuration** workspace to view a list of the node groups you created.
- 2 Navigate to the **My Network** Node Group; then click **Open**.
- 3 Click the **Child Node Groups** tab.
- 4 Click the **New** icon.
- 5 In the **Child Node Group** attribute, click the **Lookup** icon and select **Quick Find**.



Use **Quick Find** to select an object, such as a node group, when it already exists.

- 6 Select **USA** as the child node group.
- 7 Click **OK**.
- 8 Click **Save and Close** to save your changes and close the **Node Group Hierarchy** form.
- 9 Click **Save and Close** to save your changes and close the **Node Group** form.

Step 6: Establish the Node Group Hierarchy for the USA Node Group

Next, establish **Colorado** as a child node group of the **USA** node group. Repeat the same steps described in [Step 5: Set Up the Node Group Hierarchy for the My Network Node Group](#) to make the Colorado node group a child of the USA Node Group.

You are ready to create the node group maps for each node group that you created.

Configure the Node Group Maps

Step 1: Create the Node Group Maps

To create node group maps for each node group, use the **Actions** menu.

- 1 Open the node group for which you want to create a map:
 - a Return to the **Node Groups** option in the **Configuration** workspace to view a list of the node groups you created.
 - b Navigate to the node group you want and click the **Open** icon.
- 2 Select the **Actions->Maps-> Node Group Map** to display a node group map.
- 3 Position the nodes and node group map icons.
- 4 Click the **Save Layout** icon to create the node group map.



Always use **Save Layout** to create the node group map, even if you do not change the node positions. **Save Layout** creates the node group map.

A dialog box appears confirming you successfully created the node group map.

- 5 Click **OK**.
- 6 Repeat steps 1 through 5 for each node group you created.

Step 2: View the Node Group Maps

Now that you have created the node group maps, view the maps to check the contents.

- 1 Navigate to the **Topology Maps** workspace.
- 2 Select **Node Group Overview**.
- 3 Select the top level map: **My Network**.
- 4 Navigate to the child node group maps by double-clicking its icon.
- 5 Use the breadcrumb trail above the toolbar to return to the previous map.

Step 3: Configure Node Group Status

NNMi enables you to configure how status is calculated for a node group. When you configure node group status, you determine which of the following method NNMi should use:

- Use the most severe status of the nodes in the node group.
- Specify the percentage calculation NNMi should use.



Status Configuration is a global configuration. By default, NNMi uses the most severe status of the nodes in the node group.

- 1 Navigate to the **Configuration** workspace.
- 2 Select **Status Configuration**.
- 3 Examine the **Status Configuration** form to become familiar with the default percentages. To use percentages, you must deselect the **Propagate Most Severe Status** option, then save your changes.

Step 4: Configure Node Group Map Ordering

Node group map ordering is used to help determine in what order a map is displayed under the **Topology Maps** workspace.

In this example, use node group map ordering to specify that the **My Network** node group map should appear first in the list in the **Topology Maps** workspace.

- 1 Navigate to the **Configuration** workspace.
- 2 Select **User Interface > Node Group Map Settings**.



As shown in the following example, the default **Topology Maps Ordering** value is 50 for all user-defined maps.

To indicate that NNMi should list **My Network** as the first map under the **Topology Maps** workspace, change the **Topology Maps Ordering** value to a number that is less than the **Topology Maps Ordering** value for any other maps in the list; for example **5**.

- 3 Open the **My Network** Node Group map.
- 4 In the **Topology Maps Ordering** attribute, change the value to **5**.
- 5 Click **Save and Close** to save your changes and close the form.

You can also specify whether the map is initially displayed in the NNMi console. To do so, use the **User Interface Configuration** option from the **Configuration** workspace.

- 1 Navigate to the **Configuration** workspace.
- 2 Click **User Interface Configuration**.
- 3 In the **Initial View** attribute, use the drop-down menu to select **First Node Group in Topology Maps** workspace.
- 4 Click **Save and Close** to save your changes and close the form.

This will make the **My Network** map the initial view.

To verify the initial view, sign out of NNMi and sign back in. The **My Network** map should be the view you see in the NNMi console.

Step 5: Add a Background Image to a Node Group Map

To include a background graphic on a map, use the **Node Group Map Settings** form for the selected node group map.

- 1 Navigate to the **Configuration** workspace.
- 2 Click **User Interface > Node Group Map Settings**.
- 3 Open the **My Network** node group map.
- 4 Navigate to the **Background Image** tab.
- 5 Click <http://MACHINE:PORT/nnmdocs/images/>.

NNMi displays a list of HP supplied graphics.

- 6 Right-click the **world.png** link.
- 7 Select **Copy Link Location**.
- 8 Close the directory listing window.

Paste the copied link into the **Background Image** attribute.



Note the **Background Image Scale** value in case you want to change it later.

- 9 Click **Save and Close** to save your changes.
- 10 Navigate to the **Topology Maps** workspace and select **My Network** to view your new map with the background graphic.

Remove Node Groups

Suppose we want to remove a node group. Say, for example, we would like to delete the Colorado node group created earlier in this example.

Step 1: Navigate to the Node Group

- 1 In the **Configuration** workspace, click **Node Groups**.
- 2 Select the **Colorado** node group in the list and click the **Open** button.

Step 2: Delete the Node Group

- 1 Click the **Delete Node Group** button.
- 2 A dialog box appears, warning you that all contained objects and references will also be deleted by deleting the node group.
- 3 Click **OK** to delete the node group.

Reducing the Maximum Number of Nodes Displayed in a Network Overview Map

The **Network Overview** map displays a map containing up to 250 of the most highly connected nodes in the layer 3 network. If this map contains too many nodes, the map might respond slowly when moving nodes or become too complex for practical viewing.

You can increase or reduce the maximum number of nodes displayed in the **Network Overview** map by editing the Maximum Number of Displayed Nodes attribute on the **Default Map Settings** tab on the User Interface Configuration form.

You can also increase or reduce the maximum number of nodes displayed in the **Network Overview** map by performing the steps shown in the following example.

Suppose you want to change the maximum number of nodes displayed in the **Network Overview** map from 250 to 100. To do this, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-ui.properties
 - *UNIX*: \$NNM_PROPS/nms-ui.properties

- 2 Look for text similar to following line:

```
#!com.hp.nnm.ui.networkOverviewMaxNodes = 250
```

Change the line as follows:

```
com.hp.nnm.ui.networkOverviewMaxNodes = 100
```



Make sure to remove the **#!** characters located at the beginning of the line.

- 3 Save your changes.

Reducing the Number of Displayed Nodes on a Node Group Map

If you configure a node group map to contain hundreds of nodes, the map showing the node group might show many small node icons instead of the detailed node icons you expect. To view the map with better detail, you would need to use the zoom feature. Using the zoom feature might slow the NNMi console performance when displaying maps.

The remedy is to limit the number of displayed nodes, displayed end points, or both, by doing the following:

- 1 In the NNMi console, click **Configuration**.
- 2 Click **User Interface Configuration** located beneath **User Interface**.
- 3 Select the **Default Map Settings** tab.
- 4 Modify the value shown in the **Maximum Number of Displayed Nodes** field.
- 5 Modify the value shown in the **Maximum Number of Displayed End Points** field.
- 6 Click **Save and Close**.

See *Define Default Map Settings* in the NNMi help for more information.

Configuring Gauges in the Analysis Pane

The Gauges tab in the analysis pane shows real-time SNMP gauges that display State Poller and Custom Poller SNMP data. These gauges display data for nodes, interfaces, custom node collections, and for node components of type CPU, Memory, Buffers, or Backplane.

You can configure the gauges by editing the following properties file:

- *Windows*: %NNM_PROPS%\nms-ui.properties
- *UNIX*: \$NNM_PROPS/nms-ui.properties

For each property that you want to set, if present, be sure to remove the comment characters (!) located at the beginning of the line.



The properties discussed in the sections that follow apply to ALL nodes (in other words, it is not possible to apply the properties to separate Node Groups).



Make a backup copy of the `nms-ui.properties` file before making any changes. Be sure to place the backup copy in a directory other than the directory containing the properties file you are editing.

See also the comments within the `nms-ui.properties` file for more information.

Limiting the Number of Gauges Displayed

Set the maximum number of gauges to be displayed by editing the following line and providing the desired value:

```
com.hp.nnm.ui.maxGaugePerAnalysisPanel =
```



A higher number of gauges affects performance when the analysis pane is displayed. A fewer number of gauges results in larger size gauges.

Setting the Refresh Rate for Gauges in the Analysis Pane

Set the refresh interval (in seconds) for gauges displayed in the analysis pane by editing the following property value:

```
com.hp.nnm.ui.analysisGaugeRefreshSecs =
```



Setting the value to “0” results in gauges never refreshing. A refresh rate faster than 10 seconds causes some SNMP agents to cache their values for short periods of time, causing repeated results.

Eliminating Gauges from the Display

Define the gauges that you do NOT want displayed (for all gauge views) by editing the following line and providing a list of gauges to eliminate from the display:

```
com.hp.nnm.ui.analysisGaugeNoDisplayKeyPatterns =
```

Be sure to uncomment all related lines; you cannot have comments within a list of gauges. Also, ensure that no blank lines exist within the list of gauges, as a blank line terminates the entries at the location of the blank line.

The default settings for this property are those in the comments. These settings must be included if this configuration is being extended or amended; otherwise, an unexpected amount of gauges will appear.

Controlling the Order of Displayed Node Gauges

You can control the order in which node gauges are displayed by editing the following line:

```
com.hp.nnm.ui.analysisGaugeNodeComponentKeys =
```

Wildcards are not supported in this property setting. Ensure that the list does not contain comments or empty lines.

The default settings for this property are those in the comments. These settings must be included if this configuration is being extended or amended; otherwise, the order will not match what was anticipated.

Controlling the Order of Displayed Interface Gauges

You can control the order in which interface gauges are displayed by editing the following line:

```
com.hp.nnm.ui.analysisGaugeInterfaceKeys =
```

Wildcards are not supported in this property setting. Ensure that the list does not contain comments or empty lines.

The default settings for this property are those in the comments. These settings must be included if this configuration is being extended or amended; otherwise, the order will not match what was anticipated.

Controlling the Order of Displayed Custom Poller Gauges

You can control the order in which Custom Poller gauges are displayed by editing the following line:

```
com.hp.ov.nnm.ui.analysisGaugeCustomPolledInstanceKeys =
```

There is no default setting for this attribute.

Understanding how Gauge Properties are Applied

Gauge properties are applied in the following order:

- 1 The list of all possible gauges is retrieved from State Poller.
- 2 The `analysisGaugeNoDisplayKeyPatterns` is first applied to remove the specified gauges from the list.
- 3 The `analysisGaugeNodeComponentKeys`, `analysisGaugeInterfaceKeys`, or `analysisGaugeCustomPolledInstanceKeys` is applied, as appropriate, to order the list of displayed gauges.
- 4 Finally the `maxGaugePerAnalysisPanel` is applied to truncate the displayed list.

Determining the Names of Gauges

One needs to know the names of gauges to include, suppress, order, or troubleshoot them. Determine gauge names as follows:

- 1 Bring up the jmx-console: ***http://<nnmiHost>/jmx-console***
- 2 Search the page for either of the following:
 - `com.hp.ov.nms.statePoller` (for Node and Interface gauges)
 - `com.hp.ov.nms.customPoller` (for Custom Poller gauges)
- 3 Click the **Collector** mbean.
- 4 Search for the function `dumpCollectionsMatchingTopologyObjectAndPolicy` and click **Invoke** below it without entering any parameter values. This creates a file in the `tmp` directory on the NNMi system.
- 5 Open this file and then search for the node in question and look for the collection information associated with it. For example:

```
columnsToCollect:
```

```
    Type: SNMPInstrumentationVariable, Name: sysUpTime, Value:
.1.3.6.1.2.1.1.3
```

```
    Type: SNMPInstrumentationVariable, Name: cpu5s, Value:
.1.3.6.1.4.1.9.9.109.1.1.1.1.3
```

```
Type: SNMPInstrumentationVariable, Name: cpu1m, Value:  
.1.3.6.1.4.1.9.9.109.1.1.1.1.4
```

```
Type: SNMPInstrumentationVariable, Name: cpu5m, Value:  
.1.3.6.1.4.1.9.9.109.1.1.1.1.5
```

The names of the collections, and, therefore, the gauges, can be seen in the listing.

Troubleshooting Gauge Problems

Too Many Gauges

Either limit the number of gauges displayed using the `maxGaugePerAnalysisPanel` property, or use the `analysisGaugeNoDisplayKeyPatterns` property to remove the gauges that are not wanted.

Missing Gauges

Use the `jmx-console` to dump the State Poller collections and verify which collections are being performed on the device and their names. It is possible that the device does not support a collection; for example, `cpu1m` is not available for certain devices.

Customizing Device Profile Icons

NNMi enables you to customize icons associated with a Device Profile or specific Nodes. These icons appear in table views, menu items, and as foreground images on an NNMi topology map. You can customize one or many icons using the `nmmicons.ovpl` command. For more information, see the *nmmicons.ovpl* reference page, or the UNIX manpage. See also the NNMi help.

Advanced Configuration

This section contains the following chapters:

- [Licensing NNMi](#)
- [Working with Certificates for NNMi](#)
- [Using Single Sign-On with NNMi](#)
- [Configuring the Telnet and SSH Protocols for Use by NNMi](#)
- [Integrating NNMi with a Directory Service through LDAP](#)
- [NNMi Security and Multi-Tenancy](#)
- [Global Network Management](#)
- [Configuring NNMi Advanced for IPv6](#)
- [Running NNMi in a Solaris Zones Environment](#)
- [NNMi Northbound Interface](#)

Licensing NNMi

If you do not have a permanent license key installed, the NNMi product includes a temporary Instant-On license key that is valid for 60 days after you install NNMi. This temporary Instant-On license key enables you to use NNMi Advanced features. You should obtain and install a permanent license key as soon as possible.

To view a list of the features included with an NNMi Advanced license, see the licensing section of the *HP NNMi Software Release Notes*.

Preparing to Install a Permanent License Key

The temporary Instant-On license has a 250 node limit. If you have been running NNMi with the Instant-On license key, you might be managing more nodes than your permanent license supports. When the permanent license takes effect, NNMi automatically unmanages nodes of its choosing to achieve the license limit.

If you want to control which nodes are no longer managed with the permanent license, use the NNMi console to delete less important nodes before installing your new license key.

Checking the License Type and the Number of Managed Nodes

To determine the type of license that NNMi is using, follow these steps:

- 1 In the NNMi console, click **Help > About Network Node Manager**.
- 2 In the **About Network Node Manager** window, click **View Licensing Information**.
(**View Licensing Information** is also available on the NNMi console sign-in page.)
- 3 Look for the value shown in the **Consumption** field. This is the number of nodes that NNMi is currently managing.
- 4 If your permanent license supports fewer nodes than NNMi is currently managing, use the NNMi console to delete less important nodes. For more information, see *Delete a Node* in the NNMi help.

Obtaining and Installing a Permanent License Key

To request a permanent license key, gather the following information:

- The Entitlement Certificate, which contains the HP product number and order number
- The IP address of one of the NNMi management servers
- If the license is for NNMi running under HA, the virtual IP address of the NNMi HA resource group
- Your company or organization information

Using Autopass and your HP Order Number (not possible behind a firewall)

To obtain and install a permanent license key, follow these steps:

- 1 At a command prompt, enter the following command to open the Autopass user interface:

```
nnmlicense.ovpl NNM -gui
```
- 2 On the left side of the Autopass window, click **License Management**.
- 3 Click **Install License Key**.
- 4 Click **Retrieve/Install License Key**.
- 5 Enter your HP Order Number and follow the Autopass prompts to complete the license key retrieval process.
- 6 NNMi automatically completes the installation.

From the Command Line

If the automated process does not run to completion (for example, if the NNMi management server is behind a firewall), follow these steps:

- 1 To obtain a license key, go to the HP password delivery service at

```
https://webware.hp.com/welcome.asp
```
- 2 At a command prompt on the NNMi management server, enter the following command to update the system and to store license data files:

```
nnmlicense.ovpl NNM -f license_file
```


(The product license ID (NNM) is case-sensitive.)
See the *nnmlicense.ovpl* reference page, or the UNIX manpage, for more information.
- 3 NNMi automatically completes the installation.

Obtaining Additional License Keys

Contact your HP Sales Representative or your Authorized Hewlett-Packard Reseller for information about the NNMi licensing structure, and to learn how to add license tiers for enterprise installations.

To obtain additional license keys, go to the HP License Key Delivery Service:

```
https://webware.hp.com/welcome.asp
```

See *Extend a Licensed Capacity* in the NNMi help for more information.

Note to Developers: With the NNMi Developer Toolkit, you can enhance the capabilities of NNMi by integrating custom web-service clients. After you install an NNMi Developer license, NNMi creates the `sdk-dev-kit.jar` file located in the `doc` folder. Unpack the `sdk-dev-kit.jar` file to view the NNMi Developer Toolkit documentation and samples.

Working with Certificates for NNMi

A certificate identifies the web server to the browser. This certificate can be self-signed or signed by a CA (Certificate Authority). The `nmn.keystore` file stores private keys and certificates with their corresponding public keys. The `nmn.truststore` file contains certificates from other parties that you expect to communicate with, or from Certificate Authorities that you trust to identify other parties. NNMi includes a self-signed certificate in both of the `nmn.keystore` and `nmn.truststore` files..

To use certain NNMi features, NNMi management servers need to share their certificates with one another. This chapter contains configuration instructions for copying these certificates among NNMi management servers and using the `nmncertmerge.ovpl` script to merge these certificates into the `nmn.keystore` and `nmn.truststore` files.

An administrator can disable HTTP and other unencrypted access from the network to NNMi. See [Configuring NNMi to Require Encryption for Remote Access](#) on page 401.

This chapter contains the following topics:

- [Putting it All Together](#)
- [Generating a Certificate Authority Certificate](#)
- [Configuring Application Failover to use Self-Signed Certificates](#)
- [Configuring Application Failover to use a Certificate Authority](#)
- [Configuring High Availability to use Self-Signed or Certificate Authority Certificates](#)
- [Configuring the Global Network Management Feature to use Self-Signed Certificates](#)
- [Configuring the Global Network Management Feature to use a Certificate Authority](#)
- [Configuring Global Network Management with Application Failover to use Self-Signed Certificates](#)
- [Configuring an SSL Connection to the Directory Service](#)

Putting it All Together

Use the following information to guide you in configuring certificates for your special needs:

- If you are using CA certificates, follow the instructions shown in [Generating a Certificate Authority Certificate](#) on page 127.
- If you configured your global, regional, or both NNMi management servers to use the application failover feature there are some additional configuration steps. Merge the NNMi management servers' `nmm.keystore` and `nmm.truststore` files for each cluster before completing the global network management configuration, as described in the [Configuring Application Failover to use Self-Signed Certificates](#) on page 130.
- If you must use a Certificate Authority, and you configured your global, regional, or both NNMi management servers to use the application failover feature, there are some additional configuration steps. First, follow the instructions shown in [Generating a Certificate Authority Certificate](#) on page 127; then merge the NNMi management servers' `nmm.keystore` and `nmm.truststore` files for each cluster before completing the global network management configuration, as described in the [Configuring Application Failover to use a Certificate Authority](#) on page 132.
- If you configured your global, regional, or both NNMi management servers to use High Availability, create the self-signed certificate in the `nmm.keystore` and `nmm.truststore` files for the virtual hostname before completing the global network management configuration, as described in [Configuring High Availability to use Self-Signed or Certificate Authority Certificates](#) on page 134.
- After you have each HA or application failover cluster properly configured, enable the global network management feature by copying the `nmm.truststore` file from the active regional node to the active global node, then merging the truststore. You must do this for each active regional node. Review the information shown in [Configuring Global Network Management with Application Failover to use Self-Signed Certificates](#) on page 137. If the NNMi management servers use CA certificates generated using the procedure shown in [Generating a Certificate Authority Certificate](#) on page 127, then those CA certificates are the only certificates you must merge into the global truststore.
- If you configure your NNMi management servers in a global network management configuration, then decide later to change the regional, global, or both to be in an application failover cluster, follow the instructions shown in [Configuring Application Failover to use Self-Signed Certificates](#) on page 130. You must use the commands shown in that section to configure your `nmm.keystore` and `nmm.truststore` files correctly; then copy the modified `nmm.truststore` file to the global NNMi management server and merge it into its `nmm.truststore` file.
- If you configure your NNMi management servers in a global network management configuration, then decide later to change the regional, global, or both to use HA, follow the instructions shown in [Configuring High Availability to use Self-Signed or Certificate Authority Certificates](#) on page 134.
- After directory service communications are enabled, NNMi uses the LDAP protocol for retrieving data from a directory service. If the directory service requires an SSL connection, follow the instructions show in [Configuring an SSL Connection to the Directory Service](#) on page 138.

Generating a Certificate Authority Certificate

If you plan to use a CA (Certificate Authority), complete the following steps to generate a CA certificate.



If you plan to use a CA with NNMI, sign the certificate using the RSA algorithm. The DSA algorithm is not supported.

1 Change to the directory on the NNMI management server that contains the `nnm.keystore` and `nnm.truststore` files:

- **Windows:** `%NNM_DATA%\shared\nnm\certificates`
- **UNIX:** `$NNM_DATA/shared/nnm/certificates`

2 Save a backup copy of the `nnm.keystore` file.

3 Generate a private key from your system. Use the `keytool` command to generate this private key:

a Run the following command *exactly as shown*:

- **Windows:** `%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -genkeypair -validity 3650 -keyalg rsa -keystore nnm.keystore -storepass nnmkeypass -alias myserver.mydomain`
- **UNIX:** `$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -genkeypair -validity 3650 -keyalg rsa -keystore nnm.keystore -storepass nnmkeypass -alias myserver.mydomain`



The alias, referred to as `myserver.mydomain` in this example, identifies this newly-created key. Although the alias can be any string, HP recommends you use the fully-qualified domain name of your system for the `myserver.mydomain` alias variable.



Linux operating systems have a `keytool` command that is not compatible with the `keytool` command or command options used in this step.

b Enter the requested information.



Important: When prompted for your first and last name, enter the FQDN (fully-qualified domain name) of your system.

4 Run the following command *exactly as shown* to create a CSR (Certificate Signing Request) file:

- **Windows:** `%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -keystore nnm.keystore -certreq -storepass nnmkeypass -alias myserver.mydomain -file CERTREQFILE`
- **UNIX:** `$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -keystore nnm.keystore -certreq -storepass nnmkeypass -alias myserver.mydomain -file CERTREQFILE`



For more information about the `keytool` command, search for “Key and Certificate Management Tool” at <http://www.oracle.com/technetwork/java/index.html>.

5 Send the CSR to your CA signing authority. They should provide you with one of the following:

- A signed certificate, referred to as `myserver.crt`. The `myserver.crt` file contains both the server certificate (the top certificate contained in the file) and one or more CA (Certified Authority) certificates (the last certificates contained in the file). Copy the CA certificate into a new file, the `myca.crt` file. Use the `myserver.crt` file when importing the server certificate into the `nnm.keystore` file and the `myca.crt` file when importing the CA certificate into the `nnm.truststore` file.
- Two files, referred to `myserver.crt` and `CA.crt` in this procedure. Add the `CA.crt` file content to the end of the `myserver.crt` file. Use the `myserver.crt` file when importing the server certificate into the `nnm.keystore` file and the `myca.crt` file when importing the CA certificate into the `nnm.truststore` file.

The following examples show you what the files you receive from your CA signing authority might look like:

Separate server and CA certificate files:

```
-----BEGIN CERTIFICATE-----
Sample/AVQKKExnQU0EgQ29ycG9yYXRpb24gTHRkMRAwDgYDVQQLEwd0ZXR3b3Js
eGVVSXZvY2F0aW9uTG1zdD9iYXNlP29iamVjdENsYXNzPWNSRERpc3RyaWJ1dG1w
.....
TZImiZPyLQQBGRYDaW50MRIwEAYKZCImiZPyLQQBGRYCC2cxEzARBgNVBAMTCmNb
pSo6o/76yShtT7Vr1fz+mXjWyEHaIy/QLCpPebYhejHEg4dZgzWWT/1Qt==
-----END CERTIFICATE-----
```

Combined server and CA certificates in one file:

```
-----BEGIN CERTIFICATE-----
Sample1/VQKKExnQU0EgQ29ycG9yYXRpb24gTHRkMRAwDgYDVQQLEwd0ZXR3b3Js
eGVVSXZvY2F0aW9uTG1zdD9iYXNlP29iamVjdENsYXNzPWNSRERpc3RyaWJ1dG1w
.....
TZImiZPyLQQBGRYDaW50MRIwEAYKZCImiZPyLQQBGRYCC2cxEzARBgNVBAMTCmNb
pSo6o/76yShtT7Vr1fz+mXjWyEHaIy/QLCpPebYhejHEg4dZgzWWT/1Qt==
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
Sample2/Gh0dHA6Ly9jb3JwMWRjc2cyLnNnLmludC5wc2Fnbg9iYWwuY29tL0Nlc
RaOCapwwggKYMBOGA1UdDgQWBBSqaWZzCRcpvJWOFpZ/Be9b+QSPyDAfBgNVHSMC
.....
Wp5Lz1ZJAou1VHbPVdQnXn1Bkx7V65niLoat90Eqd6laliVlJHj7GBriJ90uvVGu
BQagggEChoG9bGRhcDovLy9DTj1jb3JwMWRjc2cyL==
-----END CERTIFICATE-----
```

6 Copy the files containing these certificates to a location on the NNMi management server. For this example, copy the files to the following location:

- **Windows:** `%NNM_DATA%\shared\nnm\certificates`
- **UNIX:** `$NNM_DATA/shared/nnm/certificates`

Use the certificates you generated in the previous steps to replace the self-signed certificate:

- 1 Change to the directory on the NNMi management server that contains the `nnm.keystore` and `nnm.truststore` files:
 - *Windows*: `%NNM_DATA%\shared\nnm\certificates`
 - *UNIX*: `$NNM_DATA/shared/nnm/certificates`
- 2 Run the following command to import the server certificate and the CA certificate into the NNMi `nnm.keystore` file:

Windows:

- `%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -importcert -trustcacerts -keystore nnm.keystore -storepass nnmkeypass -alias myserver.mydomain -file myserver.crt`

UNIX:

- `$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -importcert -trustcacerts -keystore nnm.keystore -storepass nnmkeypass -alias myserver.mydomain -file myserver.crt`



If you use the `-storepass` option and provide the password, the keystore program does not prompt you for the keystore password. If you do not use the `-storepass` option, enter `nnmkeypass` when prompted for the keystore password.

- 3 When prompted to trust the certificate, enter: **y**

The output from this command is of the form:

```
Owner: CN=NNMi_server.example.com
Issuer: CN=NNMi_server.example.com
Serial number: 494440748e5
Valid from: Tue Oct 28 10:16:21 MST 2008 until: Thu Oct 04
11:16:21 MDT 2108
Certificate fingerprints:
MD5: 29:02:D7:D7:D7:D7:29:02:29:02:29:02:29:02:29:02
SHA1: C4:03:7E:C4:03:7E:C4:03:7E:C4:03:7E:C4:03:7E:C4:03
Trust this certificate? [no]: y
Certificate was added to keystore
```

- 4 Run the following commands to import the CA certificate into the NNMi `nnm.truststore` file:

— *Windows*:

```
%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -import -alias myca -keystore nnm.truststore -file myca.crt
```

— *UNIX*:

```
$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -import -alias myca -keystore nnm.truststore -file myca.crt
```

- 5 When prompted for the truststore password, enter: **ovpass**.

Example output for importing a certificate into the keystore

- 6 Examine the contents of the trust store:
 - *Windows:*
`%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool -list \-keystore nnm.truststore`
 - *UNIX:*
`$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -list \-keystore nnm.truststore`

When prompted for the truststore password, enter: **ovpass**

Example trust store output

The trust store output is of the form:

```
Keystore type: jks
Keystore provider: SUN
Your keystore contains 1 entry
nnmi_ldap, Nov 14, 2008, trustedCertEntry,
Certificate fingerprint (MD5):
29:02:D7:D7:D7:D7:29:02:29:02:29:02:29:02:29:02
```



The trust store can include multiple certificates.

- 7 Edit the following file:
 - *Windows:* %NNM_CONF%\nnm\props\nms-local.properties
 - *UNIX:* \$NNM_CONF/nnm/props/nms-local.properties
- 8 Update the `com.hp.ov.nms.ssl.KEY_ALIAS` variable to the value you used for `myserver.mydomain`. Make sure to save your work.
- 9 Restart NNMi by running these commands:
 - a **ovstop**
 - b **ovstart**
- 10 Test HTTPS access to the NNMi console using the following syntax:
https://<fully_qualified_domain_name>:<port_number>/nnm/. If the browser trusts the CA, it will trust the HTTPS connection to the NNMi console.

Configuring Application Failover to use Self-Signed Certificates

Figure 3 Using Self-Signed Certificates with Application Failover



When configuring the application failover feature, you must merge the `nmn.keystore` and `nmn.truststore` file content for both nodes into a single `nmn.keystore` and `nmn.truststore` file. Complete the following steps to configure the application failover feature to use self-signed certificates based on the above diagram.



If you are using self-signed certificates with NNMI along with the application failover feature, and do not complete the following steps, NNMI processes will not start correctly on the standby NNMI management server (Server Y in this example).

- 1 Change to the following directory on Server Y before completing [step 2](#):
 - *Windows*: %NNM_DATA%\shared\nnm\certificates
 - *UNIX*: \$NNM_DATA/shared/nnm/certificates
- 2 Copy the `nmn.keystore` and `nmn.truststore` files from Server Y to some temporary location on Server X. The remaining steps refer to these file locations as `<keystore>` and `<truststore>`.
- 3 Run the following command on Server X to merge Server Y's certificates into Server X's `nmn.keystore` and `nmn.truststore` files.

Windows:

```
nmncertmerge.ovpl -keystore <keystore> -truststore  
<truststore>
```

UNIX:

```
nmncertmerge.ovpl -keystore <keystore> -truststore  
<truststore>
```

- 4 Copy the merged `nmn.keystore` and `nmn.truststore` files from server X to server Y, so that both nodes have the merged files. The location of these files is as follows:
 - *Windows*: %NNM_DATA%\shared\nnm\certificates
 - *UNIX*: \$NNM_DATA/shared/nnm/certificates
- 5 Run the following command on both Server X and Server Y. Verify that the displayed results from both servers, including the fully-qualified-domain names, match. If they do not match do not continue, rather redo [step 1](#) through [step 7](#).

Windows:

```
%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -list  
-keystore %NnmDataDir%\shared\nnm\certificates\nnm.keystore  
-storepass nnmkeypass
```

UNIX:

```
$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -list -keystore  
$NnmDataDir/shared/nnm/certificates/nnm.keystore -storepass  
nnmkeypass
```

- 6 Run the following command on both Server X and Server Y. Verify that the displayed results from both servers, including the fully-qualified-domain names, match. If they do not match do not continue, rather redo [step 1](#) through [step 7](#).

Windows:

```
%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -list  
-keystore  
%NnmDataDir%\shared\nnm\certificates\nnm.truststore  
-storepass ovpass
```

UNIX:

```
$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -list -keystore
$NnmDataDir/shared/nnm/certificates/nnm.truststore
-storepass ovpass
```

- 7 Continue configuring the application failover feature at [step 4](#) on page 422.



Although you manually completed the following automatic action during [step 4](#), after you start the application failover feature, NNMi automatically replicates the merged keystore and truststore information from NNMi_active to NNM_standby.

Configuring Application Failover to use a Certificate Authority

Figure 4 Using CA Certificates with Application Failover



When configuring the application failover feature, you must merge the `nmn.keystore` and `nmn.truststore` file content for both nodes into a single `nmn.keystore` and `nmn.truststore` file. Complete the following steps to configure the application failover feature to use CA certificates based on the above diagram.



If you are using CA certificates with NNMi along with the application failover feature, and do not complete the following steps, NNMi processes will not start correctly on the standby NNMi management server (Server Y in this example).

- 1 Follow the instructions shown in [Generating a Certificate Authority Certificate](#) on page 127 for NNMi_standby.
- 2 Change to the following directory on Server Y before completing [step 3](#):
 - *Windows*: %NNM_DATA%\shared\nnm\certificates
 - *UNIX*: \$NNM_DATA/shared/nnm/certificates
- 3 Copy the `nmn.keystore` and `nmn.truststore` files from Server Y to some temporary location on Server X. The remaining steps refer to these file locations as `<keystore>` and `<truststore>`.
- 4 Run the following command on Server X to merge Server Y's certificates into Server X's `nmn.keystore` and `nmn.truststore` files.

Windows:

```
nmncertmerge.ovpl -keystore <keystore> -truststore
<truststore>
```

UNIX:

```
nmncertmerge.ovpl -keystore <keystore> -truststore
<truststore>
```

- 5 Copy the merged `nnm.keystore` and `nnm.truststore` files from server X to server Y, so that both nodes have the merged files. The location of these files is as follows:

- *Windows:* `%NNM_DATA%\shared\nnm\certificates`
- *UNIX:* `$NNM_DATA/shared/nnm/certificates`

- 6 Run the following command on both Server X and Server Y. Verify that the displayed results from both servers, including the `hp.com` fully-qualified-domain name, match. If they do not match do not continue, rather redo [step 1](#) through [step 7](#).

Windows:

```
%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -list
-keystore %NnmDataDir%\shared\nnm\certificates\nnm.keystore
-storepass nnmkeypass
```

UNIX:

```
$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -list -keystore
$NnmDataDir/shared/nnm/certificates/nnm.keystore -storepass
nnmkeypass
```

- 7 Run the following command on both Server X and Server Y. Verify that the displayed results from both servers, including the `hp.com` fully-qualified-domain name, match. If they do not match do not continue, rather redo [step 1](#) through [step 7](#).

Windows:

```
%NnmInstallDir%\nonOV\jdk\nnm\bin\keytool.exe -list
-keystore
%NnmDataDir%\shared\nnm\certificates\nnm.truststore
-storepass ovpass
```

UNIX:

```
$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -list -keystore
$NnmDataDir/shared/nnm/certificates/nnm.truststore
-storepass ovpass
```

- 8 Continue configuring the application failover feature at [step 4](#) on page 422.



Although you manually completed the following automatic action during [step 5](#) on page 133, after you start the application failover feature, NNMi automatically replicates the merged keystore and truststore information from Server X to Server Y

Configuring High Availability to use Self-Signed or Certificate Authority Certificates

Figure 5 Using Certificates with HA



Configuring High Availability to use Self-Signed Certificates

The process for configuring NNMi for HA correctly shares the self-signed certificate among the primary and secondary cluster nodes. You do not need to take any extra steps to use the default certificate with NNMi running under HA.

Configuring High Availability for a New Certificate

Suppose you create a new self-signed or CA certificate, referred to as `newcert`. Complete the following steps to configure HA with this new CA or self-signed certificate.



You can complete this procedure before or after configuring NNMi for HA, as described in [Shared NNMi Data](#) on page 331.

- 1 Change to the following directory on NNMi_HA1 before completing [step 2](#):
 - *Windows*: %NNM_DATA%\shared\nnm\certificates
 - *UNIX*: \$NNM_DATA/shared/nnm/certificates
- 2 On NNMi_HA1, run the following commands to import `newcert` into the `nnm.keystore` file:
 - *Windows*: %NnmInstallDir%\nonOV\jdk\nnm\bin\keytool -import -alias `newcert_Alias` -keystore `nnm.keystore` -file `newcert`
 - *UNIX*: \$NnmInstallDir/nonOV/jdk/nnm/bin/keytool -import -alias `newcert_Alias` -keystore `nnm.keystore` -file `newcert`
- 3 Edit the following file on both the active (NNMi_HA1) and the standby (NNMi_HA2) nodes:
 - *Windows*: %NNM_DATA%\conf\nnm\props\nms-local.properties
 - *UNIX*: \$NNM_DATA/conf/nnm/props/nms-local.properties
- 4 Change the following line in the `nms-local.properties` file on both NNMi_HA1 and NNMi_HA2.



```
com.hp.ov.nms.ssl.KEY_ALIAS = newcert_Alias
```
- 5 Save your changes.

Configuring the Global Network Management Feature to use Self-Signed Certificates

During NNMi installation, the installation script creates a self-signed certificate for the NNMi management server. This certificate contains an alias that includes the fully-qualified domain name of the node. The installation script adds this self-signed certificate to the NNMi management server's `nmm.keystore` and `nmm.truststore` files.

Suppose you want your global network management configuration to model [Figure 6](#).

Figure 6 Global Network Management



Complete the following steps to configure the global network management feature to use self-signed certificates based on [Figure 6](#).

- 1 Change to the following directory on `regional1` and `regional2` before completing [step 2](#):
 - *Windows*: `%NNM_DATA%\shared\nnm\certificates`
 - *UNIX*: `$NNM_DATA/shared/nnm/certificates`
- 2 Copy the `nmm.truststore` files from the above locations on `regional1` and `regional2` to some temporary location on `global1`.
- 3 Run the following command on `global1` to merge the `regional1` and `regional2` certificates into `global1`'s `nmm.truststore` file.

Windows:

- a `nnmcertmerge.ovpl -truststore regional1_nmm.truststore_location`
- b `nnmcertmerge.ovpl -truststore regional2_nmm.truststore_location`

UNIX

- a `nnmcertmerge.ovpl -truststore regional1_nmm.truststore_location`
- b `nnmcertmerge.ovpl -truststore regional2_nmm.truststore_location`


- 4 Run the following command sequence on `global1`:
 - a `ovstop`
 - b `ovstart`

Configuring the Global Network Management Feature to use a Certificate Authority

During NNMi installation, the installation script creates a self-signed certificate for the NNMi management server. This certificate contains an alias that includes the fully-qualified domain name of the node. The installation script adds this self-signed certificate to the NNMi management server's `nnm.keystore` and `nnm.truststore` files.

Suppose you want your global network management configuration to model [Figure 7](#).

Figure 7 Using Certificates with Global Network Management



- 1 Follow the instructions shown in [Generating a Certificate Authority Certificate](#) on page 127 for `regional1` and `regional2`.
- 2 Change to the following directory on `regional1` and `regional2` before completing [step 3](#).
 - *Windows*: `%NNM_DATA%\shared\nnm\certificates`
 - *UNIX*: `$NNM_DATA/shared/nnm/certificates`
- 3 Copy the `nnm.truststore` files from the above locations on `regional1` and `regional2` to some temporary location on `global1`.
- 4 Run the following command on `global1` to merge the `regional1` and `regional2` certificates into `global1`'s `nnm.truststore` file.

Windows:

- a `nnmcertmerge.ovpl -truststore regional1_nnm.truststore_location`
- b `nnmcertmerge.ovpl -truststore regional2_nnm.truststore_location`

UNIX

- a `nnmcertmerge.ovpl -truststore regional1_nnm.truststore_location`
- b `nnmcertmerge.ovpl -truststore regional2_nnm.truststore_location`


- 5 Run the following command sequence on `global1`:
 - a `ovstop`
 - b `ovstart`

Configuring Global Network Management with Application Failover to use Self-Signed Certificates

As discussed above, during NNMI installation the installation script creates a self-signed certificate for the NNMI management server. This certificate contains an alias that includes the fully-qualified domain name of the node. The installation script adds this self-signed certificate to the NNMI management server's `nmm.keystore` and `nmm.truststore` files.

Suppose you want your global network management configuration to model the application failover feature as shown in [Figure 8](#).

Figure 8 Global Network Management with Application Failover



Complete the following steps to configure the global network management feature to work with application failover based on the above diagram:

- 1 Follow the instructions shown in [Configuring Application Failover to use Self-Signed Certificates](#) on page 130 for each application failover cluster shown in the above diagram.
- 2 Complete the configuration for application failover shown in [Application Failover Basic Setup](#) on page 284.
- 3 Follow the instructions shown in [Configuring the Global Network Management Feature to use Self-Signed Certificates](#) on page 135 for `regional1_active` and `regional2_active`.

Configuring an SSL Connection to the Directory Service

By default, when directory service communications are enabled, NNMi uses the LDAP protocol for retrieving data from a directory service. If your directory service requires an SSL connection, you must enable the SSL protocol to encrypt the data that flows between NNMi and the directory service.

SSL requires a trust relationship between the directory service host and the NNMi management server. To create this trust relationship, add a certificate to the NNMi trust store. The certificate confirms the identity of the directory service host to the NNMi management server.

To install a trust store certificate for SSL communications, follow these steps:

- 1 Obtain your company's trust store certificate from the directory server. The directory service administrator should be able to give you a copy of this text file.
- 2 Change to the directory that contains the NNMi trust store:
 - *Windows*: %NNM_DATA%\shared\nnm\certificates
 - *UNIX*: \$NNM_DATA/shared/nnm/certificates

Run all commands in this procedure from the `certificates` directory.

- 3 Import your company's trust store certificate into the NNMi trust store:
 - a Run the following command:

```

— Windows:
%NnmInstallDir%\nonOV\jdk\b\bin\keytool -import
-alias nmi_ldap -keystore nnm.truststore
-file <Directory_Server_Certificate.txt>

— UNIX:
$NnmInstallDir/nonOV/jdk/b/bin/keytool -import \
-alias nmi_ldap -keystore nnm.truststore \
-file <Directory_Server_Certificate.txt>

```

Where `<Directory_Server_Certificate.txt>` is your company's trust store certificate.

- b When prompted for the keystore password, enter: **ovpass**
- c When prompted to trust the certificate, enter: **y**

The output from this command is of the form:

```

Owner: CN=NNMi_server.example.com
Issuer: CN=NNMi_server.example.com
Serial number: 494440748e5
Valid from: Tue Oct 28 10:16:21 MST 2008 until: Thu Oct 04
11:16:21 MDT 2108
Certificate fingerprints:
MD5: 29:02:D7:D7:D7:D7:29:02:29:02:29:02:29:02:29:02
SHA1: C4:03:7E:C4:03:7E:C4:03:7E:C4:03:7E:C4:03:7E:C4:03
Trust this certificate? [no]: y
Certificate was added to keystore

```

Example output for
importing a
certificate into the
trust store

- 4 Examine the contents of the trust store:
 - *Windows:*
`%NnmInstallDir%\nonOV\jdk\b\bin\keytool.exe -list -keystore nnm.truststore`
 - *UNIX:*
`$NnmInstallDir/nonOV/jdk/b/bin/keytool -list -keystore nnm.truststore`

When prompted for the keystore password, enter: **ovpass**

Example trust store
output

The trust store output is of the form:

```
Keystore type: jks
Keystore provider: SUN
Your keystore contains 1 entry
nnmi_ldap, Nov 14, 2008, trustedCertEntry,
Certificate fingerprint (MD5):
29:02:D7:D7:D7:D7:29:02:29:02:29:02:29:02:29:02
```



The trust store can include multiple certificates.

- 5 Restart NNMi by running these commands:
 - a **ovstop**
 - b **ovstart**

For more information about the `keytool` command, search for “Key and Certificate Management Tool” at <http://www.oracle.com/technetwork/java/index.html>.

Using Single Sign-On with NNMi

You can configure HP Network Node Manager i Software (NNMi) single sign-on (SSO) to facilitate access to NNM iSPIs from the NNMi console. With SSO, when you log on to the NNMi console, you receive access to NNM iSPIs and other HP applications without needing to log on again. SSO provides easier access to NNM iSPIs and other HP applications while maintaining a secure level of access. After you sign out of the NNMi console (or the NNMi console session times out), you must re-enter your sign-in credentials to access NNM iSPI and other HP application URLs outside the NNMi console.

SSO is not enabled during installation. If it was, browsing from one NNMi management server to another logs you out of the first one, providing little benefit. To keep this from happening, SSO is initially disabled so you can coordinate setting the `initString` and `protectedDomains` parameter among the NNMi management servers, as explained in this chapter.

This chapter contains the following topics:

- [SSO Access for NNMi](#) on page 142
- [Enabling SSO for a Single Domain](#) on page 143
- [Enabling SSO for NNMi Management Servers Located in Different Domains](#) on page 143
- [SSO Access for NNMi and the NNM iSPIs](#) on page 144
- [Disabling SSO](#) on page 146
- [SSO Security Notes](#) on page 146

SSO Access for NNMi

To browse among several NNMi management servers, you must do one of the following:

- Edit the `nms-ui.properties` file and make the parameter values for `com.hp.nms.ui.sso.initString` and `com.hp.nms.ui.sso.protectedDomains` the same among the NNMi management servers. Make sure to set the `com.hp.nms.ui.sso.domain` parameter to match the domain an NNMi management server resides in.
 - If you have NNMi management servers residing in only one network domain, follow the instructions show in [Enabling SSO for a Single Domain](#) on page 143.
 - If you have NNMi management servers residing in more than one network domain, follow the instructions shown in [Enabling SSO for NNMi Management Servers Located in Different Domains](#) on page 143 for more information.
- Edit the `nms-ui.properties` file and make sure you have SSO disabled. See [Disabling SSO](#) on page 146 for more information.

If you choose to not complete one of these actions, each time you browse to a different NNMi management server, you will be automatically signed out of the previous NNMi management server.

There are special considerations for using SSO with the NNMi global network management feature. See [SSO and the Actions Menu](#) on page 237 and [Configuring Single Sign-On for Global Network Management](#) on page 237 for more information.

If the domain name of the NNMi management server is short, as in `mycompany`, without any period (`.`), the NNMi console will immediately sign you out. The restrictions for SSO browser cookies require a domain name to contain at least one period, such as `mycompany.com`. To remedy this situation, complete the following steps:

- 1 Open the following file in a text editor:
 - *Windows*: `%NNM_PROPS%/nms-ui.properties`
 - *UNIX*: `$NNM_PROPS/nms-ui.properties`
- 2 For this example, search for the following string:


```
com.hp.nms.ui.sso.domain = mycompany
```

 and replace it with the following string:


```
com.hp.nms.ui.sso.domain = mycompany.com
```
- 3 Run the following command to commit the changes:


```
nnmssso.ovpl -reload
```

 See the `nnmssso.ovpl` reference page, or the UNIX manpage, for more information.

Enabling SSO for a Single Domain

To enable SSO for use in a single domain, complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-ui.properties
 - *UNIX*: \$NNM_PROPS/nms-ui.properties
- 2 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.isEnabled = false
```

Change this as follows:

```
com.hp.nms.ui.sso.isEnabled = true
```
- 3 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.domain = mycompany.com
```

Change *mycompany.com* to the domain the NNMi management server resides in. Make sure there is only one domain listed when enabling SSO in a single domain.
- 4 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.protectedDomains = mycompany.com
```

Change *mycompany.com* to the domain the NNMi management server resides in. Make sure there is only one protected domain listed when enabling SSO in a single protected domain.
- 5 Run the following command to commit the changes:

```
nnmssso.ovpl -reload
```

See the *nnmssso.ovpl* reference page, or the UNIX manpage, for more information.

Enabling SSO for NNMi Management Servers Located in Different Domains

You can configure two or more NNMi management servers for SSO. This example explains how to configure SSO for three NNMi management servers located in different domains. If you must configure two or more NNMi management servers for SSO and these systems reside in different domains, complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-ui.properties
 - *UNIX*: \$NNM_PROPS/nms-ui.properties
- 2 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.isEnabled = false
```

Change this as follows:

```
com.hp.nms.ui.sso.isEnabled = true
```
- 3 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.domain = group1.mycompany.com
```

Make sure the domain name contains at least one dot.

- 4 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.protectedDomains=group1.mycompany.com
```

Change this as follows:

```
com.hp.nms.ui.sso.protectedDomains=group1.mycompany.com,
group2.yourcompany.com, group3.yourcompany.com
```

- 5 Look for a section in the file that resembles the following:

```
com.hp.nms.ui.sso.initString = Initialization String
```

NNMi management servers must share the same initialization string to work in an SSO configuration. Change the initialization string the same value on all NNMi management servers included in the SSO configuration.

- 6 Run the following command to commit the changes:

```
nnmssso.ovpl -reload
```

See the *nnmssso.ovpl* reference page, or the UNIX manpage, for more information.

- 7 Repeat [step 1](#) through [step 6](#) two more times, configuring the remaining two NNMi management servers. For each remaining NNMi management server, substitute *group2* or *group3* for *group1* during [step 3](#).

SSO Access for NNMi and the NNM iSPIs

After SSO is enabled, SSO between NNMi and the NNM iSPIs does *not* require `initString` configuration.

To use SSO, access NNMi as follows:

- Use the correct URL in the following form:


```
<protocol>://<fully_qualified_domain_name>:<port_number>/nnm/
```

<protocol> represents either `http` or `https`.
<fully_qualified_domain_name> represents the official fully-qualified domain name (FQDN) of the NNMi management server.
<port_number> is the port for connecting to the NNMi console, is assigned during NNMi installation, and is specified in the following file:

 - *Windows*: %NnmDataDir%\conf\nnm\props\nms-local.properties
 - *UNIX*: \$NnmDataDir/conf/nnm/props/nms-local.properties
- Log on to NNMi using a valid account.

For SSO to work, URL access to NNMi and the NNM iSPIs must share a common network domain name. Additionally, the URL must not include an IP address. If you do not have a FQDN for the NNMi management server, you can substitute the IP address of the NNMi management server. However, doing so disables single sign-on for NNM iSPIs, and you must log on again the next time you access any NNM iSPI.

To determine the official FQDN of the NNMi management server, use one of the following methods:

- Use the **`nnmofficialfqdn.ovpl`** command to display the value of the official FQDN set during installation. See the *nnmofficialfqdn.ovpl* reference page, or the UNIX manpage, for more information.
- In the NNMi console, click **Help > System Information**. On the **Server** tab, look for the official FQDN statement.

If you must change the official FQDN set during installation, use the **`nnmsetofficialfqdn.ovpl`** command. See the *nnmsetofficialfqdn.ovpl* reference page, or the UNIX manpage, for more information.



After installation, the system account is still valid. Use the system account only for command-line security and for recovery purposes.

SSO to NNM iSPIs require that users access the NNMi console through a URL that contains the official FQDN. You can configure NNMi to redirect NNMi URLs to the official FQDN when the NNMi console is accessed through a non-official domain name, such as an IP address or a shortened version of the domain name. Before configuring NNMi to redirect URLs, an appropriate official FQDN must be configured. For information, see the NNMi help.

After you enable NNMi to redirect URLs, note the following:

- You can log on to the NNMi console using any hostname that is valid for the NNMi management server you want to access. For example, if you request `http://localhost/nnm`, NNMi redirects you to a URL such as `http://host.mydomain.com/nnm`.
- If you cannot access the NNMi console using `http://host.mydomain.com/nnm`, use the following to directly access the NNMi console:

`<protocol>://`

`<fully_qualified_domain_name>:<port_number>launch?cmd=showMain.`

`<protocol>` represents either `http` or `https`.

`<fully_qualified_domain_name>` represents the official fully-qualified domain name (FQDN) of the NNMi management server.

`<port_number>` is the port for connecting to the NNMi console, is assigned during NNMi installation, and is specified in the following file:

- *Windows:* `%NnmDataDir%\conf\nnm\props\nms-local.properties`
- *UNIX:* `$NnmDataDir/conf/nnm/props/nms-local.properties`

Disabling SSO

If you have a need to disable SSO, complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-ui.properties
 - *UNIX*: \$NNM_PROPS/nms-ui.properties
- 2 Look for a section in the file that resembles the following:


```
com.hp.nms.ui.sso.isEnabled = true
```

 Change this as follows:


```
com.hp.nms.ui.sso.isEnabled = false
```
- 3 Run the following command to commit the changes:


```
nnmssso.ovpl -reload
```

 See the *nnmssso.ovpl* reference page, or the UNIX manpage, for more information.

SSO Security Notes

- 1 Confidential `initString` parameter in SSO security.

SSO uses *Symmetric Encryption* to validate and create an SSO token. The `initString` parameter within the configuration is used for initialization of the secret key. An application creates a token, and each application that uses the same `initString` parameter validates the token.

The following information is very important:

 - It is not possible to use SSO without setting the `initString` parameter.
 - The `initString` parameter is confidential information and should be treated as such in terms of publishing, transporting, and persistency.
 - Applications that integrate with each other can share the `initString` using SSO.
 - The minimum length of the `initString` is 12 characters.
- 2 Disable SSO unless it is specifically required.
- 3 The application that uses the weakest authentication framework, and issues an SSO token that is trusted by other integrated applications, determines the level of authentication security for all the applications.

HP recommends that only applications using strong and secure authentication frameworks issue an SSO token.
- 4 Symmetric encryption implication:

SSO uses symmetric cryptography for issuing and validating SSO tokens. Therefore, any application using SSO can issue a token to be trusted by all other applications sharing the same `initString`.

This potential risk is relevant when an application sharing the `initString` either resides or is accessible in an untrusted location.

5 User roles:

SSO does not share user roles between integrated applications. Therefore, the integrated application must monitor user roles. HP recommends you share the same user registry (as LDAP/AD) among all integrated applications.

Failure to manage user roles might cause security breaches and negative application behavior. For example, the same user name might be assigned to different roles in the integrated applications.

There could be situations when a user logs on to application A, then accesses application B that uses container or application authentication. The failure to manage the user role will force the user to manually log on to application B and enter a username. If the user enters a different user name than the one used to log on to application A, the following unexpected behavior can arise: If the user subsequently accesses a third application, application C, from application A or application B, then the user will access it using the user names that were used to log on to application A or application B respectively.

6 Identity Manager is used for an authentication:

All unprotected resources in the Identity Manager must be configured as nonsecure URL settings in the SSO configuration.

7 SSO demonstration mode:

- Use the SSO demonstration mode for demonstrative purposes only.
- Only use the demonstration mode in unsecured networks.
- Do not use the demonstration mode in production. Any combination of the demonstration mode with the production mode should not be used.

Configuring the Telnet and SSH Protocols for Use by NNMi

The **Actions > Telnet... (from client)** menu item invokes the telnet command to the selected node (from the web browser in which the NNMi console is currently running). The **Actions > Secure Shell... (from client)** menu item invokes the secure shell (SSH) command to the selected node (from the web browser in which the NNMi console is currently running). By default, neither Microsoft Internet Explorer nor Mozilla Firefox defines the telnet command nor the SSH command, so using either of these menu items produces an error message. You can configure the telnet, SSH, or both protocols for each NNMi user (on a per-system basis), and you can change the NNMi console menu items.


This chapter contains the following topics:

- [Disable the Telnet or SSH Menu Item](#) on page 149
- [Configure a Telnet or SSH Client for the Browser on Windows](#) on page 150
- [Configure Firefox to use Telnet or SSH on Linux](#) on page 157
- [Example Files for Changing the Windows Registry](#) on page 158

Disable the Telnet or SSH Menu Item

If the NNMi users in your deployment environment do not require telnet or SSH connections from the NNMi console, you can disable the respective menu item to remove it from the NNMi console.

Disabling a menu item in the NNMi console applies to all users who log on to the NNMi console on this NNMi management server. To disable the **Telnet** or **Secure Shell** menu item, follow these steps:

- 1 In the **Configuration** workspace, expand **User Interface**, and then elect **Menu Items**.
- 2 In the **Menu Items** view, select the **Telnet... (from Client)** row or the **Secure Shell... (from client)** row, and then click **Open** .

- 3 On the **Menu Item** form, clear the **Enabled** check box, and then set the **Author** field to an appropriate value.

Changing the author value ensures that this menu item remains disabled when you upgrade NNMi.

- 4 Save and close the form.

For more information, see *Control the Actions Menu* in the NNMi help.

Configure a Telnet or SSH Client for the Browser on Windows

Configure the operating-system provided telnet command for an NNMi user's web browser. This procedure must be done for each computer and web browser from which an NNMi user needs to run the **Actions > Telnet... (from Client)** menu item.

Configure a third-party ssh command for an NNMi user's web browser. This procedure must be done for each computer and web browser from which an NNMi user needs to run the **Actions > Secure Shell... (from Client)** menu item.

To complete any of the procedures in this section, you must have administrative privileges on the computer. The specific steps depend on the version (32-bit or 64-bit) of the browser and the operating system.

To determine the version of Internet Explorer, click **Help > About Internet Explorer**. If the version information does not include the text **64-bit Edition**, this Internet Explorer is 32-bit.

Firefox is only available in a 32-bit version.

[Table 7](#) identifies the procedure to use for each browser and operating system combination.

Table 7 Matrix of Telnet and SSH Configuration Procedures on Windows

Web Browser	Windows Operating System Architecture	Applicable Procedures
Internet Explorer 32-bit	32-bit	<ul style="list-style-type: none"> Windows Operating System-Provided Telnet Client on page 152 Third-Party Telnet Client (Standard Windows) on page 153 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155
	64-bit Windows 7	<ul style="list-style-type: none"> Third-Party Telnet Client (Standard Windows) on page 153 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155
	64-bit other than Windows 7	<ul style="list-style-type: none"> Third-Party Telnet Client (Windows on Windows) on page 154 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155

Table 7 Matrix of Telnet and SSH Configuration Procedures on Windows (cont'd)

Web Browser	Windows Operating System Architecture	Applicable Procedures
Internet Explorer 64-bit	64-bit	<ul style="list-style-type: none"> Windows Operating System-Provided Telnet Client on page 152 Third-Party Telnet Client (Standard Windows) on page 153 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155
Firefox	32-bit	<ul style="list-style-type: none"> Windows Operating System-Provided Telnet Client on page 152 Third-Party Telnet Client (Standard Windows) on page 153 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155
	64-bit Windows 7	<ul style="list-style-type: none"> Third-Party Telnet Client (Standard Windows) on page 153 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155
	64-bit other than Windows 7	<ul style="list-style-type: none"> Third-Party Telnet Client (Windows on Windows) on page 154 Third-Party SSH Client (Standard Windows and Windows on Windows) on page 155



Many of the tasks in this section involve editing the Windows registry. Instead of editing the registry directly, you can create a .reg file that each user can run on their system. For example .reg files, see [Example Files for Changing the Windows Registry](#) on page 158.

For more information about the tasks described in this section, see the following Microsoft articles:

- Installing the Microsoft-provided telnet client:
<http://technet.microsoft.com/en-us/library/cc771275%28WS.10%29.aspx>
- Introduction to the Windows registry:
<http://support.microsoft.com/kb/256986>
- Backing up and restoring the Windows registry:
<http://support.microsoft.com/kb/322756>

Windows Operating System-Provided Telnet Client

This procedure applies to the following cases:

- 32-bit Internet Explorer on a 32-bit operating system
- 32-bit Firefox on a 32-bit operating system
- 64-bit Internet Explorer on a 64-bit operating system



The telnet client provided with the Windows operating system does not work with a 32-bit version of Internet Explorer running on 64-bit Windows operating system. To remedy this, use a 64-bit version of Internet Explorer. Windows 64-bit operating systems include both the 32-bit and 64-bit versions of Internet Explorer. Look for these Internet Explorer versions in the following directories:

- *64-bit Version:* %ProgramFiles%/Internet Explorer
- *32-bit Version:* %ProgramFiles(x86)%/Internet Explorer

To configure the operating system-provided telnet client for use by a web browser, follow these steps:

- 1 (Microsoft Windows 7, Microsoft Vista, or Microsoft Windows Server 2008 only) Install the operating system telnet client on the computer by following the steps appropriate to the operating system.

Windows 7 or Vista:

- a In the Control Panel, click **Programs**, and then click **Programs and Features**.
- b Under Tasks, click **Turn Windows features on or off**.
- c In the Windows Features dialog box, select the **Telnet Client** check box, and then click **OK**.

Windows Server 2008:

- a In the Server Manager, under Features Summary, click **Add Features**.
- b In the Add Features Wizard, select the **Telnet Client** check box, click **Next**, and then click **Install**.

- 2 (Internet Explorer only) Enable Internet Explorer to use the telnet protocol.

- a Back up the Windows registry.
- b Use the Windows registry editor to add the [HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\MAIN\FeatureControl\FEATURE_DISABLE_TELNET_PROTOCOL] key with the following values:

Name	Type	Data
iexplore.exe	REG_DWORD	0

- 3 Set file association for the URL:Telnet Protocol file type.

- a Back up the Windows registry.

- b Use the Windows registry editor to modify the [HKEY_CLASSES_ROOT\telnet\shell\open\command] key with the following value:

Name	Type	Data
(default)	REG_SZ	rundll32.exe url.dll,TelnetProtocolHandler %l

- 4 %l (with a lowercase L) is the argument passed to telnet, usually an IP address or the fully-qualified domain name of a node.



For tighter control, you can encode the paths to the binaries in the key (as a single line). For example:

```
"C:\Windows\system32\rundll32.exe"  
"C:\Windows\system32\url.dll",TelnetProtocolHandler %l
```

- 5 Restart the web browser, and then, in the browser address bar, enter the telnet command:

```
telnet://<node>
```

<node> is the IP address or fully-qualified domain name of a node that runs the telnet server.

If you are prompted with a security warning, permit the action.

In Firefox, select the **Remember my choice for telnet links** check box.

Third-Party Telnet Client (Standard Windows)

This procedure applies to the following cases:

- 32-bit Internet Explorer on a 32-bit operating system
- 32-bit Internet Explorer on a 64-bit Windows 7 operating system
- 32-bit Firefox on a 32-bit operating system
- 64-bit Internet Explorer on a 64-bit operating system

To configure a third-party telnet client for use by a web browser, follow these steps:

- 1 Obtain and install a third-party telnet client.

This procedure gives examples for the PuTTY client installed to C:\Program Files\PuTTY\putty.exe. The PuTTY client is available from <http://www.putty.org>.

- 2 (Internet Explorer only) Enable Internet Explorer to use the telnet protocol.

- a Back up the Windows registry.
- b Use the Windows registry editor to add the [HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\MAIN\FeatureControl\FEATURE_DISABLE_TELNET_PROTOCOL] key with the following values:

Name	Type	Data
iexplore.exe	REG_DWORD	0

- 3 Set file association for the URL:Telnet Protocol file type.

- a Back up the Windows registry.
- b Use the Windows registry editor to modify the [HKEY_CLASSES_ROOT\ telnet\shell\open\command] key with the following value:

Name	Type	Data
(default)	REG_SZ	"C:\Program Files\PuTTY\putty.exe" %1

%1 (with a lowercase L) is the argument passed to telnet, usually an IP address or the fully-qualified domain name of a node.



In a .reg file, escape each quotation mark (") and backslash (\) character with a backslash (\) character.

- 4 Restart the web browser, and then, in the browser address bar, enter the telnet command:

telnet://<node>

<node> is the IP address or fully-qualified domain name of a node that runs the telnet server.

If you are prompted with a security warning, permit the action.

In Firefox, select the **Remember my choice for telnet links** check box.

Third-Party Telnet Client (Windows on Windows)

This procedure applies to the following cases:

- 32-bit Internet Explorer on a 64-bit operating system (other than Windows 7)
- 32-bit Firefox on a 64-bit operating system

To configure a third-party telnet client for use by a web browser, follow these steps:

- 1 Obtain and install a third-party telnet client.

This procedure gives examples for the PuTTY client installed to C:\Program Files\PuTTY\putty.exe. The PuTTY client is available from **http://www.putty.org**.

- 2 (Internet Explorer only) Enable Internet Explorer to use the telnet protocol.

- a Back up the Windows registry.
- b Use the Windows registry editor to add the [HKEY_LOCAL_MACHINE\ SOFTWARE\ Wow6432Node\Microsoft\Internet Explorer\MAIN\ FeatureControl\FEATURE_DISABLE_TELNET_PROTOCOL] key with the following values:

Name	Type	Data
iexplore.exe	REG_DWORD	0

- 3 Set file association for the URL:Telnet Protocol file type.
 - a Back up the Windows registry.
 - b Use the Windows registry editor to modify the [HKEY_CLASSES_ROOT\Wow6432Node\telnet\shell\open\command] key with the following value:

Name	Type	Data
(default)	REG_SZ	"C:\Program Files\PuTTY\putty.exe" %l

%l (with a lowercase L) is the argument passed to telnet, usually an IP address or the fully-qualified domain name of a node.



In a .reg file, escape each quotation mark (") and backslash (\) character with a backslash (\) character.

- 4 Restart the web browser, and then, in the browser address bar, enter the telnet command:

```
telnet://<node>
```

<node> is the IP address or fully-qualified domain name of a node that runs the telnet server.

If you are prompted with a security warning, permit the action.

In Firefox, select the **Remember my choice for telnet links** check box.

Third-Party SSH Client (Standard Windows and Windows on Windows)

This procedure applies to the following cases:

- 32-bit Internet Explorer on a 32-bit or 64-bit operating system
- 32-bit Firefox on a 32-bit or 64-bit operating system
- 64-bit Internet Explorer on a 64-bit operating system

To configure a third-party SSH client for use by a web browser, follow these steps:

- 1 Obtain and install a third-party SSH client.

This procedure gives examples for the PuTTY client installed to C:\Program Files\PuTTY\putty.exe. The PuTTY client is available from <http://www.putty.org>.

- 2 Because PuTTY cannot correctly parse the "ssh://<node>" input, this example includes a script that strips the "ssh://" from the input argument. The script C:\Program Files\PuTTY\ssh.js contains the following commands:

```
host = WScript.Arguments(0).replace(/ssh:/, "").replace(/\/\/g, "");
shell = WScript.CreateObject("WScript.Shell");
shell.Run("\"c:\\Program Files\\PuTTY\\putty.exe\" -ssh " + host);
```



This script was created for this example and is not included with PuTTY.

- 3 Define the ssh protocol.
 - a Back up the Windows registry.

- b Use the Windows registry editor to add the [HKEY_CLASSES_ROOT\ssh] key with the following values:

Name	Type	Data
(default)	REG_SZ	URL:ssh Protocol
EditFlags	REG_DWORD	2
FriendlyTypeName	REG_SZ	Secure Shell
URL Protocol	REG_SZ	<i>no value</i>

- 4 Set file association for the URL:ssh Protocol file type.
 - a Back up the Windows registry.
 - b Use the Windows registry editor to modify the [HKEY_CLASSES_ROOT\ssh\shell\open\command] key with the following value:

Name	Type	Data
(default)	REG_SZ	"C:\Windows\System32\WScript.exe" "C:\Program Files\PuTTY\ssh.js" %1

%1 (with a lowercase L) is the complete ssh argument, including the protocol specification. The ssh.js script passes the ssh target to PuTTY.



In a .reg file, escape each quotation mark (") and backslash (\) character with a backslash (\) character.

- 5 Restart the web browser, and then, in the browser address bar, enter the ssh command:

ssh: // <node>

<node> is the IP address or fully-qualified domain name of a node that runs the telnet server.

If you are prompted with a security warning, permit the action.

In Firefox, select the **Remember my choice for ssh links** check box.

Configure Firefox to use Telnet or SSH on Linux

On the Linux operating system, define the telnet or ssh protocol, and then configure Firefox to use the new protocol.

To complete any of the procedures in this section, you must have administrative privileges on the computer.

For more information, see http://kb.mozillazine.org/Register_protocol.


Telnet on Linux

To configure Firefox on the Linux operating system to use the telnet protocol, follow these steps:

- 1 Define the telnet protocol.
 - a Create the `/usr/local/bin/nmmtelnet` file with the following contents:


```
#!/bin/bash
#
# Linux shell script called by Firefox in response to
# telnet:// URLs for the NNMi telnet menu.
#
address=`echo $1 | cut -d : -f 2 | sed 's/;/;/g'`
port=`echo $1 | cut -d : -f 3`
exec /usr/bin/xterm -e telnet $address $port
```
 - b Set the script permissions to be executable by everyone:


```
chmod 755 /usr/local/bin/nmmtelnet
```
- 2 Configure Firefox preferences for telnet.
 - a In the Firefox address bar, enter: **about:config**
 - b In the preference list, right-click, click **New**, and then click **Boolean**.
 - c Enter the preference name: **network.protocol-handler.expose.telnet**
 - d Select the preference value: **false**
- 3 Configure Firefox to use the newly-defined protocol.
 - a Browse to a telnet link.

 You can create a simple HTML file containing the link, or you can use **Actions > Telnet... (from Client)** in the NNMi console. Typing the link directly into the address bar does not have the same effect.
 - b In the Launch Application window, click **Choose**, and then select `/usr/local/bin/nmmtelnet`.
 - c Select the **Remember my choice for telnet links** check box.


Secure Shell on Linux

To configure Firefox on the Linux operating system to use the ssh protocol, follow these steps:

- 1 Define the ssh protocol.
 - a Create the `/usr/local/bin/nmssh` file with the following contents:


```
#!/bin/bash
#
# Linux shell script called by Firefox in response to
# ssh:// URLs for the NNMi SSH menu.
#
address=`echo $1 | cut -d : -f 2 | sed 's/;/;/g'`
port=`echo $1 | cut -d : -f 3`
exec /usr/bin/xterm -e ssh $address $port
```
 - b Set the script permissions to be executable by everyone:


```
chmod 755 /usr/local/bin/nmssh
```
- 2 Configure Firefox preferences for SSH.
 - a In the Firefox address bar, enter: **about:config**
 - b In the preference list, right-click, click **New**, and then click **Boolean**.
 - c Enter the preference name: **network.protocol-handler.expose.ssh**
 - d Select the preference value: **false**
- 3 Configure Firefox to use the newly-defined protocol.
 - a Browse to an SSH link.

 You can create a simple HTML file containing the link, or you can use the new SSH menu item that you defined in the NNMi console. Typing the link directly into the address bar does not have the same effect.
 - b In the Launch Application window, click **Choose**, and then select `/usr/local/bin/nmssh`.
 - c Select the **Remember my choice for ssh links** check box.

Example Files for Changing the Windows Registry

If many NNMi users need to use the telnet or ssh protocols to access managed nodes from the NNMi console, you might be able to automate the Windows registry updates with one or more `.reg` files. This section contains example `.reg` files on which you can base the creation of your own `.reg` files. Note that the registry keys are located in a different path for running 32-bit applications on 64-bit versions of Windows than they are for when the application and operating system match.

For more information, see the Microsoft article at <http://support.microsoft.com/kb/310516>.

Example nntelnet.reg

This registry content example applies to [Windows Operating System-Provided Telnet Client](#) on page 152.

Windows Registry Editor Version 5.00

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\MAIN\FeatureControl\FEATURE_DISABLE_TELNET_PROTOCOL]
"iexplore.exe"=dword:00000000
```

```
[HKEY_CLASSES_ROOT\telnet\shell\open\command]
@="\"C:\\Windows\\system32\\rundll32.exe\"
\"C:\\Windows\\system32\\url.dll\",TelnetProtocolHandler %1"
```

Example nnpullytelnet.reg

This registry content example applies to [Third-Party Telnet Client \(Standard Windows\)](#) on page 153.

Windows Registry Editor Version 5.00

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Internet Explorer\MAIN\FeatureControl\FEATURE_DISABLE_TELNET_PROTOCOL]
"iexplore.exe"=dword:0c000000
```

```
[HKEY_CLASSES_ROOT\telnet\shell\open\command]
@="\"C:\\Program Files\\PuTTY\\putty.exe\" %1"
```

Example nntelnet32on64.reg

This registry content example applies to [Third-Party Telnet Client \(Windows on Windows\)](#) on page 154.

Windows Registry Editor Version 5.00

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Microsoft\Internet Explorer\MAIN\FeatureControl\FEATURE_DISABLE_TELNET_PROTOCOL]
"iexplore.exe"=dword:00000000
```

```
[HKEY_CLASSES_ROOT\Wow6432Node\telnet\shell\open\command]
@="\"C:\\Program Files\\PuTTY\\putty.exe\" %1"
```

Example nnmssh.reg


This registry content example applies to [Third-Party SSH Client \(Standard Windows and Windows on Windows\)](#) on page 155.

Windows Registry Editor Version 5.00

```
[HKEY_CLASSES_ROOT\ssh]
@="URL:ssh Protocol"
"EditFlags"=dword:00000002
"FriendlyTypeName"="Secure Shell"
"URL Protocol"=""
```

```
[HKEY_CLASSES_ROOT\ssh\shell\open\command]
@="\"C:\\Windows\\System32\\WScript.exe\" \"c:\\Program Files\\PuTTY\\ssh.js\" %1"
```


Integrating NNMi with a Directory Service through LDAP



This chapter contains information about integrating NNMi with a directory service for consolidating the storage of user names, passwords, and, optionally, NNMi user group assignments. It contains the following topics:

- [NNMi User Access Information and Configuration Options](#) on page 161
- [Configuring NNMi to Access a Directory Service](#) on page 165
- [Changing the Directory Service Access Configuration to Support the NNMi Security Model](#) on page 173
- [Directory Service Queries](#) on page 175
- [Directory Service Configuration for Storing NNMi User Groups](#) on page 185
- [Troubleshooting the Directory Service Integration](#) on page 186
- [ldap.properties Configuration File Reference](#) on page 187

NNMi User Access Information and Configuration Options

Together, the following items define an NNMi user:

- The **user name** uniquely identifies the NNMi user. The user name provides access to NNMi and receives incident assignments.
- The **password** is associated with the user name to control access to the NNMi console or NNMi command.
- **NNMi user group** membership controls the information available and the type of actions that a user can take in the NNMi console. User group membership also controls the availability of NNMi commands to the user.

NNMi provides several options for where the NNMi user access information is stored, as described in the following topics. [Table 8](#) indicates the databases that store the NNMi user access information for each configuration option.

Table 8 Options for Storing User Information

Option	User Name	Password	User Group	User Group Membership
1	NNMi	NNMi	NNMi	NNMi
2	Both	Directory Service	NNMi	NNMi
3	Directory Service	Directory Service	Both	Directory Service

When NNMi is integrated with a directory service for some or all of the user access information, the user account and user group definition statement on the **Server** tab of the **System Information** window indicates the type of information that was obtained through LDAP queries.

Single sign-on (SSO) between NNMi and other applications is not dependent on how the NNMi user access information is configured or where this information is stored.

Option 1: All NNMi User Information in the NNMi Database

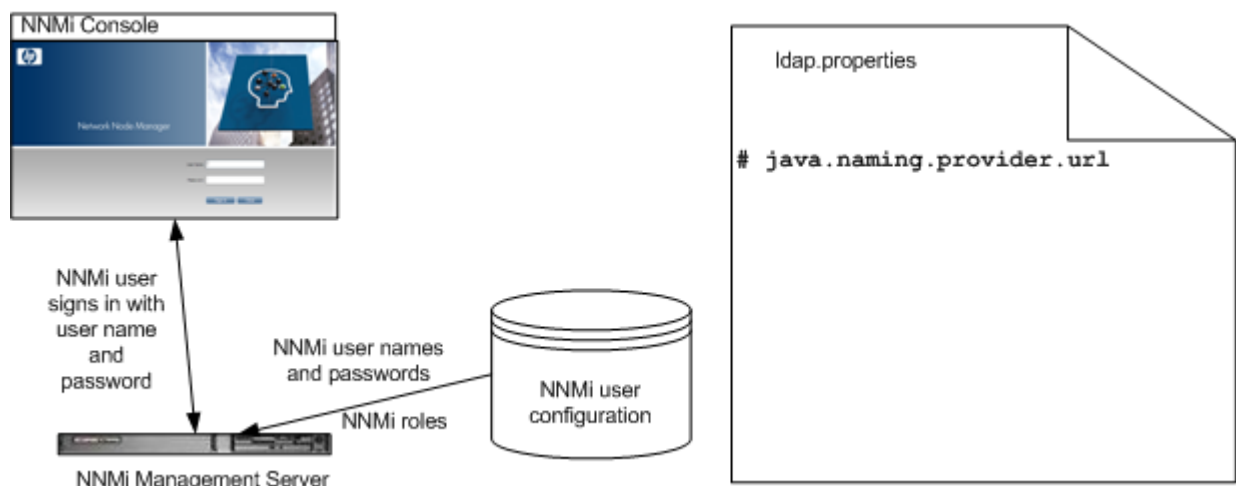
With configuration option 1, NNMi accesses the NNMi database for all user access information, which the NNMi administrator defines and maintains in the NNMi console. The user access information is local to NNMi. NNMi does not access a directory service, and NNMi ignores the `ldap.properties` file (as indicated by the commented line in [Figure 9](#)).

[Figure 9](#) shows the information flow for this option, which is appropriate in the following situations:

- The number of NNMi users is small.
- No directory service is available.

For information about setting up all user information in the NNMi database, see *Control Access with NNMi Accounts* in the NNMi help. You do not need to read this chapter.

Figure 9 NNMi User Sign-in Information Flow for Option 1



Option 2: Some NNMi User Information in the NNMi Database and Some NNMi User Information in the Directory Service

With configuration option 2, NNMi accesses a directory service for the user name and password, which are defined externally to NNMi and are also available to other applications. The mapping of users to NNMi user groups is maintained in the NNMi console. The configuration and maintenance of NNMi user access information is a joint effort as described here:

- The directory service administrator maintains the user names and password in the directory service.
- The NNMi administrator enters the user names (as defined in the directory service), user group definitions, and the user group mappings in the NNMi console.
- The NNMi administrator configures the NNMi `ldap.properties` file to describe the directory service database schema for user names to NNMi. (In [Figure 10](#), the commented line indicates that NNMi does not pull user group information from the directory service.)

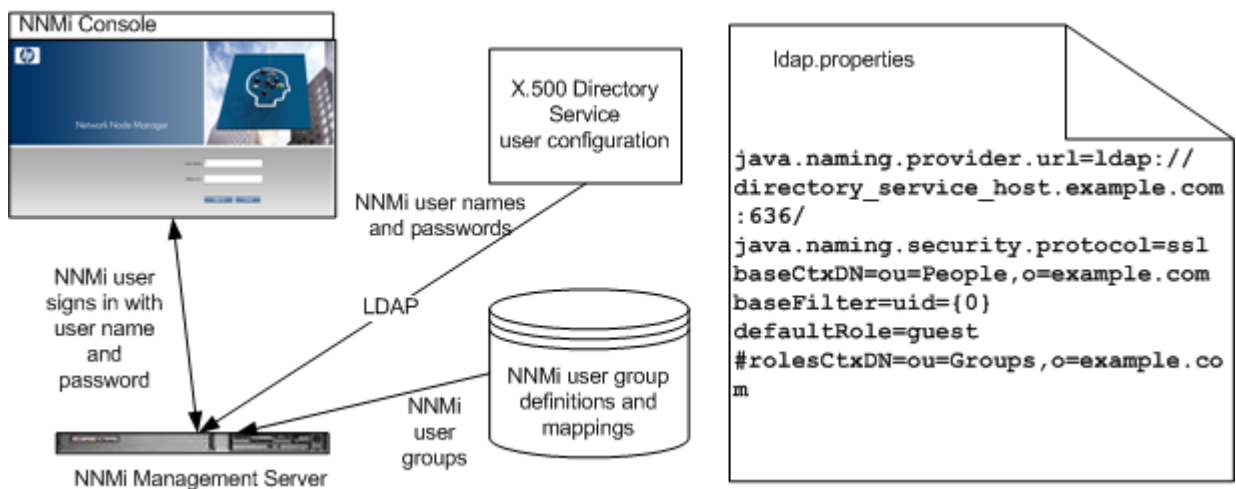
Because user names must be entered in two places, user name maintenance must be performed in both places.

[Figure 10](#) shows the information flow for this option, which is appropriate in the following situations:

- The number of NNMi users is small, and a directory service is available.
- The NNMi administrator wants to control the user groups instead of requiring a directory service change for each user group change.
- The directory service group definitions are not easily expandable.

For information about integrating with a directory service for the user name and password, see the rest of this chapter and *Control Access Using Both Directory Service and NNMi* in the NNMi help.

Figure 10 NNMi User Sign-in Information Flow for Option 2



Option 3: All NNMi User Information in the Directory Service

With configuration option 3, NNMi accesses a directory service for all user access information, which is defined externally to NNMi and is available to other applications. Membership in one or more directory service groups determines the NNMi user groups for the user.

The configuration and maintenance of NNMi user access information is a joint effort as described here:

- The directory service administrator maintains the user names, passwords, and group membership in the directory service.
- The NNMi administrator maps the directory service groups to NNMi user groups in the NNMi console.
- The NNMi administrator configures the NNMi `ldap.properties` file to describe the directory service database schema for user names and groups to NNMi.


Figure 11 shows the information flow for this option, which is appropriate for environments where the directory service can be modified to include user groups that align with the people who need access to NNMi.

Because this option is an expansion of the option 2 scenario, HP recommends the following configuration process:

- 1 Configure and verify NNMi user name and password retrieval from the directory service.
- 2 Configure NNMi user group retrieval from the directory service.

For information about integrating with a directory service for all user information, see the rest of this chapter and *Control Access with a Directory Service* in the NNMi help.

Figure 11 NNMi User Sign-in Information Flow for Option 3



Configuring NNMi to Access a Directory Service

Directory service access is configured in the following file:

- *Windows*: %NNM_SHARED_CONF%\ldap.properties
- *UNIX*: \$NNM_SHARED_CONF/ldap.properties

For information about this file, see [ldap.properties Configuration File Reference](#) on page 187. Also see [Examples](#) on page 192.

For information about the general structure of a directory service, see [Directory Service Queries](#) on page 175.

For configuration option 2, complete the following tasks:

- [Task 1: Back up the Current NNMi User Information](#)
- [Task 2: Optional. Configure Secure Communications to the Directory Service](#)
- [Task 3: Configure User Access from the Directory Service](#)
- [Task 4: Test the User Name and Password Configuration](#)
- [Task 9: Clean up to Prevent Unexpected Access to NNMi](#)
- [Task 10: Optional. Map the User Groups to Security Groups](#)

For configuration option 3, complete the following tasks:

- [Task 1: Back up the Current NNMi User Information](#)
- [Task 2: Optional. Configure Secure Communications to the Directory Service](#)
- [Task 3: Configure User Access from the Directory Service](#)
- [Task 4: Test the User Name and Password Configuration](#)
- [Task 5: \(Configuration Option 3 only\) Configure Group Retrieval from the Directory Service](#)



If you plan to store NNMi user groups in the directory service, the directory service must be configured with the NNMi user groups. For more information, see [Directory Service Configuration for Storing NNMi User Groups](#) on page 185.

- [Task 6: \(Configuration Option 3 only\) Map the Directory Service Groups to NNMi User Groups](#)
- [Task 7: \(Configuration Option 3 only\) Test the NNMi User Group Configuration](#)
- [Task 8: \(Configuration Option 3 only\) Configure NNMi User Groups for Incident Assignment](#)
- [Task 9: Clean up to Prevent Unexpected Access to NNMi](#)
- [Task 10: Optional. Map the User Groups to Security Groups](#)

Task 1: Back up the Current NNMi User Information

Back up the user information in the NNMi database:

```
nnmconfigexport.ovpl -c account -u <user> \
-p <password> -f NNMi_database_accounts.xml
```

Task 2: Optional. Configure Secure Communications to the Directory Service

If the directory service requires the use of secure sockets layer (SSL), import your company's certificate into the NNMi trust store as described in [Configuring an SSL Connection to the Directory Service](#) on page 138.

Task 3: Configure User Access from the Directory Service

Complete this task for configuration options 2 and 3. Follow the appropriate procedure for your directory service. This task includes the following sections:

- [Simple Approach for Microsoft Active Directory](#)
- [Simple Approach for Other Directory Services](#)

(For detailed configuration instructions, see [User Identification](#) on page 180.)

Simple Approach for Microsoft Active Directory

- 1 Back up the `ldap.properties` file that was shipped with NNMi, and then open the file in any text editor.
- 2 Overwrite the file contents with the following text:

```
java.naming.provider.url=ldap://<myldapserver>:389/

bindDN=<mydomain>\<myusername>
bindCredential=<mypassword>

baseCtxDN=CN=Users,DC=<myhostname>,DC=<mycompanyname>,DC=<mysuffix>
baseFilter=CN={0}

defaultRole=guest

#rolesCtxDN=CN=Users,DC=<myhostname>,DC=<mycompanyname>,DC=<mysuffix>
roleFilter=member={1}
uidAttributeID=member
userRoleFilterList=admin;level2;level1
```

- 3 Specify the URL for accessing the directory service. In the following line:

```
java.naming.provider.url=ldap://<myldapserver>:389/
```

Replace `<myldapserver>` with the fully-qualified hostname of the Active Directory server (for example: `myserver.example.com`).



To specify multiple directory service URLs, separate each URL with a single space character ().

- 4 Specify credentials for a valid directory service user. In the following lines:

```
bindDN=<mydomain>\<myusername>
bindCredential=<mypassword>
```

Make the following substitutions:

- Replace `<mydomain>` with the name of the Active Directory domain.

- Replace `<myusername>` and `<mypassword>` with a user name and password for accessing the Active Directory server.
If you plan to add the password in plain text, specify a user name with read-only access to the directory service.
If you plan to specify an encrypted password, use the following command to encrypt the plain text password before adding it to the `ldap.properties` file:
`nnmlldap.ovpl -encrypt <mypassword>`



This encrypted password only works for the NNMi instance you create it for. Do not attempt to use it for a different NNMi instance.

For more information see the `nnmlldap.ovpl` reference page, or the UNIX manpage.

- 5 Specify the portion of the directory service domain that stores user records. In the following line:

```
baseCtxDN=CN=Users,DC=<myhostname>,DC=<mycompanyname>,
DC=<mysuffix>
```

Replace `<myhostname>`, `<mycompanyname>`, and `<mysuffix>` with the components of the fully-qualified hostname of the Active Directory server (for example, for the hostname `myserver.example.com`, specify: `DC=myserver,DC=example,DC=com`).

Simple Approach for Other Directory Services

- 1 Back up the `ldap.properties` file that was shipped with NNMi, and then open the file in any text editor.
- 2 Specify the URL for accessing the directory service. In the following line:

```
#java.naming.provider.url=ldap://<myldapserver>:389/
```

Do the following:

- Uncomment the line (by deleting the # character).
- Replace `<myldapserver>` with the fully-qualified hostname of the directory server (for example: `myserver.example.com`).



To specify multiple directory service URLs, separate each URL with a single space character ().

- 3 Specify the portion of the directory service domain that stores user records. In the following line:

```
baseCtxDN=ou=People,o=myco.com
```

Replace `ou=People,o=myco.com` with the portion of the directory service domain that stores user records.

- 4 Specify the format of user names for signing in to NNMi. In the following line:

```
baseFilter=uid={0}
```

Replace `uid` with the user name attribute from the directory service domain.

Task 4: Test the User Name and Password Configuration

- 1 In the `ldap.properties` file, set `defaultRole=guest` for testing purposes. (You can change this value at any time.)
- 2 Save the `ldap.properties` file.

- 3 Force NNMi to re-read the `ldap.properties` file by running the following command:

```
nnmldap.ovpl -reload
```

- 4 Log on to the NNMi console with a user name and password that are defined in the directory service.



Run this test with a user name that is not already defined in the NNMi database.

- 5 Verify the user name and NNMi role (Guest) in the title bar of the NNMi console.
 - If user signin works correctly, continue with [step 8](#) of this task.
 - If user signin does not work correctly, continue with [step 6](#), next.



After each test, sign out of the NNMi console to clear the session credentials.

- 6 Test the configuration for one user by running the following command:

```
nnmlldap.ovpl -diagnose <NNMi_user>
```

Replace *<NNMi_user>* with the sign-in name of an NNMi user as defined in the directory service.

Examine the command output and respond appropriately. Suggestions include:

- Verify that you completed [Task 3](#) on page 166 correctly.
 - Follow the detailed configuration process in [User Identification](#) on page 180.
- 7 Repeat [step 1](#) through [step 5](#) until you see the expected result when signing in to the NNMi console.
- 8 After you can log on, choose your strategy:
- If you plan to store NNMi user group membership in the NNMi database (configuration option 2), continue with [Task 9](#) on page 172.
 - If you plan to store NNMi user group membership in the directory service (configuration option 3), continue with [Task 5](#), next.

Task 5: (Configuration Option 3 only) Configure Group Retrieval from the Directory Service

Complete this task for configuration option 3. Follow the appropriate procedure for your directory service. This task includes the following sections:

- [Simple Approach for Microsoft Active Directory](#)
- [Simple Approach for Other Directory Services](#)

(For detailed configuration instructions, see [User Group Identification](#) on page 183.)

Simple Approach for Microsoft Active Directory

- 1 Back up the `ldap.properties` file, and then open the file in any text editor.
- 2 Specify the portion of the directory service domain that stores group records. In the following line:

```
#rolesCtxDN=CN=Users,DC=<myhostname>,DC=<mycompanyname>,
DC=<mysuffix>
```

Do the following:

- Uncomment the line (by deleting the # character).
- Replace *<myhostname>*, *<mycompanyname>*, and *<mysuffix>* with the components of the fully-qualified hostname of the Active Directory server (for example, for the hostname `myserver.example.com`, specify: `DC=myserver,DC=example,DC=com`).

Simple Approach for Other Directory Services

- 1 Back up the `ldap.properties` file, and then open the file in any text editor.
- 2 Specify the portion of the directory service domain that stores group records. In the following line:

```
#rolesCtxDN=ou=Groups,o=myco.com
```


Do the following:

- Uncomment the line (by deleting the # character).
 - Replace `ou=Groups,o=myco.com` with the portion of the directory service domain that stores group records.
- 3 Specify the format of group member names in the directory service group definitions. In the following line:

```
roleFilter=member={1}
```



Replace `member` with the name of the group attribute that stores the directory service user ID in the directory service domain.

Task 6: (Configuration Option 3 only) Map the Directory Service Groups to NNMi User Groups

- 1 In the NNMi console, map the predefined NNMi user groups to their counterparts in the directory service:
 - a Open the **User Groups** view.
In the **Configuration** workspace, expand **Security**, and then click **User Groups**.
 - b Double-click the **admin** row.
 - c In the **Directory Service Name** field, enter the full distinguished name of the directory service group for NNMi administrators.
 - d Click  **Save and Close**.
 - e Repeat **step b** through **step d** for each of the **guest**, **level1**, and **level2** rows.



These mappings provide NNMi console access. Every user who will access the NNMi console must be in a directory service group that is mapped to one of the predefined NNMi user groups named in this step.

- 2 For other groups containing one or more NNMi users in the directory service, create a new user group in the NNMi console:
 - a Open the **User Groups** view.
In the **Configuration** workspace, expand **Security**, and then click **User Groups**.
 - b Click  **New**, and then enter the information for the group:
 - Set **Unique Name** to any unique value. Short names are recommended.
 - Set **Display Name** to the value users should see.
 - Set **Directory Service Name** to the full distinguished name of the directory service group.
 - Set **Description** to text that describes the purpose of this NNMi user group.
 - c Click  **Save and Close**.

- d Repeat [step b](#) and [step c](#) for each additional directory service group of NNMi users.



These mappings provide topology object access in the NNMi console. Each directory service group can be mapped to multiple NNMi user groups.

Task 7: (Configuration Option 3 only) Test the NNMi User Group Configuration

- 1 Save the `ldap.properties` file.
- 2 Force NNMi to re-read the `ldap.properties` file by running the following command:


```
nnmlldap.ovpl -reload
```
- 3 Log on to the NNMi console with a user name and password that are defined in the directory service.



Run this test with a user name that is not already defined in the NNMi database and is a member of a directory service group that is mapped to the admin, level1, or level2 NNMi user group.

- 4 Verify the user name and NNMi role (as configured in the **Display Name** field in the **User Group** view) in the title bar of the NNMi console.
 - If user signin works correctly, continue with [Task 8](#) on page 171.
 - If user signin does not work correctly, continue with [step 5](#), next.



After each test, sign out of the NNMi console to clear the session credentials.

- 5 Test the configuration for one user by running the following command:

```
nnmlldap.ovpl -diagnose <NNMi_user>
```

Replace `<NNMi_user>` with the sign-in name of an NNMi user as defined in the directory service.

Examine the command output and respond appropriately. Suggestions include:

- Verify that you completed [Task 5](#) on page 169 correctly.
 - Verify that you completed [Task 6](#) on page 170 correctly for each of the predefined NNMi user groups.
 - Follow the detailed configuration process in [User Group Identification](#) on page 183.
- 6 Repeat [step 1](#) through [step 4](#) until you see the expected result when signing in to the NNMi console.

Task 8: (Configuration Option 3 only) Configure NNMi User Groups for Incident Assignment

- 1 Back up the `ldap.properties` file, and then open the file in any text editor.
- 2 Modify the `userRoleFilterList` parameter value to specify the NNMi roles to which NNMi operators can assign incidents.



The format is a semicolon-separated list of the unique names for one or more of the predefined NNMi user group names (as defined in [Table 11](#) on page 183).

- 3 Save the `ldap.properties` file.

- 4 Force NNMi to re-read the `ldap.properties` file by running the following command:


```
nnmldap.ovpl -reload
```
- 5 Log on to the NNMi console with a user name and password that are defined in the directory service.
- 6 In any incident view, select an incident, and then click **Actions > Assign > Assign Incident**. Verify that you can assign the incident to a user in each of the NNMi roles specified by the `userRoleFilterList` parameter.
- 7 Repeat [step 1](#) through [step 6](#) until you can assign an incident to each configured NNMi role.

Task 9: [Clean up to Prevent Unexpected Access to NNMi](#)

- 1 Optional. Change the value of, or comment out, the `defaultRole` parameter in the `ldap.properties` file.
- 2 (Configuration Option 2 only) To store user group membership in the NNMi database, reset the user access information in the NNMi database as follows:
 - a Remove any pre-existing user access information. (Delete all rows in the **User Accounts** view.)

For instructions, see *Delete a User Account* in the NNMi help.
 - b For each NNMi user, create a new object in the **User Accounts** view for the user name.
 - For the **Name** field, enter the user name as defined in the directory service.
 - Select the **Directory Service Account** check box.
 - Do not specify a password.

For more information, see *User Account Tasks* in the NNMi help.
 - c For each NNMi user, map the user account to one or more NNMi user groups.

For instructions, see *User Account Mapping Tasks* in the NNMi help.
 - d Update incident ownership so that each assigned incident is associated with a valid user name.

For instructions, see *Manage Incident Assignments* in the NNMi help.
- 3 (Configuration Option 3 only) To rely on the user group membership in the directory service, reset the user access information in the NNMi database as follows:
 - a Remove any pre-existing user access information. (Delete all rows in the **User Accounts** view.)

For instructions, see *Delete a User Account* in the NNMi help.
 - b Update incident ownership so that each assigned incident is associated with a valid user name.

For instructions, see *Manage Incident Assignments* in the NNMi help.

Task 10: [Optional. Map the User Groups to Security Groups](#)

For instructions, see *Security Group Mapping Tasks* in the NNMi help.

Changing the Directory Service Access Configuration to Support the NNMi Security Model

The information in this section describes how to revise an `ldap.properties` file from NNMi 8.1x or 9.0x to support multiple NNMi user groups per user. This revision is necessary under *both* of the following conditions:

- The `ldap.properties` file currently enables NNMi user access configuration option 3 (all NNMi user information in the directory service).
- NNMi has been or will be configured with custom security groups.

In NNMi 8.1x and 9.0x, NNMi users were assigned to one of the predefined NNMi roles. Each user had access to all objects in the NNMi topology.

In NNMi 9.10, the predefined NNMi user groups replace NNMi roles. Each NNMi user must belong to at least one predefined NNMi user group, which defines what an NNMi user can do in the NNMi console. Additional user groups, if they exist, limit access to NNMi topology objects as follows:

- If no custom user groups exist, all NNMi console users can access all topology objects.
- If one or more custom user groups exist, each of these user groups provide access to a subset of objects in the NNMi topology.

NNMi 8.1x and 9.0x required each directory service group definition to include a group attribute that named the NNMi role. In the `ldap.properties` configuration file, the following parameters specified this group attribute:

- `roleAttributeID`
- `roleAttributeIsDN`
- `roleNameAttributeID`



NNMi 9.10 deprecates these parameters. They will become unsupported in a future release.

In NNMi 9.10, each user group must be defined in the NNMi console. The user group definition includes an external name, which is the distinguished name of the group in the directory service.

To change the directory service access configuration to support the NNMi security model, follow these steps:

- 1 Back up the user information in the NNMi database:


```
nnmconfigexport.ovpl -c account -u <user> \
-p <password> -f NNMi_database_accounts.xml
```
- 2 Back up the `ldap.properties` file, and then open the file in any text editor.



For information about the `ldap.properties` file, see [ldap.properties Configuration File Reference](#) on page 187. For information about the deprecated parameters, see the *NNMi Deployment Reference* for the previous version of NNMi.

3 Comment out or delete the following parameters (if they exist):

- roleAttributeID
- roleAttributeIsDN
- roleNameAttributeID



The `roleAttributeID` parameter is the flag that tells NNMi which method to use for identifying NNMi user groups. When `roleAttributeID` is set, NNMi uses the NNMi 8.1x and 9.0x approach. When `roleAttributeID` is not set, NNMi uses the NNMi 9.10 approach.

4 In the NNMi console, map the predefined NNMi user groups to their counterparts in the directory service:

- a Open the **User Groups** view.

In the **Configuration** workspace, expand **Security**, and then click **User Groups**.

- b Double-click the **admin** row.

- c In the **Directory Service Name** field, enter the full distinguished name of the directory service group for NNMi administrators.

- d Click  **Save and Close**.

- e Repeat [step b](#) through [step d](#) for each of the **guest**, **level1**, and **level2** rows.



These mappings provide NNMi console access. Every user who will access the NNMi console must be in a directory service group that is mapped to one of the predefined NNMi user groups named in this step.

5 In the directory service, identify additional groups of NNMi users. Define new groups as needed.

6 For each new group added in [step 5](#), create a new user group in the NNMi console:

- a Open the **User Groups** view.

In the **Configuration** workspace, expand **Security**, and then click **User Groups**.

- b Click  **New**, and then enter the information for the group:

— Set **Unique Name** to any unique value. Short names are recommended.

— Set **Display Name** to the value users should see.

— Set **Directory Service Name** to the full distinguished name of the directory service group.

— Set **Description** to text that describes the purpose of this NNMi user group.

- c Click  **Save and Close**.

- d Repeat [step b](#) and [step c](#) for each new directory service group of NNMi users.



These mappings provide topology object access in the NNMi console. Each directory service group can be mapped to multiple NNMi user groups.

7 Optional. Map the user groups to security groups.

For information, see *Configuring Security* in the NNMi help.

Directory Service Queries

NNMi uses LDAP to communicate with a directory service. NNMi sends a request, and the directory service returns stored information. NNMi cannot alter the information that is stored in the directory service.

This section contains the following topics:

- [Directory Service Access](#)
- [Directory Service Content](#)
- [Information Owned by the Directory Service Administrator](#)
- [User Identification](#)
- [User Group Identification](#)

Directory Service Access

LDAP queries to a directory service use the following format:

ldap://<directory_service_host>:<port>/<search_string>

- `ldap` is the protocol indicator. Use this indicator for both standard connections and SSL connections to the directory service.
- `<directory_service_host>` is the fully-qualified name of the computer that hosts the directory service.
- `<port>` is the port that the directory service uses for LDAP communication. The default port for non-SSL connections is 389. The default port for SSL connections is 636.
- `<search_string>` contains the information request. For more information, see [Directory Service Content](#) and RFC 1959, *An LDAP URL Format*, which is available at:
labs.apache.org/webarch/uri/rfc/rfc1959.txt

You can enter an LDAP query as a URL in a web browser to verify that you have the correct access information and the correct structure for the search string.



If the directory service (for example, Active Directory) does not permit anonymous access, the directory service denies LDAP queries from a web browser. In this case, you can use a third-party LDAP browser (for example, the LDAP browser included in Apache Directory Studio) to validate your configuration parameters.

Directory Service Content

A directory service stores information such as user names, passwords, and group membership. To access the information in a directory service, you must know the distinguished name that references the storage location of the information. For sign-in applications, the distinguished name is a combination of variable information (such as a user name) and fixed information (such as the storage location of user names). The elements that make up a distinguished name depend on the structure and content of the directory service.

The following examples show possible definitions for a group of users called USERS-NNMi-Admin. This group lists the directory service user IDs that have administrative access to NNMi. The following information pertains to these examples:

- The Active Directory example is for the Windows operating system.
- The other directory services example is for UNIX operating systems.
- The file shown in each example is a portion of a lightweight directory interchange format (LDIF) file. LDIF files provide for sharing directory service information.
- The figure shown in each example is a graphical representation of the directory service domain that provides an expanded view of the information in the LDIF file excerpt.

Example content
structure for Active
Directory

In this example, the following items are of interest:

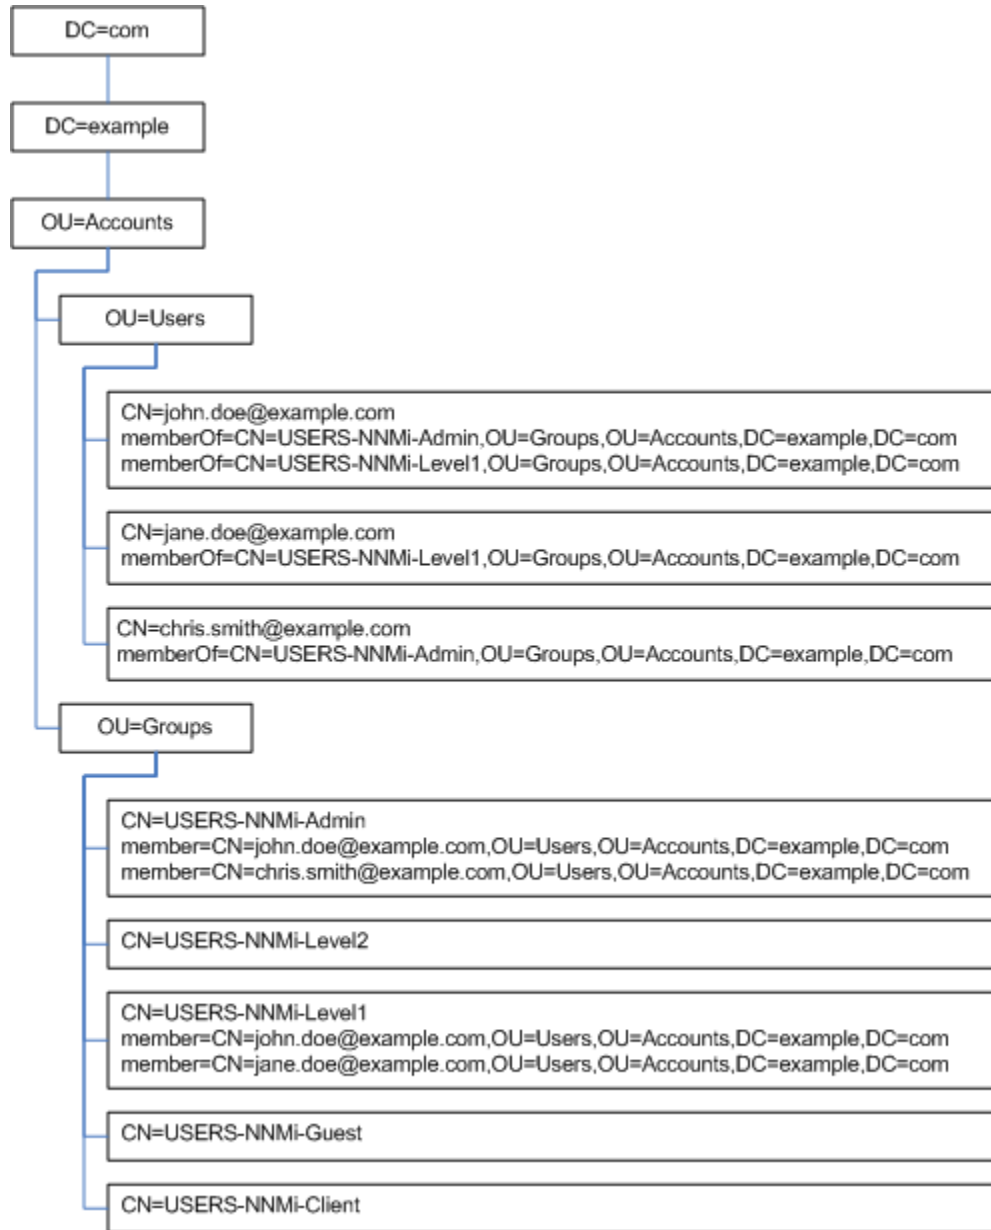
- The distinguished name of the user John Doe is:
CN=john.doe@example.com,OU=Users,OU=Accounts,DC=example,DC=com
- The distinguished name of the group USERS-NNMi-Admin is:
CN=USERS-NNMi-Admin,OU=Groups,OU=Accounts,DC=example,DC=com
- The group attribute that stores the directory service user ID is: member

Example LDIF file excerpt:

```
groups |USERS-NNMi-Admin
dn: CN=USERS-NNMi-Admin,OU=Groups,OU=Accounts,DC=example,DC=com
cn: USERS-NNMi-Admin
description: Group of users for NNMi administration.
member: CN=john.doe@example.com,OU=Users,OU=Accounts,
DC=example,DC=com
member: CN=chris.smith@example.com,OU=Users,OU=Accounts,
DC=example,DC=com
```

[Figure 12](#) on page 177 illustrates this directory service domain.

Figure 12 Example Domain for Active Directory



Example content structure for other directory services

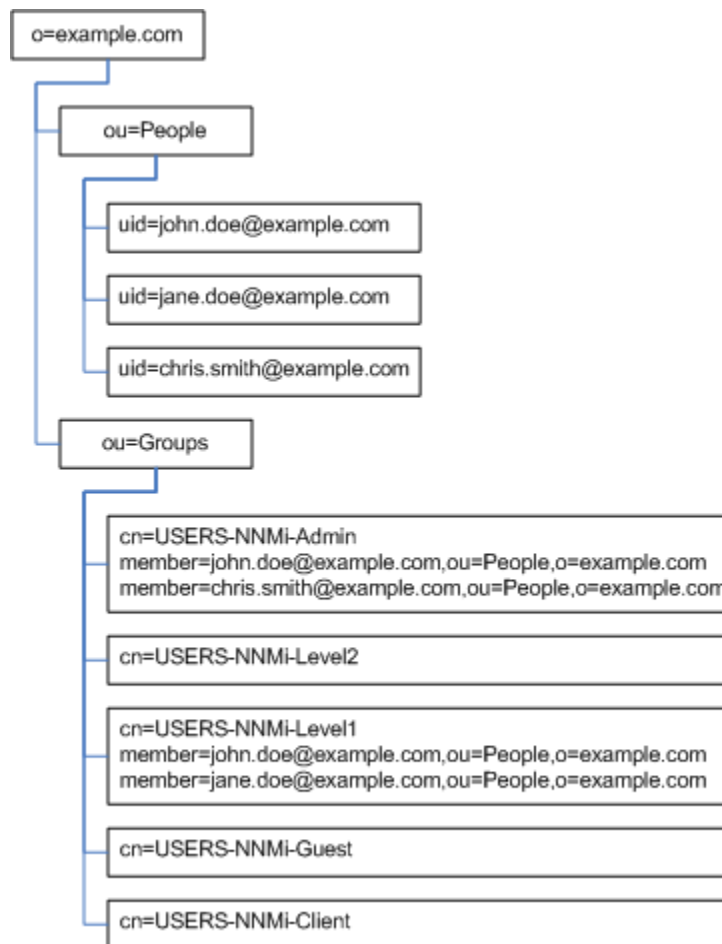
In this example, the following items are of interest:

- The distinguished name of the user John Doe is:
uid=john.doe@example.com,ou=People,o=example.com
- The distinguished name of the group USERS-NNMi-Admin is:
cn=USERS-NNMi-Admin,ou=Groups,o=example.com
- The group attribute that stores the directory service user ID is: member

Example LDIF file excerpt:

```
groups |USERS-NNMi-Admin
dn: cn=USERS-NNMi-Admin,ou=Groups,o=example.com
cn: USERS-NNMi-Admin
description: Group of users for NNMi administration.
member: uid=john.doe@example.com,ou=People,o=example.com
member: uid=chris.smith@example.com,ou=People,o=example.com
```

Figure 13 Example Domain for Other Directory Services



Information Owned by the Directory Service Administrator

Table 9 and Table 10 list the information to obtain from the directory service administrator before configuring NNMi for LDAP access to a directory service.

- If you plan to use the directory service for user names and passwords only (configuration option 2), gather the information for Table 9.
- If you plan to use the directory service for all NNMi access information (configuration option 3), gather the information for Table 9 and Table 10.

Table 9 Information for Retrieving User Names and Passwords from a Directory Service

Information	Active Directory Example	Other Directory Services Example
The fully-qualified name of the computer that hosts the directory service	directory_service_host.example.com	
The port that the directory service uses for LDAP communication	<ul style="list-style-type: none"> • 389 for non-SSL connections • 636 for SSL connections 	
Does the directory service require an SSL connection?	If yes, obtain a copy of your company's trust store certificate and see Configuring an SSL Connection to the Directory Service on page 138.	
The distinguished name for one user name that is stored in the directory service (to demonstrate the directory service domain)	CN=john.doe@example.com, OU=Users,OU=Accounts, DC=example,DC=com	uid=john.doe@example.com, ou=People,o=example.com

Table 10 Information for Retrieving Group Membership from a Directory Service

Information	Active Directory Example	Other Directory Services Example
The distinguished name for identifying the groups to which a user is assigned	The memberOf user attribute identifies the groups.	<ul style="list-style-type: none"> • ou=Groups,o=example.com • cn=USERS-NNMi-*, ou=Groups,o=example.com
The method of identifying a user within a group	<ul style="list-style-type: none"> • CN=john.doe@example.com, OU=Users,OU=Accounts, DC=example,DC=com • CN=john.doe@example.com 	<ul style="list-style-type: none"> • cn=john.doe@example.com, ou=People,o=example.com • cn=john.doe@example.com
The group attribute that stores the directory service user ID	member	member

Table 10 Information for Retrieving Group Membership from a Directory Service (cont'd)

Information	Active Directory Example	Other Directory Services Example
The names of the groups in the directory service that apply to NNMi access	<ul style="list-style-type: none"> • CN=USERS-NNMi-Admin, OU=Groups, OU=Accounts, DC=example, DC=com • CN=USERS-NNMi-Level2, OU=Groups, OU=Accounts, DC=example, DC=com • CN=USERS-NNMi-Level1, OU=Groups, OU=Accounts, DC=example, DC=com • CN=USERS-NNMi-Client, OU=Groups, OU=Accounts, DC=example, DC=com • CN=USERS-NNMi-Guest, OU=Groups, OU=Accounts, DC=example, DC=com 	<ul style="list-style-type: none"> • cn=USERS-NNMi-Admin, ou=Groups, o=example.com • cn=USERS-NNMi-Level2, ou=Groups, o=example.com • cn=USERS-NNMi-Level1, ou=Groups, o=example.com • cn=USERS-NNMi-Client, ou=Groups, o=example.com • cn=USERS-NNMi-Guest, ou=Groups, o=example.com

User Identification

User identification applies to configuration options 2 and 3.

The distinguished name for user identification is the fully-qualified method of locating one user in the directory service. NNMi passes the user distinguished name in an LDAP request to the directory service.

In the `ldap.properties` file, the user distinguished name is the concatenation of the `baseFilter` value and the `baseCtxDN` value. If the password returned by the directory service matches the sign-in password the user entered into the NNMi console, user sign-in continues.

For configuration option 2, the following information applies:

- For NNMi console access, NNMi examines the following information and grants the user the highest possible privileges:
 - The value of the `defaultRole` parameter in the `ldap.properties` file
 - This user's membership in the predefined NNMi user groups in the NNMi console
- For NNMi topology object access, NNMi grants access according to the security group mappings for the NNMi user groups to which this user belongs in the NNMi console.

For configuration option 3, the following information applies:

- For NNMi console access, NNMi examines the following information and grants the user the highest possible privileges:
 - The value of the `defaultRole` parameter in the `ldap.properties` file
 - This user's membership in the directory service groups that are mapped (with the **Directory Service Name** field) to the predefined NNMi user groups in the NNMi console

- For NNMi topology object access, NNMi grants access according to the security group mappings for the groups to which this user belongs in the directory service (as mapped to NNMi user groups in the NNMi console).

Active Directory user identification example

If `baseFilter` is set to `CN={0}`, `baseCtxDN` is set to `OU=Users,OU=Accounts,DC=example,DC=com`, and a user signs in to NNMi as `john.doe`, the string passed to the directory service is:

```
CN=john.doe,OU=Users,OU=Accounts,DC=example,DC=com
```

Other directory services user identification example

If `baseFilter` is set to `uid={0}@example.com`, `baseCtxDN` is set to `ou=People,o=example.com`, and a user signs in to NNMi as `john.doe`, the string passed to the directory service is:

```
uid=john.doe@example.com,ou=People,o=example.com
```

Configuring NNMi User Access from the Directory Service (Detailed Approach)

If the simple approach described in [Task 3](#) on page 166 did not work correctly, follow these steps:

- 1 Obtain the information listed in [Table 9](#) on page 179 from the directory service administrator.
- 2 Verify the format of user names in the directory service by completing the appropriate procedure:
 - *LDAP browser approach for Active Directory and other directory services:* See [Determining How the Directory Service Identifies a User \(LDAP Browser Approach\)](#) on page 182.
 - *Web browser approach for other directory services:* See [Determining How the Directory Service Identifies a User \(Web Browser Approach\)](#) on page 182.
- 3 Open the `ldap.properties` file in any text editor.



For information about the `ldap.properties` file, see [ldap.properties Configuration File Reference](#) on page 187.


- 4 Set the `java.naming.provider.url` parameter to the URL for accessing the directory service through LDAP.
 - *LDAP browser approach:* Obtain this information from the LDAP browser configuration.
 - *Web browser approach:* Include the values of `<directory_service_host>` and `<port>` from [Determining How the Directory Service Identifies a User \(Web Browser Approach\)](#) on page 182.



To specify multiple directory service URLs, separate each URL with a single space character.

- 5 If you configured secure communications to the directory service, uncomment (or add) the following line:

```
java.naming.security.protocol=ssl
```

- 6 (Active Directory only) Set the `bindDN` and `bindCredential` parameters as follows:
- Replace `<mydomain>` with the name of Active Directory domain.
 - Replace `<myusername>` and `<mypassword>` with a user name and password for accessing the Active Directory server.
If you plan to add the password in plain text, specify a user name with read-only access to the directory service.
If you plan to specify an encrypted password, use the following command to encrypt the plain text password before adding it to the `ldap.properties` file:
`nnmldap.ovpl -encrypt <mypassword>`
-  This encrypted password only works for the NNMi instance you create it for. Do not attempt to use it for a different NNMi instance.
- For more information see the `nnmldap.ovpl` reference page, or the UNIX manpage.
- 7 Set the `baseCtxDN` parameter to the elements of the distinguished user name that are the same for multiple users.
- 8 Set the `baseFilter` parameter to correlate user names as they are entered for NNMi signin to the way user names are stored in the directory service.
This value is the element of the distinguished user name that changes for each user. Replace the actual user name with the expression `{0}`.
- 9 Test the configuration as described in [Task 4](#) on page 167.

Determining How the Directory Service Identifies a User (LDAP Browser Approach)

In a third-party LDAP browser, do the following:

- 1 Navigate to the portion of the directory service domain that stores group information.
- 2 Identify a group of users, and then examine the format of the distinguished names for the users associated with that group.

Determining How the Directory Service Identifies a User (Web Browser Approach)

- 1 In a supported web browser, enter the following URL:

`ldap://<directory_service_host>:<port>/<user_search_string>`

 - `<directory_service_host>` is the fully-qualified name of the computer that hosts the directory service.
 - `<port>` is the port that the directory service uses for LDAP communication.
 - `<user_search_string>` is the distinguished name for one user name that is stored in the directory service.
- 2 Evaluate the results of the directory service access test.
 - If the request times out or you see a message that the directory service could not be reached, verify the values of `<directory_service_host>` and `<port>`, and then repeat [step 1](#).
 - If you see a message that the directory service does not contain the requested entry, verify the value of `<user_search_string>`, and then repeat [step 1](#).
 - If you see the appropriate user record, the access information is correct. The value of `<user_search_string>` is the distinguished user name.

User Group Identification

User group identification applies to configuration option 3.

NNMi determines the user groups for an NNMi user as follows:

- 1 NNMi compares the values of the external names of all user groups configured in the NNMi console with the names of the directory service groups.
- 2 For any user group match, NNMi then determines whether the NNMi user is a member of that group in the directory service.

In the NNMi console, short text strings identify the unique names of the predefined NNMi user groups that grant NNMi console access. These text strings are also required by the `defaultRole` and `userRoleFilterList` parameters in the `ldap.properties` configuration file. [Table 11](#) maps the unique names of these groups to their display names.

Table 11 NNMi User Group Name Mappings

NNMi Role Name in the NNMi Console	User Group Unique Name and Text String in NNMi Configuration Files
Administrator	admin
Global Operators	globalops
Operator Level 2	level2
Operator Level 1	level1
Guest	guest
Web Service Client	client




The NNMi Global Operators User Group (`globalops`) grants access to all topology objects only. A user must be assigned to one of the other User Groups (`level2`, `level1`, or `guest`) to access the NNMi console.

The administrator should not map the `globalops` User Group to any security group because this User Group is, by default, mapped to all security groups.

Configuring User Group Retrieval from the Directory Service (Detailed Approach)

If the simple approach described in [Task 5](#) on page 169 did not work correctly, follow these steps:

- 1 Obtain the information listed in [Table 10](#) on page 179 from the directory service administrator.
- 2 Verify the format of group names and group members in the directory service by completing the appropriate procedure:
 - *LDAP browser approach for Active Directory:* See [Determining How the Directory Service Identifies a Group and Group Membership \(LDAP Browser Approach for Active Directory\)](#) on page 184.
 - *LDAP browser approach for other directory services:* See [Determining How the Directory Service Identifies a Group and Group Membership \(LDAP Browser Approach for Other Directory Services\)](#) on page 184.

- *Web browser approach for other directory services:* See [Determining How the Directory Service Identifies a Group \(Web Browser Approach\)](#) on page 184.
- 3 Open the `ldap.properties` file in any text editor.
 For information about the `ldap.properties` file, see [ldap.properties Configuration File Reference](#) on page 187.
 - 4 Set the `rolesCtxDN` parameter to the elements of the distinguished group name that are the same for multiple groups.
 - 5 Set the `roleFilter` parameter to correlate user names to the way user names are stored for groups in the directory service. Replace the actual user name with one of the following expressions:
 - Use `{0}` to denote the user name entered for signin (for example, `john.doe`).
 - Use `{1}` to denote the distinguished name of the authenticated user as returned by the directory service (for example, `uid=john.doe@example.com,ou=People,o=example.com`).
 - 6 Set the `uidAttributeID` parameter to the name of the group attribute that stores the user ID.
 - 7 Test the configuration as described in [Task 7](#) on page 171.

[Determining How the Directory Service Identifies a Group and Group Membership \(LDAP Browser Approach for Active Directory\)](#)

In a third-party LDAP browser, do the following:

- 1 Navigate to the portion of the directory service domain that stores user information.
- 2 Identify a user who requires access to NNMi, and then examine the format of the distinguished names for the groups associated with that user.
- 3 Navigate to the portion of the directory service domain that stores group information.
- 4 Identify the groups that correspond to NNMi user groups, and then examine the format of the names for the users associated with a group.

[Determining How the Directory Service Identifies a Group and Group Membership \(LDAP Browser Approach for Other Directory Services\)](#)

In a third-party LDAP browser, do the following:

- 1 Navigate to the portion of the directory service domain that stores group information.
- 2 Identify the groups that correspond to NNMi user groups, and then examine the format of the distinguished names for those groups.
- 3 Also examine the format of the names for the users associated with a group.

[Determining How the Directory Service Identifies a Group \(Web Browser Approach\)](#)

- 1 In a supported web browser, enter the following URL:


```
ldap://<directory_service_host>:<port>/<group_search_string>
```

 - `<directory_service_host>` is the fully-qualified name of the computer that hosts the directory service.

- `<port>` is the port that the directory service uses for LDAP communication.
 - `<group_search_string>` is the distinguished name for a group name that is stored in the directory service, for example:
`cn=USERS-NNMi-Admin,ou=Groups,o=example.com`
- 2 Evaluate the results of the directory service access test.
 - If you see a message that the directory service does not contain the requested entry, verify the value of `<group_search_string>`, and then repeat [step 1](#).
 - If you see the appropriate list of groups, the access information is correct.
 - 3 Examine the group properties to determine the format of the names for the users associated with that group.

Directory Service Configuration for Storing NNMi User Groups

If you plan to store NNMi user groups in the directory service (configuration option 3), the directory service must be configured with NNMi user group information. Ideally, the directory service already contains appropriate user groups. If this is not the case, the directory service administrator can create new user groups specifically for NNMi user group assignment.

Because directory service configuration and maintenance procedures depend on the specific directory service software and your company's policies, those procedures are not documented here.

Troubleshooting the Directory Service Integration

- 1 Verify the NNMi LDAP configuration by running the following command:

```
nnmlldap.ovpl -info
```

If the reported configuration is not as expected, verify the settings in the `ldap.properties` file.

- 2 Force NNMi to re-read the `ldap.properties` file by running the following command:

```
nnmlldap.ovpl -reload
```

- 3 Test the configuration for one user by running the following command:

```
nnmlldap.ovpl -diagnose <NNMi_user>
```

Replace `<NNMi_user>` with the sign-in name of an NNMi user as defined in the directory service.

Examine the command output and respond appropriately.

- 4 Verify that the directory service contains the expected records. Use a web browser or a third-party LDAP browser (for example, the LDAP browser included in Apache Directory Studio) to examine the directory service information.

Information about the format of a query to a directory service can be found in RFC 1959, *An LDAP URL Format*, which is available at:

<http://labs.apache.org/webarch/uri/rfc/rfc1959.txt>

- 5 View the `%NnmDataDir%\log\nnm\jbossServer.log` or (*Windows*) or `/var/opt/OV/log/nnm/jbossServer.log` (*UNIX*) log file to verify that the sign-in request is correct, and to determine if any errors occurred:

- A message similar to the following line indicates that the directory service requires HTTPS communication. In this case, enable SSL as described in [Configuring an SSL Connection to the Directory Service](#) on page 138.

```
javax.naming.AuthenticationNotSupportedException: [LDAP:
error code 13 - confidentiality required]
```

- A message similar to the following line indicates that a timeout occurred while communicating with the directory service. In this case, increase the value of `searchTimeLimit` in the `nms-ldap.properties` file.

```
javax.naming.TimeLimitExceededException: [LDAP: error code 3
- Timelimit Exceeded]
```

ldap.properties Configuration File Reference

The `ldap.properties` file contains the settings for communicating with and building LDAP queries to the directory service. This file is located as follows:

- *Windows*: `%NNM_SHARED_CONF%\ldap.properties`
- *UNIX*: `$NNM_SHARED_CONF/ldap.properties`

In the `ldap.properties` file, the following conventions apply:

- To comment out a line, begin that line with a number sign character (#).
- The following rules apply to special characters:
 - To specify a backslash character (\), comma (,), semicolon (;), plus sign (+), less than sign (<), or greater than sign (>), escape the character with a backslash character. For example: `\\` or `\+`
 - To include a space character () as the *first* or *last* character in a string, escape the space character with a backslash character (\).
 - To include a number sign character (#) as the *first* character in a string, escape the number sign character with a backslash character (\).

Characters not mentioned here do not need to be escaped or quoted.

▶ After editing the `ldap.properties` file, force NNMi to re-read the LDAP configuration by running the following command:

```
nnmlldap.ovpl -reload
```

Table 12 describes the parameters in the `ldap.properties` file.

▶ The initial `ldap.properties` file might not include all parameters that are listed in Table 12. Add the parameters you need.

Table 12 Parameters in the ldap.properties File

Parameter	Description
<code>java.naming.provider.url</code>	<p>Specifies the URL for accessing the directory service.</p> <p>The format is the protocol (<code>ldap</code>), followed by the fully-qualified host name of the directory server, optionally followed by the port number. For example:</p> <pre>java.naming.provider.url=ldap://ldap.example.com:389/</pre> <p>If the port number is omitted the following defaults apply:</p> <ul style="list-style-type: none"> • For non-SSL connections, the default port is 389. • For SSL connections, the default port is 636. <p>If you specify multiple directory service URLs, NNMi uses the first directory service when possible. If that directory service is not accessible, NNMi queries the next directory service in the list, and so forth. Separate each URL with a single space character. For example:</p> <pre>java.naming.provider.url=ldap://ldap1.example.com/ ldap:// ldap2.example.com/</pre> <p>Configuring this parameter enables LDAP communication between NNMi and the directory service. To disable LDAP communication, comment out this parameter, and then save the file. NNMi ignores the configuration in the <code>ldap.properties</code> file.</p>

Table 12 Parameters in the ldap.properties File (cont'd)

Parameter	Description
java.naming.security.protocol	<p>Specifies the connection protocol specification.</p> <ul style="list-style-type: none"> If the directory service is configured to use LDAP over SSL, set this parameter to <code>ssl</code>. For example: <code>java.naming.security.protocol=ssl</code> If the directory service does not require SSL, leave this parameter commented out. <p>For more information, see Configuring an SSL Connection to the Directory Service on page 138.</p>
bindDN	<p>For a directory service (such as Active Directory) that does not permit anonymous access, specify the user name for accessing the directory service.</p> <p>For example: <code>bindDN=region1\john.doe@example.com</code></p> <ul style="list-style-type: none"> If you plan to add the password in plain text, specify a user name with read-only access to the directory service. For example: <code>bindCredential=PasswordForJohnDoe</code> If you plan to specify an encrypted password, use the following command to encrypt the plain text password before adding it to the <code>ldap.properties</code> file: <code>nnldap.ovpl -encrypt <mypassword></code> For example: <code>bindCredential={ENC}uaF22C+0CF9VozBVYj8OAw==</code> <p>This encrypted password only works for the NNMi instance you create it for. Do not attempt to use it for a different NNMi instance. For more information see the <i>nnldap.ovpl</i> reference page, or the UNIX manpage.</p>
bindCredential	<p>When <code>bindDN</code> is set, specifies the password for the user name that <code>bindDN</code> identifies. For example: <code>bindCredential=PasswordForJohnDoe</code></p>
baseCtxDN	<p>Specifies the portion of the directory service domain that stores user records.</p> <p>The format is a comma-separated list of directory service attribute names and values. For example:</p> <ul style="list-style-type: none"> <code>baseCtxDN=CN=Users,DC=ldapserver,DC=example,DC=com</code> <code>baseCtxDN=ou=People,o=example.com</code> <p>For more information, see User Identification on page 180.</p>

Table 12 Parameters in the ldap.properties File (cont'd)

Parameter	Description
baseFilter	<p>Specifies the format of user names for signing in to NNMi.</p> <p>The format is the name of the directory service user name attribute and a string that relates the entered user sign-in name to the format of names in the directory service. The user name string contains the expression {0} (to denote the user name entered for signin) and any other characters that are needed to match the directory service formatting of user names.</p> <ul style="list-style-type: none"> • If the user name entered for NNMi signin is the same as the user name stored in the directory service, the value is the replacement expression. For example: <ul style="list-style-type: none"> – baseFilter=CN={0} – baseFilter=uid={0} • If the user name entered for NNMi signin is as subset of the user name stored in the directory service, include the additional characters in the value. For example: <ul style="list-style-type: none"> – baseFilter=CN={0}@example.com – baseFilter=uid={0}@example.com <p>For more information, see User Identification on page 180.</p>
defaultRole	<p>Optional. Specifies a default role that applies to any directory service user who signs in to NNMi through LDAP. The value of this parameter applies regardless of where user group mappings are stored (in the NNMi database or in the directory service).</p> <p>If a user is directly configured for a predefined NNMi user group, NNMi grants the user the superset of privileges for the default role and the assigned user group.</p> <p>Valid values are as follows: <code>admin</code>, <code>level2</code>, <code>level1</code>, or <code>guest</code>.</p> <p>Note that although <code>admin</code> is a valid value, you should use caution and consider the implications of making <code>admin</code> a default role.</p> <p>These names are the unique names of the predefined NNMi user group names (as defined in Table 11 on page 183).</p> <p>For example:</p> <pre>defaultRole=guest</pre> <p>If commented out or omitted, NNMi does not use a default role.</p>

Table 12 Parameters in the ldap.properties File (cont'd)

Parameter	Description
rolesCtxDN	<p>Specifies the portion of the directory service domain that stores group records.</p> <p>The format is a comma-separated list of directory service attribute names and values. For example:</p> <ul style="list-style-type: none"> • <code>rolesCtxDN=CN=Users,DC=ldapserver,DC=example,DC=com</code> • <code>rolesCtxDN=ou=Groups,o=example.com</code> <p>In other directory services (not Active Directory), for a faster search, you can identify one or more directory service groups that contain NNMi user groups. If the group names form a pattern, you can specify a wildcard. For example, if the directory service includes groups named <code>USERS-NNMi-administrators</code>, <code>USERS-NNMi-level1Operators</code>, and so forth, you could use a search context similar to:</p> <pre>rolesCtxDN=cn=USERS-NNMi-*,ou=Groups,o=example.com</pre> <p>Configuring this parameter enables directory service queries for NNMi user group assignments through LDAP.</p> <p>To disable directory service queries for NNMi user group assignments through LDAP, comment out this parameter, and then save the file. NNMi ignores the remaining user group-related values in the <code>ldap.properties</code> file.</p> <p>For more information, see User Group Identification on page 183.</p>
roleFilter	<p>Specifies the format of group member names in the directory service group definitions.</p> <p>The format is the name of the directory service group attribute for user ID and a string that relates the entered user sign-in name to the format of user IDs in the directory service. The user name string contains one of the following expressions and any other characters that are needed to match the directory service formatting of group member names.</p> <ul style="list-style-type: none"> • The expression <code>{0}</code> denotes the user name entered for signin (for example, <code>john.doe</code>). An example role filter that matches on the (short) user name entered for signin is: <code>roleFilter=member={0}</code> • The expression <code>{1}</code> denotes the distinguished name of the authenticated user as returned by the directory service (for example, <code>CN=john.doe@example.com,OU=Users,OU=Accounts,DC=example,DC=com</code> or <code>uid=john.doe@example.com,ou=People,o=example.com</code>). An example role filter that matches on the (full) authenticated user name is: <code>roleFilter=member={1}</code> <p>For more information, see User Group Identification on page 183.</p>
uidAttributeID	<p>Specifies the group attribute that stores the directory service user ID.</p> <p>For example:</p> <pre>uidAttributeID=member</pre> <p>For more information, see User Group Identification on page 183.</p>

Table 12 Parameters in the ldap.properties File (cont'd)

Parameter	Description
userRoleFilterList	<p>Optional. Limits the NNMi user groups whose associated users can be assigned incidents in the NNMi console.</p> <p>The user groups in this list apply only to directory service user names authenticated through LDAP. This parameter provides functionality that is not available when NNMi user groups are assigned in the NNMi console and stored in the NNMi database.</p> <p>The format is a semicolon-separated list of the unique names for one or more predefined NNMi user group names (as defined in Table 11 on page 183).</p> <pre>userRoleFilterList=admin;globalops;level2;level1</pre>
searchTimeLimit	<p>Optional. Specifies the timeout value in milliseconds. The default value is 10000 (10 seconds). If you are encountering timeouts during NNMi user signin, increase this value.</p> <p>For example:</p> <pre>searchTimeLimit=10000</pre>

Examples

Example
ldap.properties file
for Active Directory

An example ldap.properties file follows for Active Directory:

```
java.naming.provider.url=ldap://MYldapservice.example.com:389/
bindDN=MYdomain\MYusername
bindCredential=MYpassword
baseCtxDN=CN=Users,DC=MYldapservice,DC=EXAMPLE,DC=com
baseFilter=CN={0}
defaultRole=guest
rolesCtxDN=CN=Users,DC=MYldapservice,DC=EXAMPLE,DC=com
roleFilter=member={1}
uidAttributeID=member
userRoleFilterList=admin;level2;level1
```

Example
ldap.properties file
for other directory
services

An example ldap.properties file follows for other directory services:

```
java.naming.provider.url=ldap://MYldapservice.example.com:389/
baseCtxDN=ou=People,o=EXAMPLE.com
baseFilter=uid={0}
defaultRole=guest
rolesCtxDN=ou=Groups,o=EXAMPLE.com
roleFilter=member={1}
uidAttributeID=member
userRoleFilterList=admin;level2;level1
```


Managing Overlapping IP Addresses in NAT Environments

NNMi helps you manage areas of your network that include overlapping IP addresses resulting from implementing Network Address Translation (NAT).

What is NAT?

Network Address Translation is typically used to interconnect a local network to the public (external) Internet. This technology was developed as a solution for the ever-increasing need for more IPv4 addresses. In addition, certain ranges of IP addresses (described in RFC 1918) have been designated as internal only; in other words, not routable over the Internet, further necessitating a technology such as NAT.

Specifically, NAT translates IP header information, substituting external (public) addresses for internal addresses in IP packets that need to transit the public network. NAT accomplishes this by providing a static or dynamic external address.

What are the Benefits of NAT?

Some benefits of NAT include:

- connecting a large number of hosts to the global Internet using a single dynamic public (external) IP address, thereby conserving IP address space
- reuse of private IP addresses
- enhancing security for private networks by keeping internal addressing private from the external network

What Types of NAT are Supported?

NNMi supports the following types of NAT protocols:

- **Static NAT**—A type of NAT in which an internal IP address is mapped to an external IP address, where the external address is always the same IP address (in other words, each node has a static internal/external address pair). This permits an internal host, such as a Web server, to have an unregistered (private) IP address and still be reachable over the Internet.
- **Dynamic NAT**—A NAT scheme in which bindings between external and internal addresses can change between sessions. A NAT scheme in which internal IP addresses are mapped to a public IP address, drawing from a pool of available, registered (public) IP addresses. Typically, the NAT router in a network keeps a table of registered IP addresses, and when an internal IP address requests access to the Internet, the router chooses an IP address from the table that is not currently being used by another internal IP address.
- **Dynamic Port Address Translation (PAT)** (also referred to as **Network Address and Port Translation (NAPT)**)—This type of NAT translates not only the IP address but also the port number. Translating the address and the port number allows a single external address to be used for multiple simultaneous internal address conversations over the Internet.

How is NAT Implemented in NNMi?

NNMi uses tenancy to manage NAT environments. Tenancy is a logical grouping concept, providing Node Group, mapping, and security support. An example of a tenant could be a customer within an Internet provider's network. An Internet provider might use dynamic NAT within their network to reuse internal IP addresses. In such a case, NNMi would use regional managers to manage each customer within the network to ensure proper network security. In other words, one tenant (customer) would not be able to communicate with another tenant (customer) within the regional network. For more information on tenancy, see [NNMi Security and Multi-Tenancy](#) on page 207.

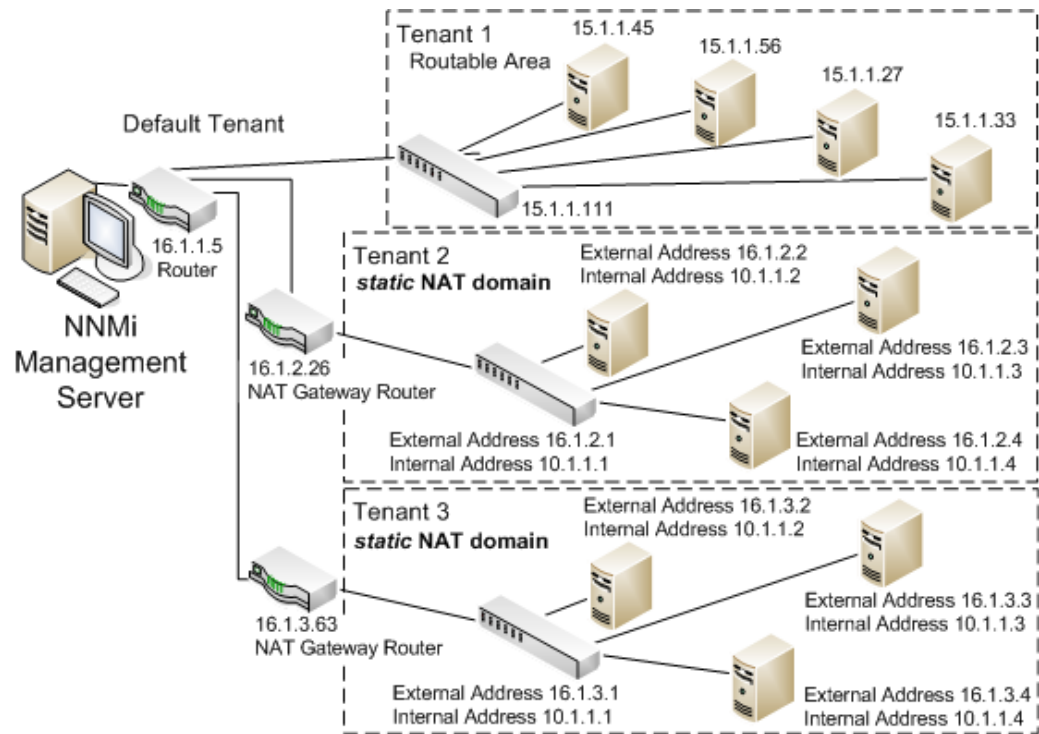
At a minimum, if your network management environment includes overlapping address domains, you must configure each domain as a unique tenant. Depending on the protocols you are using, NNMi's implementation method and requirements might look different. For example, use of dynamic NAT or PAT would require additional hardware and licenses. See the appropriate sections that follow based on your type of NAT protocol.

Static NAT Considerations

Any number of static NAT instances can be monitored by one NNMi management server, as long as each instance is configured with a unique tenant. For more information on tenancy, see [NNMi Security and Multi-Tenancy](#) on page 207 and *Configure Tenants* in the NNMi help.

See [Figure 14](#) for an example of a static NAT configuration.

Figure 14 Example Static NAT Configurations



- ▶ Nodes that belong to the default tenant can have Layer 2 connections to any node in any tenant. Nodes within any tenant other than the default tenant can have Layer 2 connections only to devices within the same tenant or the default tenant.

Subnets are tenant specific (in other words, subnets do not span tenants). The benefit here is that you can use the same subnet on different tenants.

Router Redundancy Groups (RRGs) cannot span tenants.

- ▶ Assign any infrastructure device that interconnects multiple NAT domains (such as the NAT gateway) to the default tenant. This ensures that NNMi displays the Layer 2 connections your workgroup (and customers) need to see.

- ▶ Devices within the default Security Group are visible from all views. To control access to a device, assign that device to a Security Group other than default Security Group.

Hardware and Software Requirements and Static NAT

There are no special hardware or software requirements for static NAT.

Communication and Static NAT

Administering ICMP Polling of the Management Address in a Static NAT Environment

In a NAT environment, a firewall blocks NNMi from communicating with NAT nodes using the IP addresses on the nodes (the private IP addresses). To remedy this, use the NAT address (the public IP address) for communication with NNMi.

In a NAT environment, a node's management address might be different from the IP addresses hosted on the node. For NNMi to discover a node in a NAT environment, you must add the NAT address to NNMi as a discovery seed. NNMi uses this NAT address for communication, even though it is not in the node's `ipAddressTable`.

NNMi provides this feature to avoid generating false node down incidents and a better root cause analysis.

Enabling ICMP Polling of the Management Address in a NAT Environment

By default, NNMi automatically enables ICMP management address polling for all nodes, including those nodes residing in a NAT environment. If you have a NAT environment, it is highly recommended that you do not disable this setting.

To enable ICMP management address polling (if it is disabled), do the following:

- 1 From the workspace navigation panel, select the **Configuration** workspace, expand the **Monitoring** folder, select **Monitoring Configuration**, and locate the **Default Settings** tab.
- 2 Enable ICMP Management Address Polling. See *Set Default Monitoring* in the NNMi help.

View the information NNMi displays after performing **Actions->Monitoring Settings** for SNMP Agents. The displayed information indicates whether NNMi has the management address polling enabled.

How this Changes NNMi

When ICMP Management Address Polling is enabled, NNMi changes as follows:

- The Agent ICMP State field appears in the following forms:
 - Node form
 - SNMP Agent form
 - SNMP Agent table views
- NNMi changes the display location of the management address ICMP state. NNMi also changes the way it determines the SNMP agent status.

Table 13 shows the Agent ICMP and IP Address state polling actions that NNMi takes for the ICMP Management Address Polling and ICMP Fault Polling settings. The shaded first row in **Table 13** shows the default configuration.

Table 13 ICMP Configurations and Resulting State Polling

ICMP Management Address Polling	ICMP Fault Polling	Agent ICMP State	IP Address State
Enabled	Disabled	Polled	Not Polled
Enabled	Enabled	Polled	Polled
Disabled	Disabled	Not Polled	Not Polled
Disabled	Enabled	Not Polled	Polled

Table 14 shows changes to the SNMP Agent Status determined by APA for the SNMP agent and ICMP responses.

Table 14 Determining SNMP Agent Status

SNMP Agent Response	Management Address ICMP Response	SNMP Agent Status
Responding	Responding	Normal
Responding	Not Responding	Minor
Not Responding	Responding	Critical
Not Responding	Not Responding	Critical

With ICMP polling of the management address enabled, APA now considers the management address ICMP response and the SNMP agent response when generating conclusions and generating incidents.

Discovery and Static NAT

NNMi uses tenancy to support networks with overlapping address domains that might exist within NAT areas of your network management domain. Spiral Discovery requires a Discovery seed (tenant and address pair) to identify each node before NNMi discovers and monitors that node. See the NNMi help for more information.

When adding Discovery seeds (using the `nmmloadseeds.ovpl` command or the NNMi console) in a static NAT environment, be sure to use the node's external (public) IP address. For more information, see the `nmmloadseeds.ovpl` reference page, or the UNIX manpage.



A best practice is to not have duplicate Domain Name System (DNS) names.

Traps and Static NAT

You must make changes to the managed nodes for the NNMi management server to receive SNMP traps from nodes behind the NAT gateway. This section covers two types of SNMP traps: SNMPv2c and SNMPv1.

Note that NNMi must unambiguously resolve the source address of each trap that it receives.

SNMPv2c Traps

Table 15 shows the format of an SNMPv2c trap, with the IP header forming the top section of the table and the SNMP Trap Protocol Data Unit (PDU) forming the lower section of the table.

Table 15 SNMPv2c Trap Format

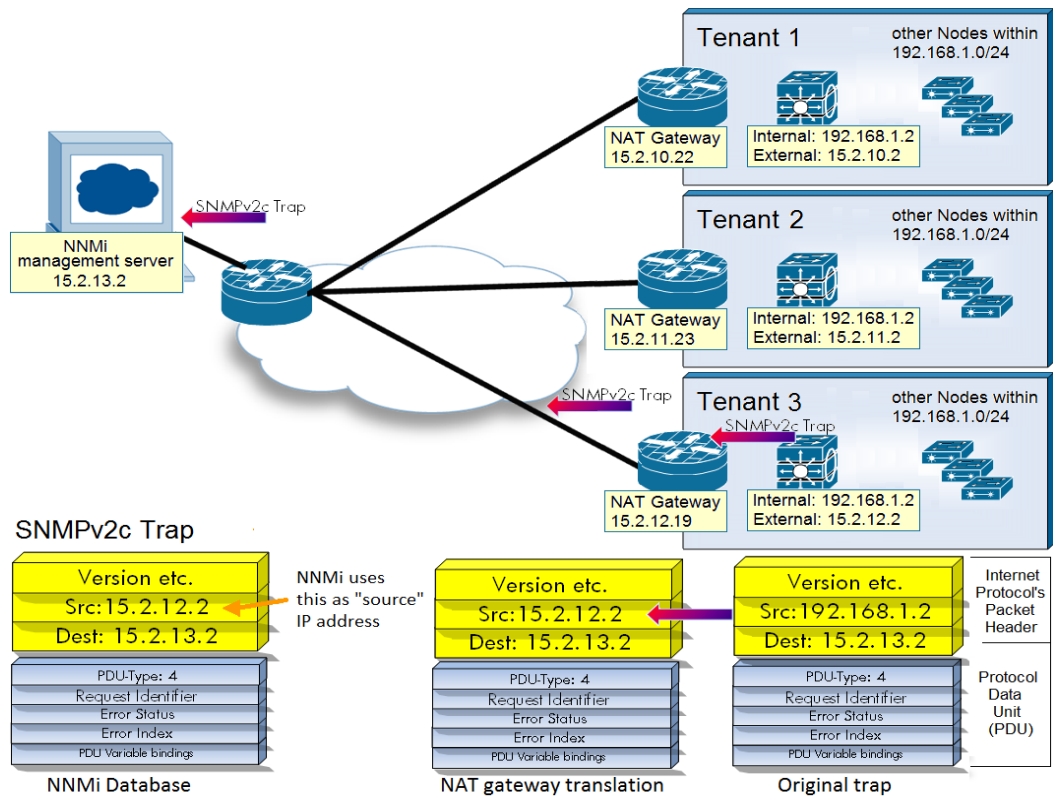
Version and other information
Source Address
Destination Address
PDU-Type: 4
Request Identifier
Error Status
Error Index
PDU Variable Bindings

SNMPv2c traps do not have an Agent Address field in the PDU; therefore, the only source field of the trap is within the IP packet header. NAT routers properly translate the source field.

On the source node, ensure that the interface associated with the private inside IP address sources all traps from devices behind the NAT router. Then, the NAT gateway can translate the trap to the correct public address.

Figure 15 shows an example of correct translation from the NAT gateway. The NAT gateway properly translates a trap that begins with the source address of 192.168.1.2 to address 15.2.13.2. Then the NNMi management server correctly resolves this address.

Figure 15 SNMPv2c Example



SNMPv1 Traps

SNMPv1 traps embed the Agent Address inside the SNMP trap PDU. Table 16 shows the format of an SNMPv1 trap, with the IP header forming the top section and the SNMP trap PDU forming the lower section.

Table 16 SNMPv1 Trap Format

Version and other information
Source Address
Destination Address
PDU-Type: 4
Enterprise
Agent Address
Generic Trap Code
Specific Trap Code
Timestamp
PDU Variable Bindings

Because the Agent Address is embedded in the PDU rather than the header, usually the NAT router will not translate this value. You can enable NNMi to note the address in the header and ignore the Agent Address in the payload by doing the following:

- 1 Edit the following file:
 - *Windows:* %NNM_PROPS%\nms-jboss.properties
 - *UNIX:* \$NNM_PROPS/nms-jboss.properties

2 Find the following line

```
#!com.hp.nnm.trapd.useUdpHeaderIpAddress=false
```

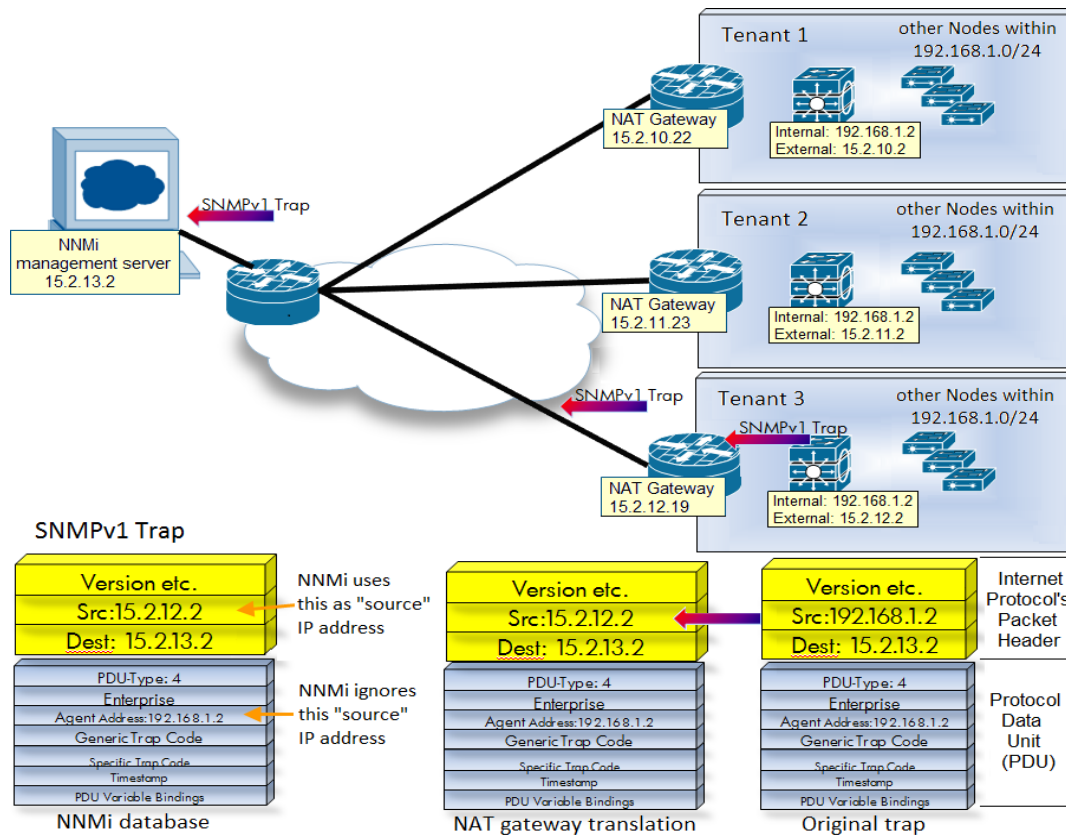
3 Change the value to **true** and remove the #! characters as shown below:

```
com.hp.nnm.trapd.useUdpHeaderIpAddress=true
```

4 Save the file; then restart NNMi.

Figure 16 shows an example of an SNMPv1 trap where NNMi ignores the conflicting IP address fields.

Figure 16 SNMPv1 Example





NNMi provides the following related Custom Incident Attributes (CIAs):

- `cia.agentAddress`—the IP address stored in the SNMPv1 trap data for the SNMP Agent that generated the trap.
- `cia.internalAddress`—If static NAT is part of your network management domain, the NNMi administrator can configure this attribute to show the internal IP address that is mapped to the external management address of the selected incident's Source Node.

The external management IP address (public address) must be mapped to this internal address (private address) using the Overlapping IP Address Mapping form. For more information, see the NNMi help.

Subnets and Static NAT

Note the following with regard to subnets and NAT:

- Subnets are tenant specific (in other words, subnets do not span tenants). The benefit here is that you can use the same subnet on different tenants.
- Subnet filters use a tenant and address pair.
- If you configure a subnet connection rule, the rule applies to all tenants. The members of the subnets must be unique across all tenants (each node assigned to only one tenant). A subnet connection rule can establish a link between the default tenant and another tenant. However, links between two tenants are not allowed unless one of them is the default tenant.

Global Network Management and Static NAT

At least one static or routable (non-translated) address must exist per regional manager. This enables NNMi management servers to communicate with each other, keeping communications internal and secure. For more information on Global Network Management, see [Global Network Management](#) on page 229.


Dynamic NAT and PAT Considerations

One NNMi management server can manage one dynamic NAT or PAT domain. All nodes in this domain must belong to the same unique tenant. The NNMi management server must participate in a Global Network Management environment as a regional manager. See the following figure for an example of a dynamic NAT configuration.



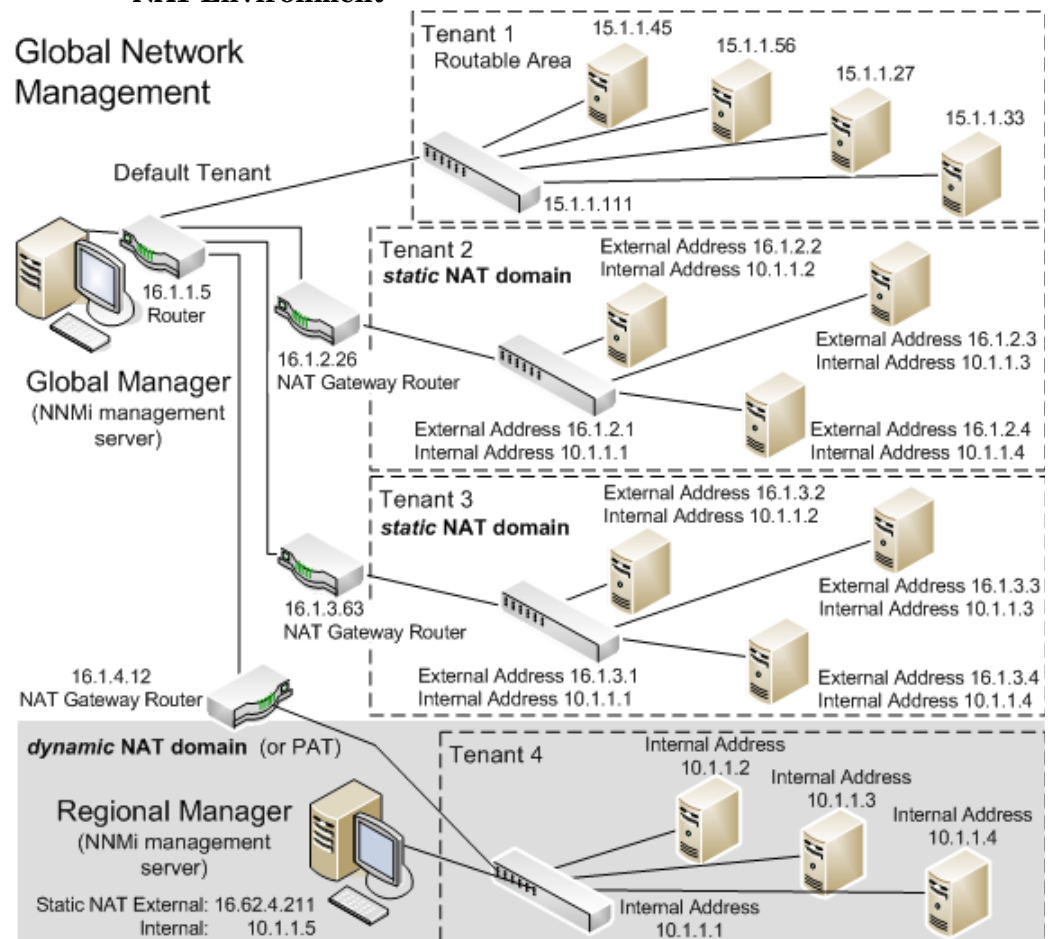
If a regional manager is located behind a NAT firewall, its external (public) address must be static.

Figure 17 Example Dynamic NAT Configurations



Use NNMi’s Global Network Management feature to monitor multiple dynamic NAT and PAT domains. Tenants must be unique within the entire NNMi Global Network Management configuration. See the following figure for an example of a Global Network Management configuration within a NAT environment.

Figure 18 Example Global Network Management Configuration within a NAT Environment



Devices that belong to the default tenant can have Layer 2 connections to any device in any tenant. Devices within any tenant other than default tenant can have Layer 2 connections only to devices within the same tenant or the default tenant.



Assign any infrastructure device that interconnects multiple NAT domains (such as the NAT gateway) to the default tenant. This ensures that NNMi displays the Layer 2 connections your workgroup (and customers) need to see.



Devices within the default Security Group are visible from all views. To control access to a device, assign that device to a Security Group other than default Security Group.

For more information on Global Network Management, see [Global Network Management](#) on page 229. For information about configuring tenants, see *Configure Tenants* in the NNMi help.

Hardware and Software Requirements and Dynamic NAT and PAT

NNMi Advanced software is required for dynamic NAT and PAT environments.

An NNMi regional manager is required for each address domain configured with dynamic NAT or PAT.

Discovery and Dynamic NAT and PAT

NNMi uses tenancy to support networks with overlapping address domains that might exist within dynamic NAT or PAT areas of your network management domain. If you have such networks, put the overlapping address domains into different tenants (this is done using seeded discovery). See the NNMi help for more information.

When adding Discovery seeds (using the `nnmloadseeds.ovpl` command or the graphical user interface) in a dynamic NAT or PAT environment, be sure to use the node's internal IP address.

For more information, see the *nnmloadseeds.ovpl* reference page, the UNIX manpage, or the NNMi help.

Subnets and Dynamic NAT and PAT

Note the following with regard to subnets and NAT:

- Subnets are tenant specific (in other words, subnets do not span tenants). The benefit here is that you can use the same subnet on different tenants.
- Subnet filters use a tenant/address pair.
- If you configure a subnet connection rule, the rule applies to all tenants. The members of the subnets must be unique across all tenants (each node assigned to only one tenant). A subnet connection rule can establish a link between the default tenant and another tenant. However, links between two tenants are not allowed unless one of them is the default tenant.

Global Network Management and Dynamic NAT and PAT

At least one static or routable (non-translated) address must exist per regional manager. This enables NNMi management servers to communicate with each other, keeping communications internal and secure.

▶ If a regional manager is behind a NAT firewall, its external address must be static.

For more information on Global Network Management, see [Global Network Management](#) on page 229. See also *Tenant Best Practices for Global Network Management* the NNMi help.

Overlapping IP Address Mapping

If your network management environment includes overlapping address domains, you must configure each domain as a unique tenant. For more information, see *Configure Tenants* in the NNMi help, and [NNMi Security and Multi-Tenancy](#) on page 207.

Optional. If static NAT is part of your network management domain, and the NNMi management server is outside of that static NAT domain, you can configure NNMi to display the NAT external IP address (public address) in the Mapped Address attribute of the IP Address form for the identified tenant / NAT internal IP address (such as private IPv4 address) pair.

▶ If you are configuring NNMi for areas of your network management domain that use dynamic NAT and PAT, do not use the Overlapping IP Address Mapping form. See [Dynamic NAT and PAT Considerations](#) on page 201.

Your network domain's static NAT configuration might apply to public IP addresses, private IP addresses, or both.

To configure NNMi to display the static NAT external IP address in the Mapped Address attribute of the IP Address form for the identified Tenant and NAT internal IP address pair, do one of the following:

- Use the Overlapping Address Mapping form in the NNMi console.
- Use the `nnmloadipmappings.ovpl` command

For more information, see the NNMi help, or the *nnmloadipmappings.ovpl* reference page, or the UNIX manpage.

Private IP Address Ranges

The Internet Engineering Task Force (IETF) and Internet Assigned Numbers Authority (IANA)'s reserved the following IP address ranges for private networks, for example enterprise local area networks (LANs), corporate offices, or residential networks.


IPv4 private address ranges (RFC 1918):

- 10.0.0.0 – 10.255.255.255 (24-bit block)
- 172.16.0.0 – 172.31.255.255 (20-bit block)
- 192.168.0.0 – 192.168.255.255 (16-bit block)

IPv6 private address ranges:

- fc00::/7 address block = RFC 4193 Unique Local Addresses (ULA)
- fec0::/10 address block = deprecated (RFC 3879)

NNMi Security and Multi-Tenancy



- ▶ NNMi uses tenancy to support networks with overlapping address domains that may exist within static Network Address Translation (NAT), dynamic NAT, or dynamic Port Address Translation (PAT) areas of your network management domain. If you have such networks, put the overlapping address domains into different tenants (this is done using seeded discovery). See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193 and the NNMi help for more information.

By default, all NNMi console users can see information for all objects in the NNMi database. If this default configuration is acceptable for your environment, you do not need to read this chapter.

In NNMi, security and multi-tenancy provide for restricting user access to information about the objects in the NNMi database. This restriction is useful for customizing the views of network operators to their areas of responsibility. It also supports service providers with per-organization configuration of NNMi.

This chapter describes the NNMi security and tenant models and gives suggestions for configuration. It contains the following topics:

- [Effects of Limiting Object Access](#) on page 208
- [The NNMi Security Model](#) on page 209
- [The NNMi Tenant Model](#) on page 214
- [NNMi Security and Multi-Tenancy Configuration](#) on page 217
- [NNMi Security, Multi-Tenancy, and Global Network Management](#) on page 225
- [Including Select Interfaces in NPS Reports](#) on page 228

See also the *HP Network Node Manager i Software Step-by-Step Guide to Using Security Groups White Paper*.

Effects of Limiting Object Access

Configuring NNMi security has the following impacts:

- Topology inventory objects:
 - Each NNMi console user sees only those nodes that match the configuration for their NNMi user account.
 - Sub-node objects, such as interfaces, inherit the access control from the node.
 - Inter-node objects, such as connections, are visible only if the NNMi console user can see at least one of the nodes involved.
 - A NNMi console user sees only those node groups for which they can access at least one node in the group.
 - For Network Performance Server (NPS) reports, the NNMi administrator can selectively override access control inheritance on interfaces. For more information, see [Including Select Interfaces in NPS Reports](#) on page 228.
- Maps and path views:
 - Maps show connections for which the NNMi console user has permission to view both of the participating nodes.
 - Path views omit or show as clouds any intermediate nodes to which the NNMi console user does not have access.
 - For the NNM iSPI for MPLS and the NNM iSPI for IP Multicast, when maps and path views include nodes to which the NNMi console user does not have access, the NNM iSPI displays only the connecting interface and the name of the node. The icons for the inaccessible nodes are white to indicate that status and detailed information are not available for these nodes.
 - For the NNM iSPI for IP Telephony, when maps and path views include nodes to which the NNMi console user does not have access, the NNM iSPI displays only the connecting interface and the name of the node. The icons for the inaccessible nodes show the NNMi status, but all attempted actions fail.
- Incidents:
 - For incidents whose source node is in the NNMi topology, an NNMi console user sees only the incidents for which the user has access to the source node.
 - Incidents that do not have a source node, such as NNMi health and licensing management event incidents, are handled as a group. The NNMi administrator determines which NNMi console users see them (by associating the users with the Unresolved Incidents security group).
 - Incidents that result from traps for which the source node is not in the NNMi topology are handled in the same way as incidents with no source node. If NNMi is configured to generate these incidents, the NNMi administrator determines which NNMi console users see them (by associating the users with the Unresolved Incidents security group).



The incident assignment action does not check user access. It is possible for an NNMi administrator to assign an incident to an NNMi console user who does not have permission to view that incident.

- NNMi console actions:
 - For actions that run without any selections, an NNMi console user sees only those actions they have permission to run.
 - For actions that run against one or more selected objects, an NNMi console user must have the correct access level to the selected objects. Depending on the security configuration, the NNMi console might present actions that are not valid on some of the objects visible in the NNMi console views. Invoking one of these actions results in an error message regarding this limitation.
 - For map views and NNM iSPI table views and forms, NNMi cannot distinguish between unknown nodes and nodes that exist in the NNMi topology but are not accessible by the current user.
- MIB browser and line grapher:
 - An NNMi console user can view MIB data and graphs for nodes to which they have access.
 - An NNMi console user can view MIB data for nodes to which they know the SNMP community string.
- NNMi console URLs:

Users must log on to NNMi before accessing an NNMi console view from a direct URL. NNMi enforces that user's access according to the NNMi security configuration and limits the available topology accordingly.

The NNMi Security Model

The NNMi security model provides user access control to the objects in the NNMi database. This model is appropriate for use by any network management organization that wants to limit NNMi user access to specific objects and incidents. The NNMi security model has the following benefits:

- Provides a way to limit an NNMi console operator's view of the network. Operators can focus on specific device types or network areas.
- Provides for customizing operator access to the NNMi topology. The level of operator access can be configured per node.
- Provides for filtering the Nodes (All Attributes) view and Network Performance Server reports by security group.
- Simplifies the configuration and maintenance of node groups that align with the security configuration.
- Can be used independently of the NNMi tenant model.

Possible use cases for NNMi security include the following:

- Provide NNMi operator focus on equipment type within a site (custom maps).
- Provide NNMi operators at different sites views that show only the nodes at a given site (custom maps).
- Stage nodes during deployment. NNMi administrators see all nodes, while NNMi operators see only the deployed nodes.

- Provide full access to all NOC operators, and limit access to NOC customers.
- Provide full network views to the central NOC operators, and limit the views of the regional NOC operators.

Security Groups

In the NNMi security model, user access to nodes is controlled indirectly through user groups and security groups. Each node in the NNMi topology is associated with only one security group. A security group can be associated with multiple user groups.

Each user account is mapped to the following user groups:

- One or more of the following preconfigured NNMi user groups:
 - NNMi Administrators
 - NNMi Global Operators
 - NNMi Level 2 Operators
 - NNMi Level 1 Operators
 - NNMi Guest Users

This mapping is required for NNMi console access and determines which actions are available within the NNMi console. If a user account is mapped to more than one of these NNMi user groups, the user receives the superset of the permitted actions.



The NNMi Web Services Clients user group does not grant access to the NNMi console; however, it does grant administrator-level access to all NNMi objects.



The NNMi Global Operators User Group (`globalops`) grants access to topology objects only. A user must be assigned to one of the other User Groups (`level2`, `level1`, or `guest`) to access the NNMi console.

The administrator should not map the `globalops` User Group to any security group because this User Group is, by default, mapped to all security groups.

- Zero or more custom user groups that are mapped to security groups.

These mappings provide access to objects in the NNMi database. Each mapping includes an object access privilege level that applies to the nodes for a security group. The object access privilege level also applies to the related database objects, such as interfaces and incidents. For example, a user with Object Operator Level 1 access to node A containing interfaces X and Y has Object Operator Level 1 access to all of the following database objects:

 - Node A
 - Interfaces X and Y
 - Incidents whose source object is node A, interface X, or interface Y

NNMi provides the following security groups:

- **Default Security Group**

In a new NNMi installation, the Default Security Group is the initial security group assignment for all nodes. By default, all users can see all objects in the Default Security Group. The NNMi administrator can configure which nodes are associated with the Default Security Group and which users can access the objects in the Default Security Group.

- **Unresolved Incidents**

The Unresolved Incidents security group provides access to incidents that NNMi creates from received traps whose source node is not in the NNMi topology. By default, all users can see all incidents associated with the Unresolved Incidents security group. The NNMi administrator can configure which users can access the incidents associated with the Unresolved Incidents security group.

All node components inherit the security group assignment of the node.

Best practice

The following best practices apply to NNMi security configuration:

- Map each user account to only one preconfigured NNMi user group.
- Do not map the preconfigured NNMi user groups to security groups.
- Because any user account mapped to the NNMi Administrators user group receives administrator-level access to all objects in the NNMi database, do not map this user account to any other user groups.
- Create a separate user account for the Web Services Client role. Because this user account has access to the entire NNMi topology, map this user account to only the NNMi Web Service Clients user group.

Example Security Group Structure

The three ovals in [Figure 19](#) indicate the primary groupings for which users need to view the nodes in this example NNMi topology. For complete user access control, each of the four unique subgroups corresponds to a unique security group. Each unique security group can be mapped to one or more user groups to represent the available levels of user access to the objects in that security group.

[Table 17](#) on page 212 lists the mappings between the security groups and the possible custom user groups for this topology. (An actual implementation of this security model might not require all of these custom user groups.) [Table 18](#) on page 213 lists the mappings for several user accounts and the user groups for this topology.

Figure 19 Example Topology for User Access Requirements




Table 17 Example Security Group Mappings

Security Group	Nodes of Security Group	User Group	Object Access Privilege
SG1	A, B, C	UG1 Administrator	Object Administrator
		UG1 Level 2	Object Operator Level 2
		UG1 Level 1	Object Operator Level 1
		UG1 Guest	Object Guest
SG2	D, E	UG2 Administrator	Object Administrator
		UG2 Level 2	Object Operator Level 2
		UG2 Level 1	Object Operator Level 1
		UG2 Guest	Object Guest

Table 17 Example Security Group Mappings (cont'd)

Security Group	Nodes of Security Group	User Group	Object Access Privilege
SG3	F, G	UG3 Administrator	Object Administrator
		UG3 Level 2	Object Operator Level 2
		UG3 Level 1	Object Operator Level 1
		UG3 Guest	Object Guest
SG4	H, I, J	UG4 Administrator	Object Administrator
		UG4 Level 2	Object Operator Level 2
		UG4 Level 1	Object Operator Level 1
		UG4 Guest	Object Guest

Table 18 Example User Account Mappings

User Account	User Groups	Node Access	Notes
User Q	NNMi Level 2 Operators	none	This user has operator level 2 access to the nodes in the pink oval (solid line).
	UG1 Level 2	A, B, C	
	UG2 Level 2	D, E	
	UG3 Level 2	F, G	
User R	NNMi Level 1 Operators	none	This user has operator level 1 access to the nodes in the orange oval (dashed line).
	UG2 Level 1	D, E	
User S	NNMi Level 2 Operators	none	This user has operator level 2 access to the nodes in the green oval (dotted line).
	UG3 Level 2	F, G	
	UG4 Level 2	H, I, J	
User T	NNMi Level 2 Operators	none	This user has access (with varying privilege levels) to all nodes in the example topology. This user has administrative access to nodes D and E but cannot see the menu items for tools that require administrative access. If this user has access to the NNMi management server, this user can run command-line tools that require administrative access against nodes D and E only.
	UG1 Guest	A, B, C	
	UG2 Administrator	D, E	
	UG3 Level 2	F, G	
	UG4 Level 1	H, I, J	

The NNMi Tenant Model

The NNMi tenant model provides strict segregation of topology discovery and data into tenants, also called organizations or customers. This model is appropriate for use by service providers, especially managed service providers, and large enterprises. The NNMi tenant model has the following benefits:

- Marks the organization to which each node belongs.
- Provides for filtering the Nodes (All Attributes) inventory view and Network Performance Server reports by tenant and security group.
- Meets regulatory requirements for separating operator access to customer data.
- Simplifies the configuration and maintenance of node groups that align with the tenant configuration.
- Simplifies configuration of NNMi security.
- Provides for management of overlapping address domains when address translation protocols are used.

Use NNMi multi-tenancy to provide different customer views for a service provider that has multiple customers (tenants) managed from the same NNMi management server.



Any number of static Network Address Translation (NAT) instances can be monitored by one NNMi management server, as long as each instance is configured with a unique tenant. See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193, and the NNMi help, for more information.

Tenants

The NNMi tenant model adds the idea of an organization to the security configuration. Each node in the NNMi topology belongs to only one tenant. The tenant provides logical separation in the NNMi database. Object access is managed through security groups.

For each node, the initial discovery tenant assignment occurs when the node is first discovered and added to the NNMi database. For seeded nodes, you can specify the tenant to assign to each node. NNMi assigns all other discovered nodes (those included in an auto-discovery rule but not seeded directly) to the Default Tenant. An NNMi administrator can change the tenant for a node at any time after discovery.

Each tenant definition includes an initial discovery security group. NNMi assigns this initial discovery security group to the node along with the initial discovery tenant. An NNMi administrator can change the security group for a node at any time after discovery.



Changing the tenant assignment of a node does not automatically change the security group assignment.

NNMi provides the Default Tenant. By default, all NNMi users have access (through the Default Security Group) to all objects associated with this tenant.

All node components inherit the tenant and security group assignments of the node.

Best practice

The following best practices apply to NNMi tenant configuration:

- For a small organization, a single security group per tenant is probably sufficient.

- You might want to subdivide a large organization into multiple security groups.
- To prevent users from accessing nodes across organizations, ensure that each security group includes nodes for only one tenant.

Example Tenant Structure

Figure 20 shows an example NNMi topology containing two tenants, represented by the rectangles. The three ovals indicate the primary groupings for which users need to view the nodes. The topology for Tenant 1 is managed as a single group, so it needs only one security group. The topology for Tenant 2 is managed in overlapping sets, so it is separated into three security groups.

Table 19 on page 216 lists the mappings between the security groups and the possible custom user groups for this topology. (An actual implementation of this security model might not require all of these custom user groups.) Table 20 on page 216 lists the mappings for several user accounts and the user groups for this topology.

Figure 20 Example Topology for Multiple Tenants




Table 19 Example Security Group Mappings for Multiple Tenants

Security Group	Nodes of Security Group	User Group	Object Access Privilege
T1 SG	A, B, C, D, E	T1 Administrator	Object Administrator
		T1 Level 2	Object Operator Level 2
		T1 Level 1	Object Operator Level 1
		T1 Guest	Object Guest
T2 SGa	F, G	T2_a Administrator	Object Administrator
		T2_a Level 2	Object Operator Level 2
		T2_a Level 1	Object Operator Level 1
		T2_a Guest	Object Guest
T2 SGb	H	T2_b Administrator	Object Administrator
		T2_b Level 2	Object Operator Level 2
		T2_b Level 1	Object Operator Level 1
		T2_b Guest	Object Guest
T2 SGc	I, J	T2_c Administrator	Object Administrator
		T2_c Level 2	Object Operator Level 2
		T2_c Level 1	Object Operator Level 1
		T2_c Guest	Object Guest

Table 20 Example User Account Mappings for Multiple Tenants

User Account	User Groups	Node Access	Notes
User L	NNMi Level 2 Operators	none	This user has operator level 2 access to the nodes in the pink oval (solid line), which groups all nodes in Tenant 1.
	T1 Level 2	A, B, C, D, E	
User M	NNMi Level 1 Operators	none	This user has operator level 1 access to the nodes in the orange oval (dashed line), which groups a subset of the nodes in Tenant 2.
	T2_a Level 1	F, G	
	T2_b Level 1	H	
User N	NNMi Level 2 Operators	none	This user has operator level 2 access to the nodes in the green oval (dotted line), which groups a subset of the nodes in Tenant 2.
	T2_b Level 2	H	
	T2_c Level 2	I, J	

NNMi Security and Multi-Tenancy Configuration



Any number of static Network Address Translation (NAT) instances can be monitored by one NNMi management server, as long as each instance is configured with a unique tenant. See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193, and the NNMi help, for more information.

NNMi security and multi-tenancy configuration applies to the entire NNMi database. Any NNMi administrator can view and configure operator access to all objects for all tenants.

After an NNMi administrator has defined at least one custom security group, the **Security Group** field is visible on all **Node** forms and as a column in the **Nodes** and **Nodes (All Attributes)** inventory views.

After an NNMi administrator has defined at least one custom tenant, the **Tenant** field is visible on all **Node** forms and as a column in the **Nodes** and **Nodes (All Attributes)** inventory views.

Node groups

To create a node group that aligns with part of the security or multi-tenancy configuration, specify a node group additional filter based on security group UUID, security group name, tenant UUID, or tenant name. Use these node groups to configure per-security group or per-tenant polling cycles for monitoring and incident lifecycle transition actions.

Best practice

Because security group and tenant names can change, specify the security group or tenant UUID in additional filters. This information is available on the configuration forms and in the `nnmsecurity.ovpl` command output.

User groups: NNMi console access

The user account mapping to one of the predefined NNMi user groups sets the NNMi role and the visibility of menu items in the NNMi console. It is recommended to grant each user account the NNMi role that matches the highest object access privilege for that user's topology objects.



The exception to this recommendation is at the administration level because NNMi administrators can access all topology objects. To configure an NNMi console user as an administrator of only some nodes in the NNMi topology, assign that user to the NNMi Level 2 Operators or NNMi Level 1 Operators user group (note that Level 1 Operators have less access privileges than Level 2 Operators). Also assign that user to a custom user group mapped with the Object Administrator object access privilege to a security group containing a subset of the nodes in the topology.

User groups: directory service

If you are storing user group membership in the NNMi database, all object access configuration occurs in the NNMi configuration areas through user groups, user account mappings, security groups, and security group mappings.

If you are storing user group membership in a directory service, object access configuration is shared between NNMi configuration (security groups and security group mappings) and the directory service content (user group membership). Do not create user accounts or user account mappings in the NNMi database. For each applicable group in the directory service, create one or more user groups in the NNMi database. In NNMi, set the **Directory Service Name** field of each user group definition to the distinguished name of that group in the directory service.

For more information, see [Integrating NNMi with a Directory Service through LDAP](#) on page 161.

Configuration Tools

NNMi provides several tools for configuring multi-tenancy and security.

Security Wizard

The **Security Wizard** in the NNMi console is useful for visualizing the security configuration. It is the easiest way to assign nodes to security groups within the NNMi console. The **View Summary of Changes** page presents a list of unsaved changes from the current wizard session. It also identifies potential problems with the security configuration.



The **Security Wizard** is for NNMi security configuration only. It does not include tenant information.

For information about using the **Security Wizard**, click the NNMi help links within the wizard.

NNMi console forms

The forms for individual security and multi-tenancy objects in the NNMi console are useful for concentrating on one aspect of the configuration at a time. For information about using these forms, see the NNMi help for each form.

The **Tenants** view contains NNMi multi-tenancy configuration information. This view is available under **Discovery** in the **Configuration** workspace. Each **Tenant** form describes one NNMi tenant and shows the nodes currently assigned to that tenant. The node assignment information is read-only.

To change the tenant or security group assignment for a node, use the **Node** form or the `nnmsecurity.ovpl` command.

The following NNMi console views are available under **Security** in the **Configuration** workspace. These views contain NNMi security configuration information:

- **User Accounts**
 - Each **User Account** form describes one NNMi user and shows the user groups to which that user belongs. The membership information is read-only.
 - If you are storing user group membership in a directory service, user accounts are not visible in the NNMi console.
- **User Groups**

Each **User Group** form describes one NNMi user group and shows the user accounts and security groups mapped to the user group. The mapping information is read-only.
- **User Account Mappings**
 - Each **User Account Mapping** form shows one user account-to-user group association.
 - Changes to user account mappings do not affect the current NNMi console users. These users receive any changes the next time they log on to the NNMi console.
 - If you are storing user group membership in a directory service, user account mappings are not visible in the NNMi console.
- **Security Groups**

Each **Security Group** form describes one NNMi security group and shows the nodes currently assigned to that security group. The node assignment information is read-only.

- **Security Group Mappings**

- Each **Security Group Mapping** form shows one user group-to-security group association.
- After initial configuration, the object access privilege associated with a security group mapping is read-only. To change the object access privilege for a security group mapping, delete that mapping and recreate it.

Command line

The `nmmsecurity.ovpl` command-line interface is useful for automation and bulk operations. The tool also provides reports of potential problems with the security configuration.

Many of the `nmmsecurity.ovpl` options support loading input data from comma-separated values (CSV) files. You can maintain configuration data in a file or system that can generate CSV output for consumption by the `nmmsecurity.ovpl` command. The command can also accept UUIDs generated outside of NNMI.

Best practice

Because security group and tenant names do not need to be unique, specify the security group or tenant UUID as input to the `nmmsecurity.ovpl` command.

The following example script uses the `nmmsecurity.ovpl` command to create the security configuration for two user accounts and five nodes.

```
#!/bin/sh
# create two users
nmmsecurity.ovpl -createUserAccount user1 -password password -role level1
nmmsecurity.ovpl -createUserAccount user2 -password password -role level2

# create two user groups
nmmsecurity.ovpl -createUserGroup local1
nmmsecurity.ovpl -createUserGroup local2

# assign the user accounts to the new user groups
nmmsecurity.ovpl -assignUserToGroup -user user1 -userGroup local1
nmmsecurity.ovpl -assignUserToGroup -user user2 -userGroup local2

# create two security groups
nmmsecurity.ovpl -createSecurityGroup secgroup1
nmmsecurity.ovpl -createSecurityGroup secgroup2

# assign the new user groups to the new security groups
nmmsecurity.ovpl -assignUserGroupToSecurityGroup -userGroup local1 \
  -securityGroup secgroup1 -role level1
nmmsecurity.ovpl -assignUserGroupToSecurityGroup -userGroup local2 \
  -securityGroup secgroup2 -role level2

# assign nodes to security groups
nmmsecurity.ovpl -assignNodeToSecurityGroup -node mplspe01 -securityGroup secgroup1
nmmsecurity.ovpl -assignNodeToSecurityGroup -node vwan_router-1 -securityGroup secgroup1
nmmsecurity.ovpl -assignNodeToSecurityGroup -node vwan_router-2 -securityGroup secgroup1
nmmsecurity.ovpl -assignNodeToSecurityGroup -node data_center_1 -securityGroup secgroup2
nmmsecurity.ovpl -assignNodeToSecurityGroup -node mplspe03 -securityGroup secgroup2
```

Configuring Tenants



Any number of static Network Address Translation (NAT) instances can be monitored by one NNMI management server, as long as each instance is configured with a unique tenant. See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193, and the NNMI help, for more information.

NNMi provides the following ways to configure multi-tenancy:

- The **Tenant** form in the NNMi console is useful for working with individual tenants.
- The `nnmsecurity.ovpl` command-line interface is useful for automation and bulk operations. The tool also provides reports of potential problems with the tenant configuration.

The process of defining and configuring NNMi multi-tenancy to assign each NNMi topology object to a tenant (organization) is a cyclical process. This high-level procedure describes one approach to configuring NNMi multi-tenancy.

Note the following about configuring NNMi multi-tenancy:

- The security group that NNMi assigns to a discovered node is set by the value of the Initial Discovery Security Group for the tenant associated with that node.
- When you use the NNMi security model without also configuring NNMi tenants, all nodes are assigned to the Default Tenant.
- When you seed a node for NNMi discovery, you can specify the tenant to which that node belongs. When NNMi discovers a node through an auto-discovery rule, NNMi assigns that node to the Default Tenant. After discovery, you can change the tenant assignment for the node.

One high-level approach to planning and configuring NNMi multi-tenancy is as follows:

- 1 Analyze your customer requirements to determine how many tenants are required in the NNMi environment.

It is recommended that tenants be used only when managing multiple separate networks with a single NNMi management server.

- 2 Analyze the managed network topology to determine which nodes belong to each tenant.
- 3 Analyze the topology of each tenant to determine the groups of nodes to which NNMi users need access.
- 4 Remove the default associations between the predefined NNMi user groups and the Default Security Group and the Unresolved Incidents security group.

Doing this step assures that users do not inadvertently obtain access to nodes they should not be managing. At this point, only NNMi administrators can access objects in the NNMi topology.

- 5 Configure the identified tenants.
 - a Create the identified security groups.
 - b Create the identified tenants.

For each tenant, set the Initial Discovery Security Group to either the Default Security Group or a tenant-specific security group with restricted access. This approach ensures that new nodes for the tenant are not generally visible until the NNMi administrator configures access.

- 6 Prepare for discovery by assigning tenants to seeds.

After discovering a group of nodes, you can change the value of the Initial Discovery Security Group. Using this approach limits the manual re-assignment of nodes to security groups.



- 7 After discovery completes, do the following:
 - Verify the tenant for each node and make changes as necessary.
 - Verify the security group for each node and make changes as necessary.
- 8 Continue with [step 4](#) on page 222.

Configuring Security Groups



If you plan to integrate NNMi with a directory service for consolidating the storage of user names, passwords, and, optionally, NNMi user group assignments, complete that configuration before configuring NNMi security.

NNMi provides the following ways to configure security:

- The **Security Wizard** in the NNMi console is useful for visualizing the security configuration. The **View Summary of Changes** page presents a list of unsaved changes from the current wizard session. It also identifies potential problems with the security configuration.
- The forms in the NNMi console for individual security objects are useful for concentrating on one aspect of the security configuration at a time.
- The `nmmsecurity.ovpl` command-line interface is useful for automation and bulk operations. The tool also provides reports of potential problems with the security configuration.

The process of defining and configuring NNMi security to limit users' access to objects in the NNMi topology is a cyclical process. This high-level procedure describes one approach to configuring NNMi security.



This example moves from security groups to user accounts. For examples of configuring NNMi security from user accounts to security groups, search for "Configure Security Example" in the NNMi help.

Note the following about configuring NNMi security:

- The security group that NNMi assigns to a discovered node is set by the value of the Initial Discovery Security Group for the tenant associated with that node.
- When you use the NNMi security model without also configuring NNMi tenants, all nodes are assigned to the Default Tenant.

One high-level approach to planning and configuring NNMi security is as follows:

- 1 Analyze the managed network topology to determine the groups of nodes to which NNMi users need access.
- 2 Remove the default associations between the predefined NNMi user groups and the Default Security Group and the Unresolved Incidents security group.

Doing this step assures that users do not inadvertently obtain access to nodes they should not be managing. At this point, only NNMi administrators can access objects in the NNMi topology.
- 3 Configure a security group for each subset of nodes. Remember that a given node can belong to only one security group.
 - a Create the security groups.
 - b Assign the appropriate nodes to each security group.

- 4 Configure custom user groups.
 - a For each security group, configure a user group for each level of NNMi user access.
 - If you are storing user group membership in the NNMi database, no users are mapped to these user groups yet.
 - If you are storing user group membership in a directory service, set the Directory Service Name field for each user group to the distinguished name of that group in the directory service.
 - b Map each custom user group to the correct security group. Set the appropriate object access privilege for each mapping.
- 5 Configure user accounts.
 - If you are storing user group membership in the NNMi database, do the following:
 - Create a user account object for each user who can access the NNMi console. (The process of configuring user accounts depends on whether you are using a directory service for NNMi console logon.)
 - Map each user account to one of the predefined NNMi user groups (for access to the NNMi console).
 - Map each user account to one or more custom NNMi user groups (for access to topology objects).
 - If you are storing user group membership in a directory service, verify that each user belongs to one of the predefined NNMi user groups and one or more custom user groups.
- 6 Verify the configuration as described in [Verifying the Configuration](#) on page 222.
- 7 Maintain the security configuration.
 - Watch for nodes added to the Default Security Group, and move these nodes to the correct security groups.
 - Add new NNMi console users to the correct user groups.

Verifying the Configuration

To verify that the security configuration is correct, verify each aspect of the configuration separately. This section describes some approaches to verifying the configuration. Other approaches are possible.



NNMi provides reports of possible security configuration errors. Access these reports with **Tools > Security Reports** in the NNMi console and with the `-displayConfigReport` option to the `nmmsecurity.ovpl` command.

Verify security group-to-node assignments

One approach to verifying that each node is assigned to the correct security group is to sort the **Nodes** or **Nodes (All Attributes)** inventory view by security group, and then examine the groupings.

Another approach is to use the `-listNodesInSecurityGroup` option to the `nmmsecurity.ovpl` command.

Verify user group-to-security group assignments

One approach to verifying which user groups are mapped to each security group is to sort the **Security Group Mappings** view by user group or security group, and then examine the groupings. Also verify the object access privilege for each mapping.

Alternatively, on the **Map User Groups and Security Groups** page of the **Security Wizard**, select one user group or security group at a time to see the current mappings for that object.

Another approach is to use the `-listUserGroupsForSecurityGroup` option to the `nmsecurity.ovpl` command.

Verify that each user has NNMi console access

For NNMi console access, ensure that each user is assigned to one of the predefined NNMi user groups (listed from highest to lowest):

- NNMi Administrators
- NNMi Level 2 Operators
- NNMi Level 1 Operators
- NNMi Guest Users

All other user group assignments provide access to objects in the NNMi database.



The NNMi Global Operators Users Group provides access to topology objects only. Unless a `globalops` user is also associated with a User Group with NNMi Console access (such as `level2`, `level1`, or `guest`), that user will not be able to access the NNMi Console.

Users without NNMi console access are listed on the **View Summary of Changes** page of the **Security Wizard**. The **Tools > Security Reports** menu item and the `-displayConfigReport usersWithoutRoles` option to the `nmsecurity.ovpl` command also provide this information.



Each **Tools** and **Action** menu item provided in the NNMi Console is associated with a default NNMi role. (To determine the default NNMi Role assigned to each Action menu item, see *Actions Provided by NNMi* in the NNMi help.) If you change the setting for a menu item provided by NNMi to a role that is a lower level role than the default NNMi role assigned to the menu item, NNMi ignores that change. Any User Group with the lower level role than the default NNMi role cannot access the menu item.

Verify user-to-user group assignments

One approach to verifying user group membership is to sort the **User Account Mappings** view by user account or user group, and then examine the groupings.

Alternatively, on the **Map User Accounts and User Groups** page of the **Security Wizard**, select one user account or user group at a time to see the current mappings for that object.

Another approach is to use the `-listUserGroups` and `-listUserGroupMembers` options to the `nmsecurity.ovpl` command.

Verify tenant-to-node assignments

One approach to verifying that each node is assigned to the correct tenant is to sort the **Nodes** or **Nodes (All Attributes)** inventory view by tenant, and then examine the groupings.

Verify current user settings

To verify the NNMi console access for the currently logged-on user, click **Help > System Information**. The **User Information** section on the **Product** tab lists the following information for the current NNMi session:

- User name as defined for the user account in the NNMi database or the accessed directory service.

- NNMi role, which corresponds to the most privileged of the predefined NNMi user groups (NNMi Administrators, NNMi Level 2 Operators, NNMi Level 1 Operators, and NNMi Guest Users) to which the user is mapped. This mapping determines which actions are available within the NNMi console.
- User groups mapped to this user name. This list includes predefined NNMi user group that sets the NNMi role and any other user groups that provide access to objects in the NNMi database.

Exporting the NNMi Security and Multi-Tenancy Configuration

Table 21 describes the configuration areas (available with `nnmconfigexport.ovpl -c`) for exporting the NNMi security and multi-tenancy configuration. These export areas are beneficial for maintaining the configuration across multiple NNMi management servers, especially in a Global Network Management environment.

Table 21 NNMi Security and Multi-Tenancy Configuration Export Areas

Configuration Area	Description
account	Exports user accounts, user groups, and user account-to-user group mappings. Useful for sharing user definitions across multiple NNMi databases.
security	Exports tenants and security groups. Useful for sharing security definitions across multiple NNMi databases. Importing this information creates new objects and updates existing objects but does not delete objects not included in the current export. Therefore, this option is safe to use with an NNMi database containing locally-defined objects.
securitymappings	Exports user group-to-security group mappings. For a complete export of the security and multi-tenancy configuration, perform a concurrent export of the account, security, and securitymappings configuration areas.

NNMi Security, Multi-Tenancy, and Global Network Management

In a Global Network Management (GNM) environment, a node's tenant is set on the NNMi management server that manages that node. The tenant UUID for a given node is the same on each global and regional manager in the GNM environment.

A node's security group is set on each NNMi management server whose topology contains that node. Thus, user access to objects in the topology is configured separately on each NNMi management server in the GNM environment. The global and regional managers might use the same or different security group definitions.

If you want user access to be similar on the global manager and regional managers, you can employ some configuration tricks, but you probably cannot completely avoid custom configuration on each NNMi management server.



Each group of dynamic Network Address Translation (NAT) or dynamic Port Address Translation (PAT) requires an NNMi regional manager, in addition to a tenant that is unique within the entire NNMi global network management configuration. See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193. See also the NNMi help.

Best practice Define all tenants and security groups on the global manager. Use `nmmconfigexport.ovpl -c security` to export the tenant and security group definitions. On each regional manager, use `nmmconfigimport.ovpl` to import the tenant and security group definitions. Alternatively, you can use the `nmmsecurity.ovpl` command to create tenants and security group with the same UUID as on another NNMi management server. Following this recommendation ensures that each tenant and security group has the same UUID within the GNM environment.



This best practice becomes a *required* part of the configuration if users will be launching NPS reports from the global manager.



Tenant UUIDs must be unique, but tenant names can be reused. NNMi considers two tenants with the same name and different UUIDs to be two distinct tenants with no shared configuration.

Best practice If you are setting up one regional manager per organization, all nodes on a regional manager can be in a single tenant. However, configure a unique tenant on each regional manager to ensure separation of the topology data on the global manager.

Incidents forwarded from a regional manager to a global manager might include some additional custom incident attributes (CIAs) to convey security and tenant information.

If the incident's source object belongs to a tenant other than the Default Tenant, the forwarded incident contains the following CIAs:

- `cia.tenant.name`
- `cia.tenant.uuid`

If the incident's source object belongs to a security group other than the Default Security Group, the forwarded incident contains the following CIAs:

- `cia.securityGroup.name`
- `cia.securityGroup.uuid`

This section contains the following topics:

- [Initial GNM Configuration](#) on page 226
- [GNM Maintenance](#) on page 228

Initial GNM Configuration

After GNM is first configured, the regional manager updates the global manager with information about the nodes in the regional topology (according to the GNM configuration).

For GNM environments with custom security groups and the Default Tenant, on the global manager, all nodes managed remotely are added to the global manager topology with the following configuration:

- Default Tenant
- The security group that is set as the Initial Discovery Security Group for the Default Tenant.

**Topology
synchronization with
the Default Tenant
only**

Topology synchronization with custom tenants

For GNM environments with custom security groups and custom tenants, on the global manager, all nodes managed remotely are added to the global manager topology with the UUID of the tenant assigned to the node. If that tenant UUID does not exist on the global manager, the GNM processes create that tenant in the NNMi configuration of the global manager as follows:

- The tenant UUID is the same value as on the regional manager.
- The tenant name is the same value as on the regional manager.
- The value of the Initial Discovery Security Group is set to the security group with the same name as the tenant. (NNMi creates this security group if it does not already exist on the global manager.)

As the node is added to the topology on the global manager, it is assigned to the Initial Discovery Security Group for the tenant UUID as configured on the global manager. That is, the security group association on the global manager is independent of the security group association on the regional manager.

Best practice

Suggestions for simplifying security configuration on the global manager include:

- Maintain a spreadsheet or other record of the nodes managed by each regional manager. For each node, note the expected security group on the regional manager and that on the global manager. After GNM configuration completes, use the `nmsecurity.ovpl` command to verify and update the security group assignments.
- If the GNM environment will include multiple regional managers updating a single global manager, enable the GNM configuration from one regional manager at a time to the global manager.

If appropriate, you can change the value of the Initial Discovery Security Group of the Default Tenant (or a custom tenant) before adding each regional manager to the GNM configuration. Note that this approach can have mixed results if new nodes are being added to the topology on the previously configured regional managers.

- Before enabling GNM, on the global manager, set the Initial Discovery Security Group of each tenant used on the regional manager to be a private security group that operators cannot access. An administrator on the global manager then needs to explicitly move the nodes to the appropriate security groups for other NNMi console operators.

GNM Maintenance

Table 22 describes how changes to a node's tenant or security group assignment on a regional manager affect the global manager.

Table 22 Global Manager Impact of Configuration Changes on a Regional Manager

Action	Effect
On the regional manager, assign a node to a different tenant.	The node on the global manager is changed to be assigned to the different tenant. If this tenant UUID does not exist on the global manager, it is created.
On the regional manager, assign a node to a different security group.	No change on the global manager. The NNMi administrator can choose to replicate the change manually.
On the regional manager, change the configuration (name, description, or Initial Discovery Security Group) of a tenant.	No change on the global manager. The NNMi administrator can choose to replicate the change manually.
On the regional manager, change the configuration (name or description) of a security group.	No change on the global manager. The NNMi administrator can choose to replicate the change manually.

Including Select Interfaces in NPS Reports

The Network Performance Server (NPS) is the database server installed with the NNM iSPI Performance for Metrics software.

By default, all components of a node are in the same security group as the node. For individual interfaces, you can override this default behavior and assign an interface to a different security group. The purpose of this override is to generate tenant-specific reports that include the appropriate interfaces for that tenant (customer) on shared devices. In this way, each customer can see the interface information for their interfaces but cannot see the other interfaces on the device.


- ▶ The security group override only affects NPS reports. It has no impact on what users can see and do in the NNMi console.

To change the security group assignment for an interface, on the **Custom Attributes** tab of an **Interface** form or with the `nmloadattributes.ovpl` command, add the `InterfaceSecurityGroupOverride` custom attribute to that interface. Set the value of this custom attribute to the UUID of the security group. For example:

```
InterfaceSecurityGroupOverride=0826c95c-5ec8-4b8c-8998-301e0cf3c1c2
```

- ▶ An interface can belong to only one security group at a time. Setting the `InterfaceSecurityGroupOverride` custom attribute on an interface breaks the association between that interface and the security group to which its node belongs.

Global Network Management



This chapter contains the following topics:

- Global Network Management Benefits
- Is Global Network Management a Good Tool for Managing my Network?
- Practical Global Network Management Examples
- Review the Requirements
- Initial Preparation
- Configuring Single Sign-On for Global Network Management
- Configuring Forwarding Filters on the Regional Managers
- Connecting a Global Manager with a Regional Manager
- Determining the Connection States from global1 to regional1 and regional2
- Reviewing global1 Inventory
- Disconnecting Communication between global1 and regional1
- Additional Information
- Configuring Application Failover for Global Network Management
- Troubleshooting Tips for Global Network Management
- Global Network Management and NNM iSPIs or Third-Party Integrations
- Global Network Management and Address Translation Protocols

Global Network Management Benefits

Suppose you have HP Network Node Manager i Software (NNMi) deployed on multiple NNMi management servers in several geographic locations. You have each NNMi management server discovering and monitoring the network to meet your discovery and monitoring needs. Using these existing NNMi management servers and

configurations, you can designate specific NNMi management servers as global managers to display combined node object data without additional discovery or monitoring configuration changes.

The NNMi global network management feature enables multiple NNMi management servers to work together while managing different geographic areas of the network. You designate specific NNMi management servers as global managers to display combined node object data from 2 or more regional managers.

The NNMi global network management feature offers the following benefits:

- A central big-picture view of your corporate-wide network from the global manager.
- Easy to set up:
 - Each regional manager administrator specifies all node object data or a specific node group for participation at the global manager level.
 - Each global manager administrator specifies which regional managers are allowed to contribute information.
- Generates and manages incidents independently on each server (generated within the context of topology available on each server).

See *NNMi's Global Network Management Feature* in the NNMi help for additional details.



Each group of dynamic Network Address Translation (NAT) or dynamic Port Address Translation (PAT) or dynamic Network Address and Port Translation (NAPT) requires an NNMi regional manager, in addition to a tenant that is unique within the entire NNMi global network management configuration. See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193. See also the NNMi help.

Is Global Network Management a Good Tool for Managing my Network?

Ask the following questions to determine if NNMi's global network management feature can help you better manage your network.

Do I Need Continuous Multi-Site Network Monitoring?

Does your information technology group manage network equipment located at multiple sites on a 24 by 7 basis? If so, your group can use NNMi's global network management feature to observe combined topology and incident views.

Can my Critical Devices be Visible?

From one NNMi management server, can I view device status and incidents for critical devices located at multiple locations? Yes. You configure forwarding filters on the regional managers. This enables you to select the node object data you want regional managers to send to global managers. For example, you can set up forwarding filters on the regional managers so that they only forward information about critical devices to the global manager.

Licensing Considerations

For information about obtaining and installing NNMi license keys, see [Licensing NNMi](#) on page 121.

Do I need an NNMi Advanced license on both the global and regional managers? You must purchase and install an NNMi Advanced license on the NNMi management server you plan to use as a global manager. NNMi management servers do not need an NNMi Advanced license to function as regional managers.

I currently have adequate NNMi licenses for single geographies. Can I use the global network management feature and limit the new licenses I need on the global manager? Yes. If your information technology group needs to monitor critical equipment located at multiple sites, you can configure a forwarding filter on the regional manager to make sure you only forward information about critical devices to the global manager. This enables you to wisely use your NNMi investment and control your use of the license capacity you have on the global manager.

I increased the NNMi licenses for the regional managers such that the total number of licensed nodes is larger than the NNMi Advanced licenses on the global manager. Now the global manager does not have a complete inventory of all nodes in all regions. After I purchase and install enough licenses for the global manager, how can I get the global manager to synchronize with all of the regional managers in order for it to find and create the nodes it formerly skipped due to insufficient licenses? You must purchase and install enough NNMi Advanced licenses on the global manager to meet or exceed the number of total licenses you have installed on the regional managers. After you have enough licenses installed, do one of the following:

- Wait for all of the configured rediscovery intervals on all of the regional managers to elapse so that all of the nodes in all of the regions are rediscovered. After the regional managers rediscover all of the nodes in all of the regions, the regional managers send this rediscovered node information to the global manager. The global manager receives this node information and creates global nodes for each node in each region.
- Run the `nnmnoderediscover.ovpl -all` script on each regional manager.



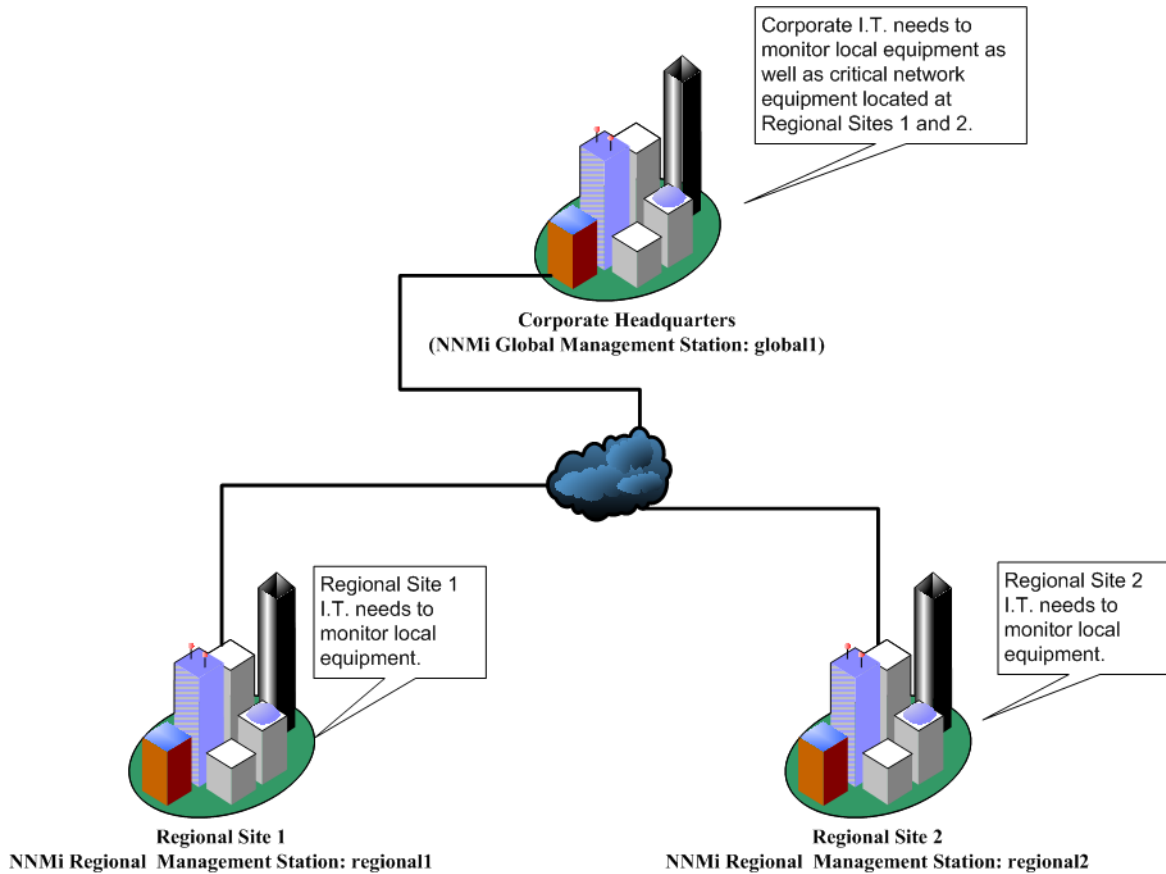
The second option causes both a lot of traffic on your network and consumes a lot of NNMi resources from the entire set of NNMi managers. This option is not as resource intensive as the initial NNMi discovery, but it is similar to doing the first discovery. The best approach is to space the running of the script for each region by some amount of time or by waiting for the current regional manager's workload to drop to normal before starting the next regional manager's rediscovery.

Practical Global Network Management Examples

See [Figure 21](#) on page 232. Suppose your company has 2 operating sites in different geographic locations. Your company's headquarters is located in a third geographic area. There are NNMi management servers functioning at all 3 locations.

From a network perspective, information technologists located at corporate headquarters need to monitor local network equipment as well as critical network equipment located at both regional sites 1 and 2. Information technologists from both regional sites 1 and 2 need to monitor the local critical network equipment located at their sites.

Figure 21 Example Network.



Review the Requirements

Suppose the NNMi management servers at corporate headquarters, regional site 1, and regional site 2 manage several routers and switches located at their individual sites. For this example, consider the NNMi management servers as `global1`, `regional1` and `regional2` respectively. Suppose you configured these NNMi management servers to discover and monitor critical switches and routers located at their own locations. There is no need to reconfigure discovery for NNMi management servers at any of these sites to use the global network management feature.

- ▶ During global network management configuration, you might be tempted to use the `nnmbackup.ovpl` script to back up one NNMi management server, use the `nnmrestore.ovpl` script to restore this backup to a second NNMi management server, then connect both of these NNMi management servers to a regional NNMi management server. Do not do this. Placing the backup data from one NNMi management server onto a second NNMi management server means that both servers have the same database UUID. After you restore NNMi on the second NNMi management server, you would need to uninstall NNMi from the original NNMi management server.

The information technology group at your corporate site wants to monitor critical equipment located at regional sites 1 and 2, but they do not want to manage every device. The following table summarizes the monitoring needs:

Table 23 Network Requirements for Global Network Management

Site	NNMi Management Servers	Critical Switches	Regional Equipment to Manage
Corporate Headquarters	global1	15 Model 3500yl HP Procurve Switches	All model 3500yl HP Procurve Switches from each regional site
Regional Site 1	regional1	15 Model 3500yl HP Procurve Switches	not applicable
Regional Site 2	regional2	15 Model 3500yl HP Procurve Switches	not applicable

To summarize, you have an NNMi management server, `global1`, monitoring the corporate headquarters. You have NNMi management servers, `regional1` and `regional2`, monitoring each of the regional sites. From corporate headquarters, you must view incidents and device information for the Model 3500yl Procurve switches located at regional sites 1 and 2. Suppose that, for this example, `regional1` and `regional2` both manage several common switches located at regional site 1.

Regional Manager and Global Manager Connections

When you configure global network management connections, consider the following information:

- Use the same NNMi version and patch level on the global manager and all regional managers. Configuring global network management using different NNMi versions is not supported.
- NNMi enables you to configure more than one global manager to communicate with a regional manager. For example, if you need a second global manager, `global2`, to communicate with `regional1`, NNMi enables you to configure both `global1` and `global2` to communicate with `regional1`. For more information see the *HP Network Node Manager i Software System and Device Support Matrix*.
- Global network management works with one connection layer. For example, the examples in this chapter discuss one connection layer: `global1` communicating with `regional1` and `global1` communicating with `regional2`. Do not configure NNMi for multiple connection levels. For example, do not configure `global1` to communicate with `regional1`, then configure `regional1` to communicate with `regional2`. The global network management feature is not designed for this three layer configuration.
- Do not configure two NNMi management servers to communicate both ways with each other. For example, do not configure `global1` to communicate with `regional1`, then configure `regional1` to communicate with `global1`.

Initial Preparation

Port Availability: Configuring the Firewall

For the global network management feature to function properly, you must verify that certain well-known ports are open for TCP access from `global1` to `regional1`, and `regional2`. The NNMi installation script sets ports 80 and 443 as defaults; however, you can change these values during installation.



In the example discussed in this section, `global1` establishes TCP access to `regional1` and `regional2`. Firewalls are usually configured based on the server initiating the connection. After `global1` establishes the connection to `regional1` and `regional2`, traffic flows in both directions.

Edit the following file to see the current values or to make port configuration changes:

- *Windows*: %NNM_CONF%\nmm\props\nms-local.properties
- *UNIX*: \$NNM_CONF/nmm/props/nms-local.properties

The following table shows the well-know ports that need to be accessible:

Table 24 Required Accessible Sockets

Security	Parameter	TCP Port
non-SSL	nmsas.server.port.web.http	80
	nmsas.server.port.hq	4457
SSL	nmsas.server.port.web.https	443
	nmsas.server.port.hq.ssl	4459

See [NNMi 9.20 and Well-Known Ports](#) for more information.

Configuring Self-Signed Certificates

If you plan to use the global network management feature with SSL (Secure Sockets Layer) between `global1` and the two regional NNMi management servers (`regional1` and `regional2`), you must do some additional work. During NNMi installation, the NNMi installation script creates a self-signed certificate on the NNMi management server so it can identify itself to other entities. You must configure the NNMi management servers you plan to use with the global network management feature with the correct certificates. Complete the steps shown in [Configuring the Global Network Management Feature to use Self-Signed Certificates](#) on page 135.

Configuring Global Network Management for Application Failover

During NNMi installation, the NNMi installation script creates a self-signed certificate on the NNMi management server so it can identify itself to other entities. If you plan to use the application failover along with the global network management feature, you must do some additional configuration. Complete the steps shown in [Configuring Global Network Management with Application Failover to use Self-Signed Certificates](#) on page 137.

NNMi Management Server Sizing Considerations

This example assumes you plan to use existing NNMi management servers in a global network management configuration. The global network management feature is different than the distributed solution used in earlier NNM products. The global network management feature avoids polling nodes being managed by regional systems, so you do not need to be as concerned about network bandwidth and computer resources.

Review the *HP Network Node Manager i Software Interactive Installation Guide*, the *NNMi Release Notes*, and the *NNMi System and Device Support Matrix*, for specific information about the size of server you need for NNMi.

Synchronizing System Clocks

It is important for you to synchronize the NNMi management server clocks for `global1`, `regional1`, and `regional1` before you connect these servers in a global network management configuration. All NNMi management servers in your network environment that participate in global network management (global managers and regional managers) or single sign-on (SSO) must have their internal time clocks synchronized in universal time. Use a Time Synchronization program, for example, the UNIX (HP-UX / Linux / Solaris) tool Network Time Protocol Daemon (NTPD) or one of the available Window operating system tools. See *Clock Synchronization Issues* or *Troubleshoot Global Network Management* in the NNMi help and [Clock Synchronization](#) on page 262 for more information.



NNMi displays a warning message at the bottom of the NNMi Console if there is a connection problem with a regional manager, such as a server clock synchronization problem.

Using the Application Failover Feature with Self-Signed Certificates in Global Network Management

If you plan to use the global network management feature using self-signed certificates in an application failover configuration, you must complete some additional steps. See [Configuring Global Network Management with Application Failover to use Self-Signed Certificates](#) on page 137.

Using Self-Signed Certificates in Global Network Management

If you plan to use the global network management feature using self-signed certificates, you must complete some additional steps. See [Configuring the Global Network Management Feature to use Self-Signed Certificates](#) on page 135.

Using a Certificate Authority in Global Network Management

If you plan to use the global network management feature using a Certificate Authority, you must complete some additional steps. See [Configuring the Global Network Management Feature to use a Certificate Authority](#) on page 136.

List the Critical Equipment you Want to Monitor

Make a list of the equipment managed by `regional1` and `regional2` that you want to monitor from `global1`. You will use this information in a forwarding filter (to be discussed later). You must carefully consider the possible outcomes of limiting the information forwarded to `global1` from `regional1` and `regional2`. Below are some things to consider during your planning:

- Be careful not to exclude too many devices, as `global1` needs a complete topology from `regional1` and `regional2` to do a complete analysis to generate accurate incidents.
- Excluding non-critical devices helps you to reduce license costs on `global1`.
- Excluding non-critical devices helps you to improve the solution's overall scalability, and reduce the network traffic required by NNMi.

Review the Global and Regional Managers' Management Domains

NNMi management servers `global1`, `regional1`, and `regional2` manage their own set of nodes. Later in this example, you will configure `regional1` and `regional2` to forward information about equipment they manage to `global1`.

Use the following procedure to understand the equipment that `global1`, `regional1`, and `regional2` currently monitor. This helps you select the critical equipment you want `regional1` and `regional2` to forward to `global1`.

For this example, complete the following steps to review this information:

- 1 Point your browser to `global1`'s, NNMi console.
- 2 Sign in.
- 3 Click **Inventory** workspace.
- 4 From here you can review the discovered inventory `global1` currently monitors.
- 5 Point your browser to `regional1`'s, NNMi console.
- 6 Sign in.
- 7 Click **Inventory** workspace.
- 8 Review the nodes that `regional1` monitors and make a list of the devices you want to monitor from `global1`.
- 9 Point your browser to `regional2`'s, NNMi console.
- 10 Sign in.
- 11 Click **Inventory** workspace.
- 12 Review the nodes that `regional2` monitors and make a list of the devices you want to monitor from `global1`.

Review NNMi Help Topics

To review all of the help topics related to global network management, complete the following steps:

- 1 From the NNMi help, click **Search**.
- 2 Type "**Global Network Management**" in the Search field.
- 3 Click **Search**.

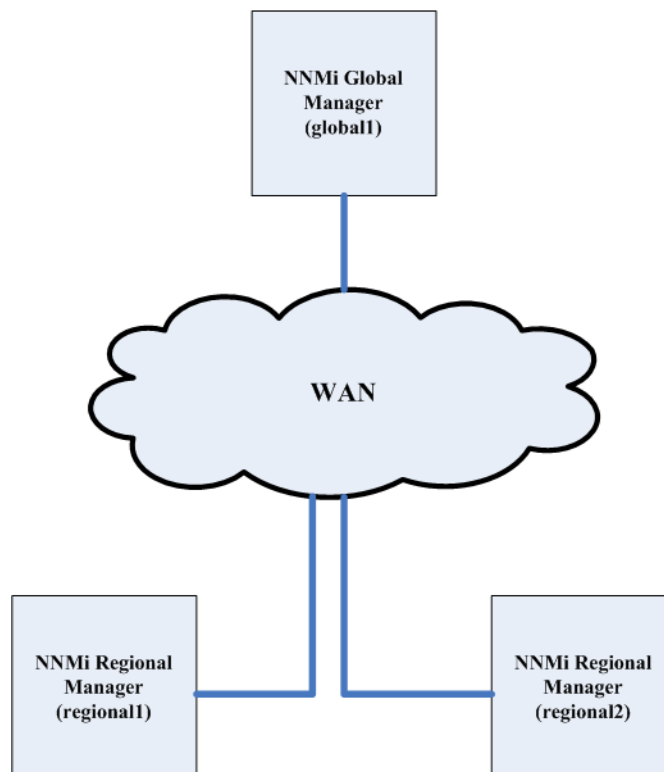
This search results in more than 50 topics related to global network management.

SSO and the Actions Menu

From an NNMi console on a global manager, suppose you select a node managed by a regional manager, then use the **Actions** menu to initiate an action on the selected node. Without having the `initString` and `domain` parameters the same among the NNMi management servers, the session information from the global manager does not get passed on to the new session and the action does not get initiated. To avoid this problem, follow the configuration steps shown in [Configuring Single Sign-On for Global Network Management](#) on page 237.

Configuring Single Sign-On for Global Network Management

You can configure NNMi single sign-on (SSO) to facilitate access to NNMi regional managers from an NNMi global manager. You must complete this step before connecting regional managers from a global manager. See [Using Single Sign-On with NNMi](#) on page 141 for more information.



The SSO feature communicates a user name among NNMi management servers, but not passwords or roles. For example, NNMi associates the same username on one NNMi management server (`global1`) with a different role on other NNMi management servers (`regional1` or `regional2`). Any of these three NNMi management servers could associate a different password with the same username.

If a global and regional manager resides in the same management domain, and you do not copy the *Initialization String* value from the global NNMi management server to the regional NNMi management server as shown in [step 4](#) on page 238, you could have NNMi console access problems. To avoid this, either configure SSO correctly using the following steps, or disable SSO as described in [Disabling SSO](#) on page 146.

To configure SSO to work with the global network management feature, complete the following steps:


- 1 Edit the following file on `global1`, `regional1`, and `regional2`:
 - *Windows*: `%NNM_PROPS%\nms-ui.properties`
 - *UNIX*: `$NNM_PROPS/nms-ui.properties`
- 2 On `global1`, `regional1`, and `regional2`, look for a section in the file that resembles the following:



```
com.hp.nms.ui.sso.isEnabled = false
```

Change this as follows:

```
com.hp.nms.ui.sso.isEnabled = true
```
- 3 Locate the SSO NNMi initialization string for `global1`. Look for a section in the `nms-ui.properties` file that resembles the following:


```
com.hp.nms.ui.sso.initString = Initialization String
```
- 4 Copy the value of *Initialization String* from the `nms-ui.properties` file on `global1` to the `nms-ui.properties` files on `regional1` and `regional2`. All of the servers must use the same value for *Initialization String*. Save your changes.

 NNMi supports copying the *Initialization String* value from the global NNMi management server to the regional NNMi management servers. In this step, you copied the *Initialization String* value from the global manager to the two regional managers. Always copy the *Initialization String* value from the global manager to the regional managers if you want to use SSO with the global network management feature.

 If a global and regional manager resides in the same management domain, and you do not copy the *Initialization String* value from the global NNMi management server to the regional NNMi management server, disable SSO to avoid NNMi console access problems. See [Disabling SSO](#) on page 146 for more information.
- 5 If `global1`, `regional1`, and `regional2` are in different domains, modify the `protectedDomains` content. To do this, look in the `nms-ui.properties` file for a section that resembles the following:


```
com.hp.nms.ui.sso.protectedDomains=group1.mycompany.com
```

Suppose `global1` is in `global1.company1.com`, `regional1` is in `regional1.company2.com` and `regional2` is in `regional2.company3.com`. Modify the `protectedDomains` section of the `nms-ui.properties` file on `global1`, `regional1` and `regional2` as follows:

```
com.hp.nms.ui.sso.protectedDomains=regional1.company1.com,  
regional2.company2.com,regional3.company3.com
```

- 6 Save your changes.
- 7 Run the following command sequence on `global1`, `regional1`, and `regional2`:
 - a `ovstop`
 - b `ovstart`



There are no manual configuration steps to perform to enable single sign-on in an application failover configuration. For example, If you plan to configure single sign-on in an application failover configuration, NNMI replicates the above changes from the active NNMI management server to the standby NNMI management server.

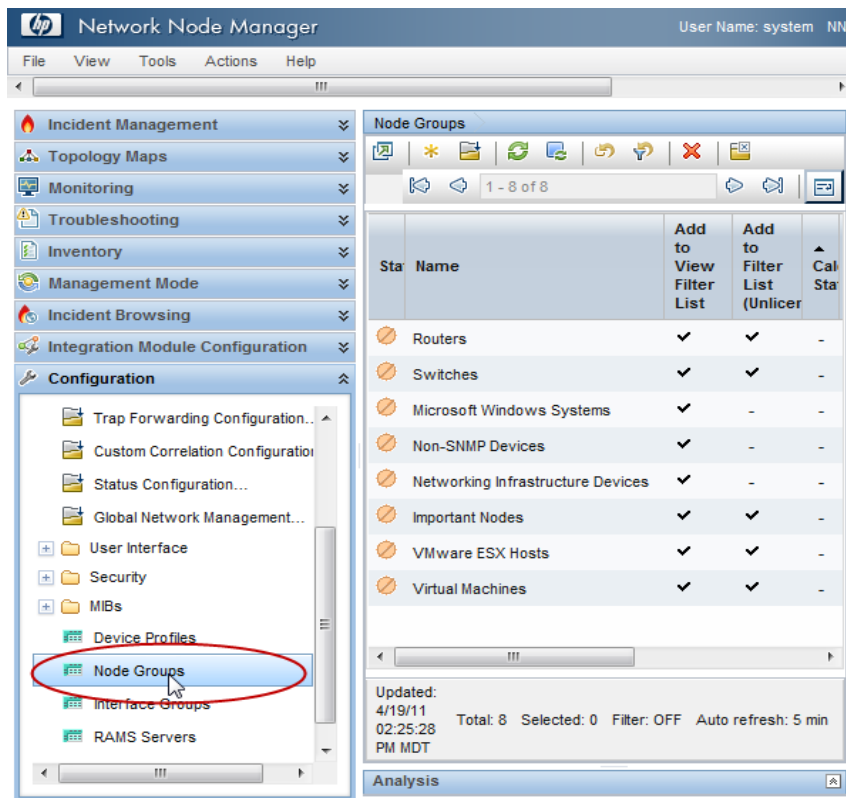
Configuring Forwarding Filters on the Regional Managers

In this example, `global1` communicates with both `regional1` and `regional2`. To control the node object data you want the global manager, `global1`, to receive from regional managers `regional1` and `regional2`, you must configure forwarding filters on both `regional1` and `regional2`.

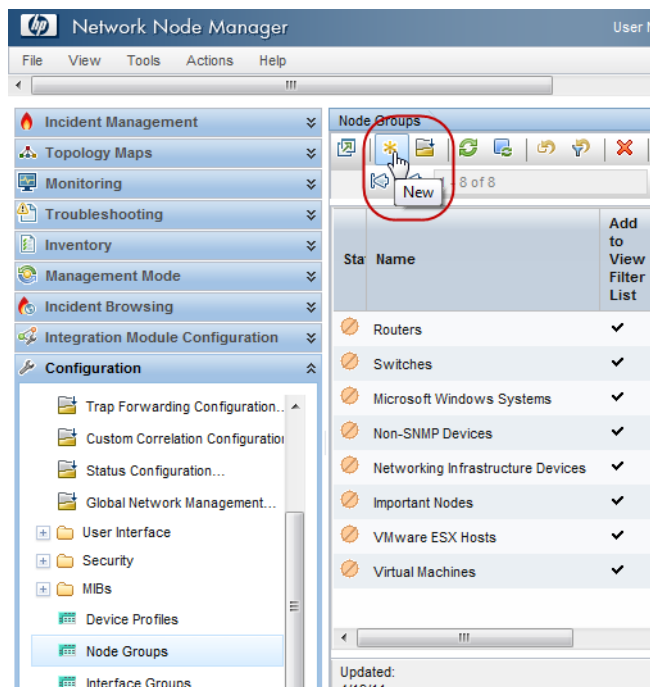
Configuring a Forwarding Filter to Limit Forwarded Nodes

Suppose you want to set up a node group to enable `regional1` to only forward node information for Procurve Model 3500yl switches to `global1`. To create a new node group and set these limits, complete the following steps:

- 1 From `regional1`'s **Configuration** workspace in the NNMI console, click **Node Groups**.




2 Click **New**.




▶ Although this example explains how to create a new node filter, then use it to create a forwarding filter from regional1 and regional2, you can use any of these existing filters to set up forwarding filters from a regional NNMi management server to a global NNMi management server.

▶ You can create a *container* node group that contains no devices or filters of its own; then use this node group to specify child node groups. Using this approach, you can forward node object data to global NNMi management servers using one *container* node group.


- 3 Click the **Device Filters** tab. Type global1 as the filter name and make any notes you need about the filter you are creating in the notes field.



- Click the **New** icon to open a Node Device Filter form.



- Using the pull-down, select the **Switch Router Device Category**, the **Hewlett-Packard Device Vendor**, and the **HP Procurve 3500 Fixed-port Switch Device Family**.
- Using the pull-down, click **Quick Find** to open a Device Profile form.



7 Find and select the profile for the HP Procurve 3500yl Switch; then click **OK**.

Device Model	SNMP Object ID	OUI	Device Family	Device Vendor
hpProCurve2900-48G	.1.3.6.1.4.1.11.2.3.7.11.68	2900	HP ProCurve 2900	Hewlett-Pi
hpProCurve2910ai-24G	.1.3.6.1.4.1.11.2.3.7.11.86	2900	HP ProCurve 2900	Hewlett-Pi
hpProCurve2910ai-24G-	.1.3.6.1.4.1.11.2.3.7.11.84	2900	HP ProCurve 2900	Hewlett-Pi
hpProCurve2910ai-48G	.1.3.6.1.4.1.11.2.3.7.11.87	2900	HP ProCurve 2900	Hewlett-Pi
hpProCurve2910ai-48G-	.1.3.6.1.4.1.11.2.3.7.11.85	2900	HP ProCurve 2900	Hewlett-Pi
hpProCurve3424	.1.3.6.1.4.1.11.2.3.7.11.42	3400	HP ProCurve 3400	Hewlett-Pi
hpProCurve3448	.1.3.6.1.4.1.11.2.3.7.11.43	3400	HP ProCurve 3400	Hewlett-Pi
hpProCurve3500-24	.1.3.6.1.4.1.11.2.3.7.11.111	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve3500-24-Pot	.1.3.6.1.4.1.11.2.3.7.11.109	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve3500-48	.1.3.6.1.4.1.11.2.3.7.11.112	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve3500-48-Pot	.1.3.6.1.4.1.11.2.3.7.11.110	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve3500yl	.1.3.6.1.4.1.11.2.3.7.11.114	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve3500yl-48G	.1.3.6.1.4.1.11.2.3.7.11.59	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve3500yl-PWR	.1.3.6.1.4.1.11.2.3.7.11.58	3500	HP ProCurve 3500	Hewlett-Pi
hpProCurve 4000M	.1.3.6.1.4.1.11.2.3.7.11.9	4000	HP ProCurve 4000	Hewlett-Pi
hpProCurve4104gl	.1.3.6.1.4.1.11.2.3.7.11.27	4100	HP ProCurve 4100	Hewlett-Pi
hpProCurve4108gl	.1.3.6.1.4.1.11.2.3.7.11.23	4100	HP ProCurve 4100	Hewlett-Pi
hpProCurve4202vl-48G	.1.3.6.1.4.1.11.2.3.7.11.56	4200	HP ProCurve 4200	Hewlett-Pi

8 Click **Save and Close** two times.

File View Tools Actions Help

Node Device Filter *

Save and Close Delete Node Device Filter

Changes are not committed until the top-level form is saved!

Device Category: Switch-Router


Device Vendor: Hewlett-Packard

Device Family: HP ProCurve 3500 Fixed-port Switch


Device Profile: hpProCurve3500yl-48G

9 To test this filter, select **global1**.

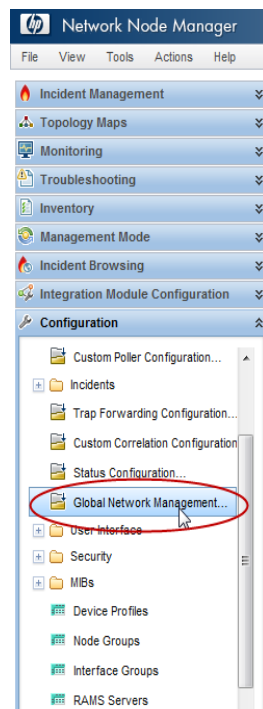
- Using the pull-down, click **Show Members**.



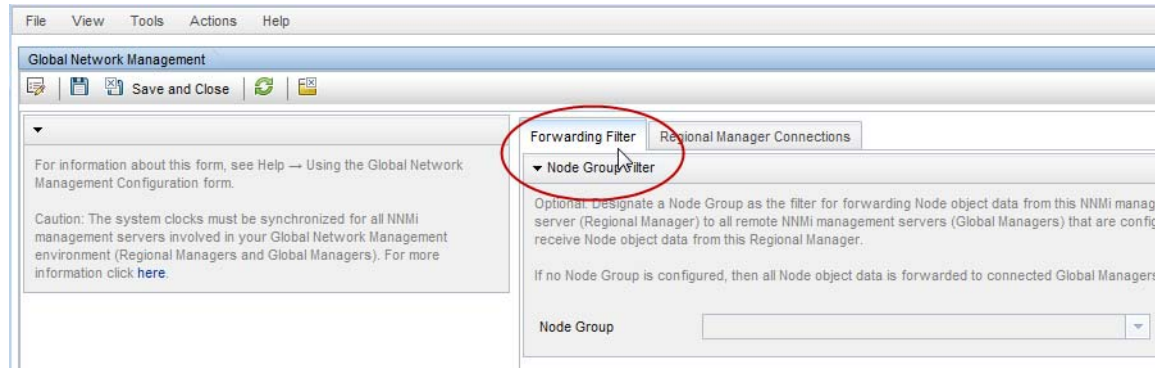
- Notice that NNMi discovered 1 HP 3500yl switch already. This shows you that the filter you created is finding the specific switch models you configured it for. The next step is to configure the forwarding filter using this node filter you just created.



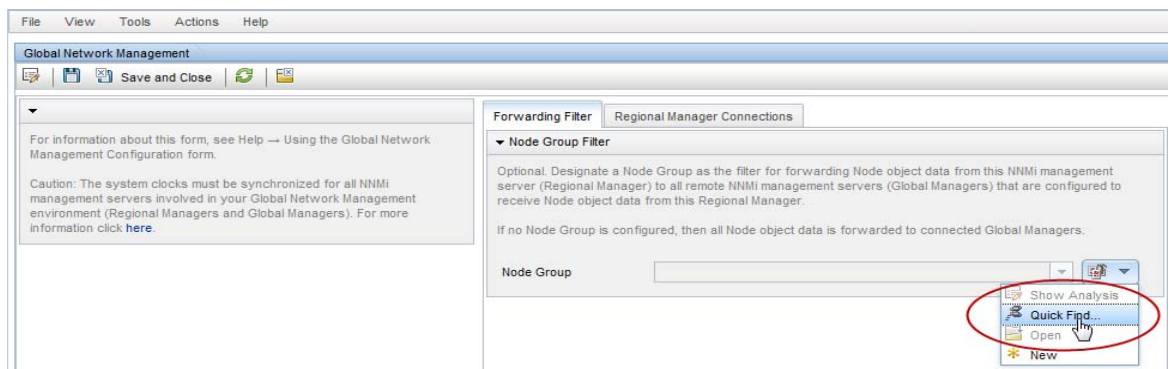
- From regional1's **Configuration** workspace in the NNMi console, click **Global Network Management**.




13 Click the **Forwarding Filter** tab.




14 Click **Quick Find**.



15 Select the **global1** filter; then click **OK**.




16 Click **Save and Close**.

This completes the task of setting up a forwarding filter on `regional1`. After you complete [step 1](#) through [step 16](#) for `regional2`, move on to the next section to connect `global1` to `regional1` and `regional2`.

Connecting a Global Manager with a Regional Manager

As mentioned earlier, suppose that `regional1` and `regional2` both manage several common switches. Suppose that you want this common switch information forwarded to `global1` from `regional1`.



To make that happen you must connect `global1` to `regional1` before connecting it to `regional2`. By using that connection sequence, `global1` considers `regional1` to be the NNMi management server monitoring these common switches. `Global1` also ignores information about these common switches that it receives from `regional2`.



HP recommends you use this feature on a small scale to better understand how it works, then expand it to meet your network management needs.

To connect `global1` first to `regional1`, then to `regional2`, complete the following steps:


- 1 As mentioned earlier, synchronize the NNMi management server clocks for `global1`, `regional1`, and `regional2` before you connect these servers in a global network management configuration. See *Clock Synchronization Issues* in the NNMi help for more information.



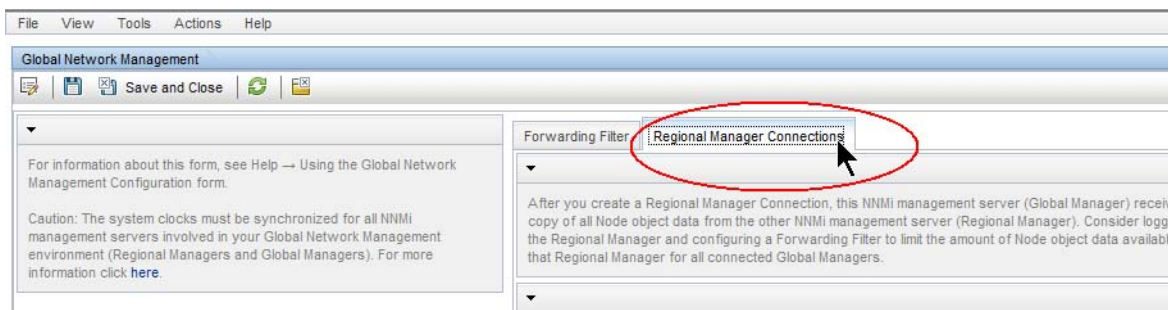
NNMi displays a warning message if there is a connection problem with a regional manager, such as a server clock synchronization problem.

- 2 Set up a connection from `global1` to `regional1`.

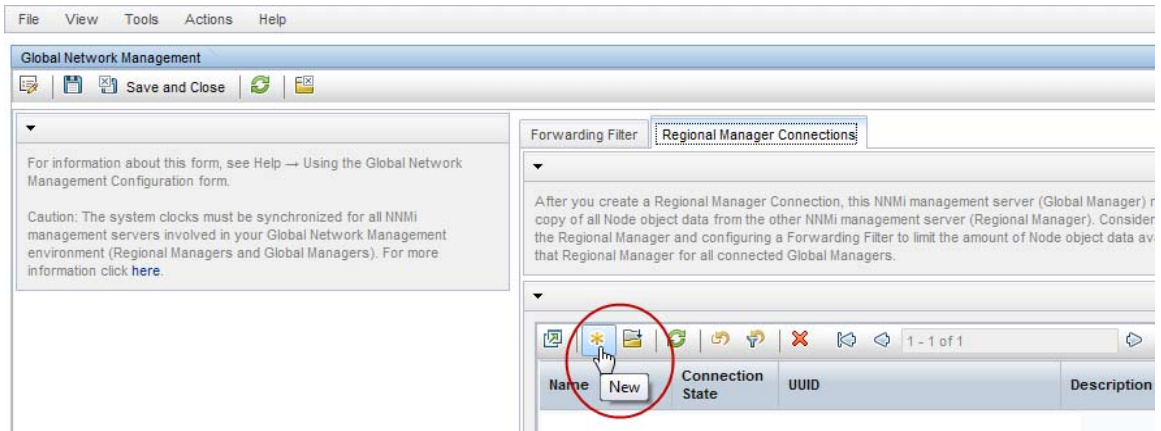
- a From the `global1` NNMi console, click **Global Network Management** in the **Configuration** workspace.



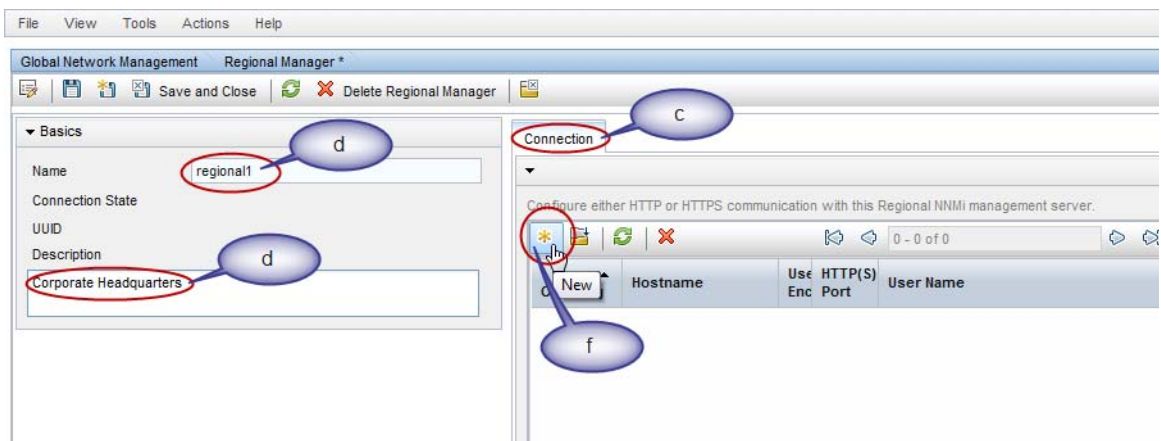
- b Click **Regional Manager Connections**.



- c Click the **New** icon to create a new regional manager.



- d Add the name and description information for regional1.
- e Click the **Connection** Tab.
- f Click the **New** icon.



g Add the connection information for regional1



See **Help->Using the Regional Manager Connection Form** in the NNMi help for specific information about the entries to make in this form

Regional Manager Connection *

Save and Close Delete Regional Manager Connection

Changes are not committed until the top-level form is saved!

Provide the fully-qualified domain name for the hostname of the remote Regional Manager server. See Help → Using the Regional Manager Connection form for more information.

Hostname regional1.example.hp.com

Use Encryption

HTTP(S) Port 80

User Name system

User Password

Ordering 20

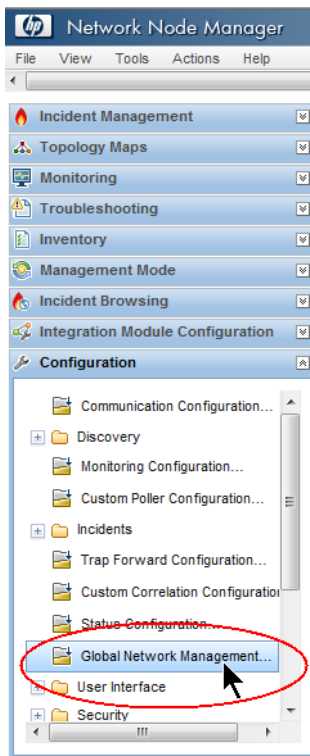
h Click **Save and Close** two times to save your work.

3 Complete [step a](#) on page 247 through [step g](#) on page 249 to establish a connection from global1 to regional2.

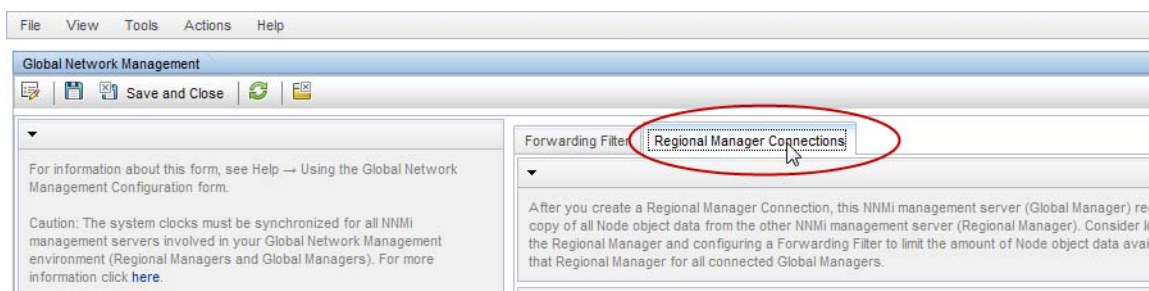
Determining the Connection States from global1 to regional1 and regional2

To check the connection states from global1 to regional1 and regional2, complete the following steps:

- 1 From the global1 NNMI console, click **Global Network Management** in the **Configuration** workspace.



- 2 Click the **Regional Managers Connections** tab.



- 3 Check the status of `regional1` and `regional2` by checking their connection states. Notice that the connection states are shown as `Connected`, which means they are functioning properly.

See *Determine the State of the Connection to a Regional Manager* in the NNMi help for more information.


Do not continue to the next section until NNMi completes a good discovery. See *Checking Discovery Progress* in the *HP Network Node Manager i Software Interactive Installation Guide* for more information.

Reviewing global1 Inventory

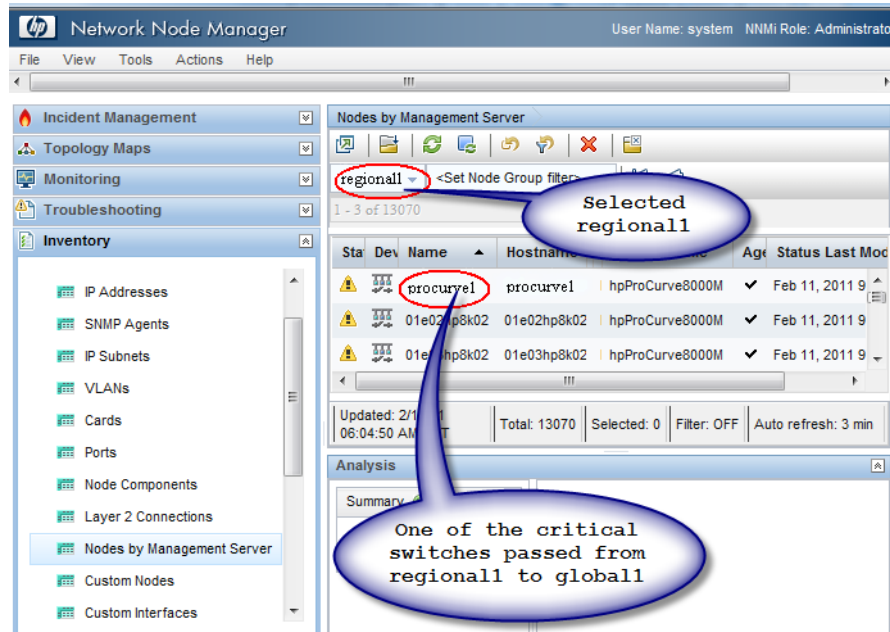
Do not complete this section until NNMi completes a good discovery. See *Checking Discovery Progress* in the *HP Network Node Manager i Software Interactive Installation Guide* for more information.

To view the node information `regional1` forwarded to `global1`, complete the following steps:

- 1 From the `global1` NNMi console, navigate to the **Nodes by Management Server** form located in the **Inventory** workspace.



- Assume that regional1 passed information about switch procurve1.x.y.z to global1. After selecting **regional1**, the inventory might look as follows:




Complete [step 1](#) through [step 2](#) to look at the device inventory passed to global1 from other connected regional managers.

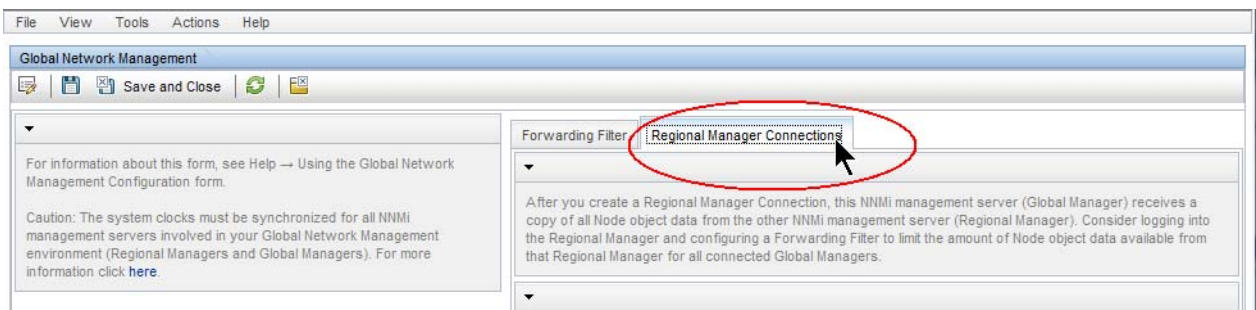
Disconnecting Communication between global1 and regional1

Suppose you plan to permanently shut down global1 or to shut it down for a number of days. Assume for this example that global1 still has active subscriptions to regional1. You must complete some additional steps to complete the shutdown:


- 1 From the global1 NNMI console, click **Global Network Management** in the **Configuration** workspace.




- 2 Click **Regional Manager Connections**.



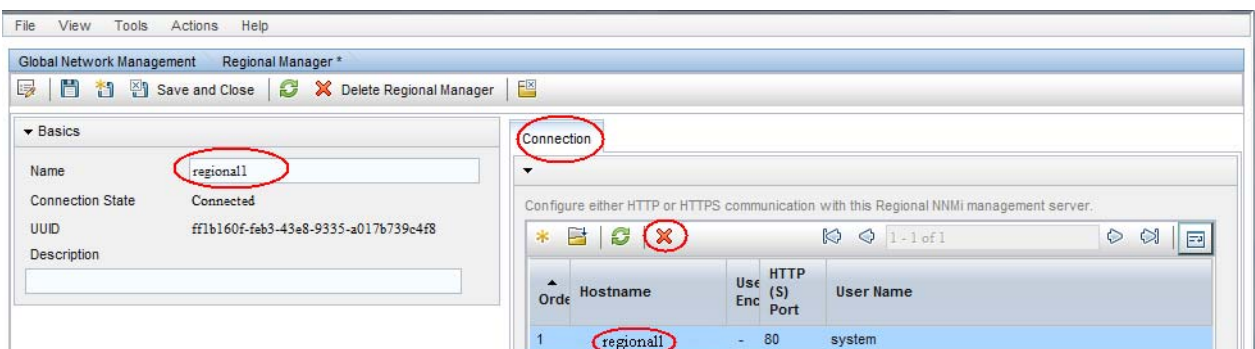
- 3 Check to make sure the status is Connected. If the status is not Connected, diagnose the problem using information from the *Troubleshoot Global Network Management* topic in the NNMi help before continuing.



- 4 Select regional1, then click the Open icon.




- 5 Click **Connection**, select **regional1.x.y.z**, then click the **Delete** icon.



- 6 Click **Save and Close**.
- 7 In the **Regional Manager Connections** tab, note the **Name** attribute value for **regional1** (case-sensitive). You need this text string for the `RemoteNNMiServerName` variable in a later step.

- 8 Click **Save and Close** again.
- 9 On `global1`, at the command line, type the following command:


```
nnmnodedelete.ovpl -rm regional1 -u NNMIadminUserName -p NNMIadminPassword
```
- 10 These commands remove the node records from `global1` that `regional1` forwarded to it. The commands also close incidents associated with the nodes forwarded to `global1` from `regional1`. For detailed information, see *Disconnect Communication with a Regional Manager* in the NNMI help.
- 11 To remove the configuration records for `regional1`, do the following.
 - a Click the **Configuration** workspace.
 - b Select the **Global Network Management** form.
 - c Select the **Regional Manager Connections** tab.
 - d Select `regional1`, then click the Delete icon.



- e Click **Save and Close** to save your deletions.
- 12 Complete [step 1](#) through [step 11](#) for other regional NNMI management servers, such as `regional2`, that are connected to `global1`.

Additional Information

Discovery and Data Synchronization

As network administrators add, delete, or modify devices on a network, regional servers, such as `regional1` and `regional2`, discover those changes and update global servers, such as `global1` in the example in this chapter, `regional1` and `regional2` also notify `global1` of changes that administrators make to the management mode of a node it manages.



To maintain consistency, as `regional1` and `regional2` discover device state changes, they continuously update `global1`, thereby maintaining identical node states on both the global and regional servers.

Any time `global1` requests information about a node that is managed by `regional1` or `regional2`, `regional1` or `regional2` responds to `global1` with the requested information. `global1` never talks directly to a node. There will not be duplicate SNMP queries to devices when `global1` performs a discovery.

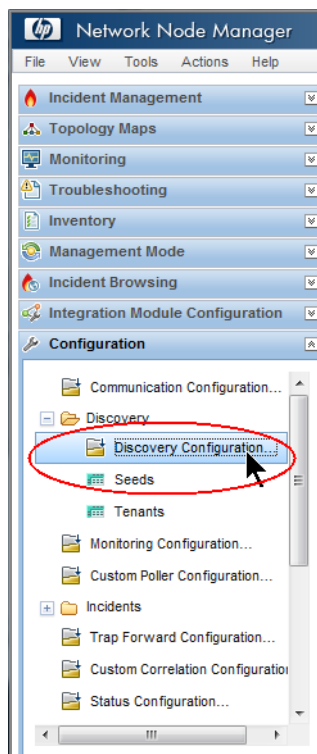
`global1` synchronizes with `regional1` and `regional2` each time `regional1` or `regional2` completes a discovery. NNMi uses FDB (Forwarding Database) data to calculate layer 2 connections. FDB data is very dynamic, and varies a lot between discoveries, especially if there are multiple regionals connected to a global.



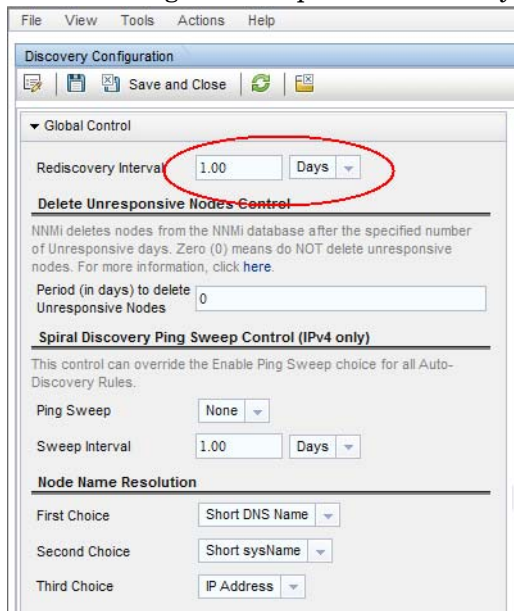
Changes to user-modified or application-modified attributes are not updated on the global during a synchronization.

The Rediscovery Interval is adjustable on each regional, and can make a difference in the discovery accuracy between `global1` and the regional managers. The shorter the Rediscovery Interval, the more accurate the discovery, and the more NNMi-generated network traffic. The longer the Rediscovery Interval, the less accurate the discovery, and the less NNMi-generated network traffic. This means that the larger your network grows, the less frequently you might want to rediscover. To set the Rediscovery Interval, do the following steps:

From the `regional1` or `regional2` NNMi console, click **Discovery Configuration** in the **Configuration** workspace.



- Adjust the Rediscovery Interval according to your how often you want the regionals to initiate a discovery. The global will initiate a discovery immediately after a regional completes a discovery.




- Click **Save and Close**.

Status Poll or Configuration Poll a Device

Suppose regional NNMi management server `regional2` discovers and manages Node X and global NNMi management server `global1` connects with regional NNMi management server `regional2`.

Figure 22 Status Poll or Configuration Poll a Node



To status poll Node X from `global1`, do the following:

- 1 From `global1`, click **Nodes** in the **Inventory** workspace.
- 2 Select Node X from the nodes inventory.
- 3 Request a status poll of Node X using the **Actions > Status Poll** menu item.
- 4 NNMi management server `global1` requests a status poll from regional NNMi management server `regional2` and shows the results on your screen. It does not matter if you initiate the status poll request from either `global1` or `regional2`. You still see the same status poll results.

If you want `global1` to have the most current discovery information for Node X, do the following to configuration poll Node X from `global1`.

- 1 From `global1`, click **Nodes** in the **Inventory** workspace.
- 2 Select Node X from the nodes inventory.
- 3 Request a configuration poll of Node X using the **Actions > Configuration Poll** menu item.
- 4 NNMi management server `global1` requests a configuration poll from regional NNMi management server `regional2` and shows the results on your screen. It does not matter if you initiate the configuration poll request from either `global1` or `regional2`. You still see the same configuration poll results.

Determining Device Status and NNMi Incident Generation using a Global Manager

NNMi management server `global1` listens for state changes coming from regional managers `regional1` and `regional2` and updates the states in its local database.

The NNMi `StatePoller` services on NNMi management servers `regional1` and `regional2` calculate state values for the devices it monitors. `global1` receives state value updates from `regional1` and `regional2`. `global1` polls nodes that it discovers, and does not poll nodes being managed by `regional1` and `regional2`.

After you change the management mode of a node being managed by `regional1`, you see that management mode change on `global1` as well. As network administrators add, remove, or modify network equipment being managed by `regional1` or `regional2`, `regional1` or `regional2` updates `global1` of these network device changes.

`global1` generates incidents using its own causal engine and topology, including the node object data forwarded to it by `regional1` and `regional2`. This means that the incidents it generates might be slightly different from the `regional1` and `regional2` incidents if there are differences in topology.

It is better to avoid using a forwarding filter on `regional1` or `regional2`, as filtering might affect the connectivity on `global1`. The result could be a difference in the root cause analysis between `global1` and the two regionals (`regional1` and `regional2`). In most cases, if you choose to avoid using forwarding filters, a global NNMi management server will have a larger topology. This helps it draw more accurate root cause analysis conclusions.

Without additional configuration, `regional1` does not forward traps to `global1`. To do this, you must configure `regional1` to forward specific traps to `global1`. HP recommends you only configure regional managers to forward low-volume, important traps to avoid excessive burden on the global manager. NNMi drops forwarded traps if the forwarded traps result in a `TrapStorm` incident. See the `TrapStorm` Management Event details in the NNMi console.

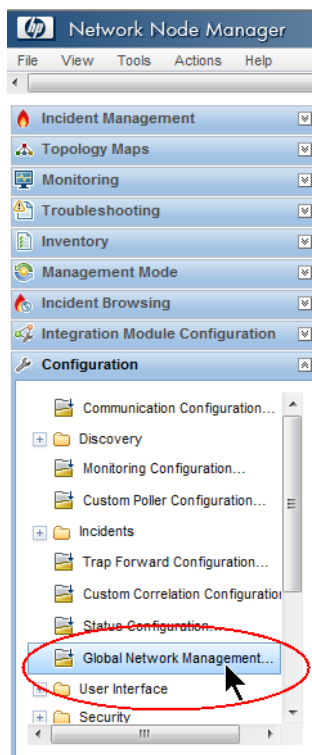
Configuring Application Failover for Global Network Management

You can configure both global and regional managers to use application failover. The global or regional manager automatically detects and connects to the active system.

Configuring Application Failover on a Global Manager

To configure `global1` to recognize the application failover do the following:


- 1 From the `global1` NNMi console, click **Global Network Management** in the **Configuration** workspace.



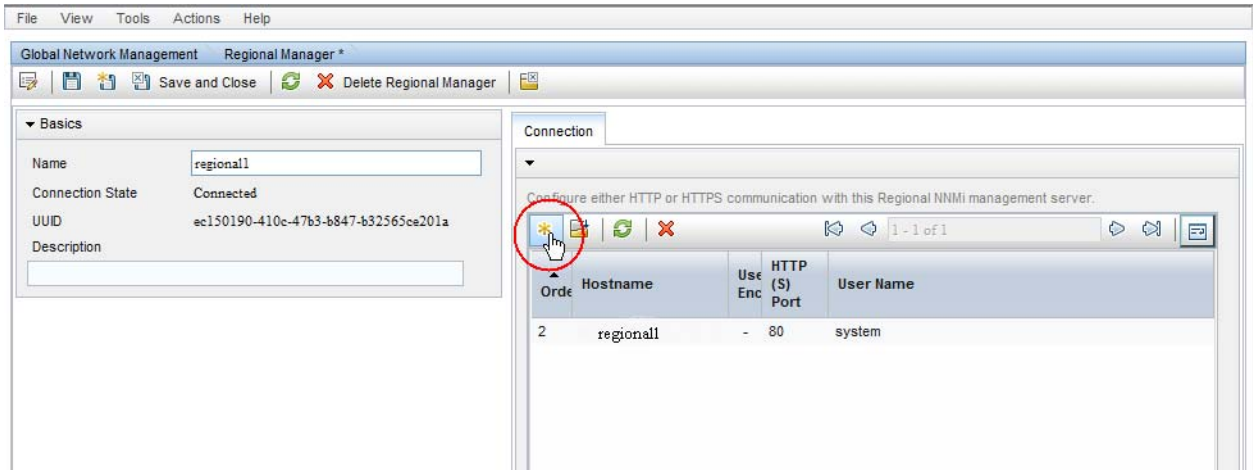
Suppose you configured `regional1` for application failover, and `regional1_backup` as the secondary server.

- 2 Click **Regional Manager Connections**.


3 Select regional1, then click the **Open** icon.



4 Click the **New** icon.



5 Add the **Hostname**, **HTTP or HTTPS Port**, **User Name**, and **Ordering** value. Set the ordering value to a value greater than the regional1 value.



- 6 Click **Save and Close** three times to save your work.

If a regional manager fails, the global manager does the following:

- a It contacts the primary.
- b If the primary does not respond, it contacts the secondary.

If the global system detects that the active system is not responding, it tries to reconnect starting with the lowest order number.

Troubleshooting Tips for Global Network Management

Troubleshooting Information in the NNMi Help

See the *Troubleshoot Global Network Management* topic in the NNMi help for global network management troubleshooting information.

Clock Synchronization

All NNMi management servers in your network environment that participate in global network management (global managers and regional managers) or single sign-on (SSO) must have their internal time clocks synchronized in universal time. Use a Time Synchronization program, for example, the UNIX (HP-UX / Linux / Solaris) tool Network Time Protocol Daemon (NTPD) or one of the available Windows operating system tools.

If you see the following message at the bottom of the NNMi console:

```
NNMi is not connected to 1 Regional Manager(s). See Help ? System
Information, Global Network Management.
```

Check the `nmn.0.0.log` file on the Global Manager for the following message:

```
WARNING: Not connecting to system <serverName> due to clock
difference of <number of seconds>. Remote time is <date/time>.
```

Perhaps the clocks have drifted apart and need to be resynchronized. Check the `nmn.0.0.log` file on the Global Manager for the following message:

```
WARNING: Not connecting to system <serverName> due to clock
difference of <number of seconds>. Remote time is <date/time>.
```

Within a few minutes of this warning, NNMi disconnects the Regional Manager Connection. And the following message appears at the bottom of the NNMi console:

```
NNMi is not connected to 1 Regional Manager(s). See Help ? System
Information, Global Network Management.
```

Global Network Management System Information

Select **Help > System Information**, then click the **Global Network Management** tab to view information about your global network management connections.

Synchronize Regional Manager Discovery from a Global Manager

Suppose you notice an information inconsistency between `global1` and `regional2`. To fix that, run the `nnmnoderediscover.ovpl` script from `global1`, causing `global1` and `regional2` to synchronize. This also results in the `regional2` updating `global1` with any new discovery results.

Consider the network shown in [Figure 22](#) on page 258. Suppose you want `regional2` to synchronize its entire set of nodes, nodes: X, Y and Z, with `global1`. Run the following command to synchronize nodes X, Y, and Z with `global1`: `nnmnoderediscover.ovpl -u username -p password -rm regional2`.



You can use the `-fullsync` flag with the `nnmnoderediscover.ovpl` command to synchronize all polled object states and status (although this takes more time and causes a greater load on the systems). For more information, see the `nnmnoderediscover.ovpl` reference page, or the UNIX manpage.



Note the following:

- Updates to Status and Incidents could be delayed as NNMi resynchronizes following a manual resynchronization.
- If you see the following message during this resynchronization, it does not indicate a problem:

The Causal Engine's large queue size is causing delayed updates to Status and Incidents. This could be due to resynchronization following an upgrade, application failover, restore from backup, or a manual resynchronization.
- Do not stop NNMi during this resynchronization. To ensure resynchronization has completed, keep NNMi running for several hours following the manual resynchronization.

Remedying a Destroyed Database on global1

If you take `global1` out of service and need to restore its database, you face several scenarios:

- 1 If you restore `global1`'s database successfully, `regional1` and `regional2` synchronize their cached information with `global1`. There are no manual steps to perform after bringing `global1` back online.
- 2 If `global1` is out of service for an extended period of time, [step 1](#) might not work successfully. To remedy this, run the `nnmnoderediscover.ovpl` script on `global1` to initiate a new discovery on `global1`, `regional1` and `regional2`. In this case you could run status polls on key devices to more quickly get updated status information.
- 3 If you cannot recover `global1`'s database then you should submit a support call to clear out the old `global1` data from the `regional1` and `regional2` databases using the `nnmsubscription.ovpl` script.

Upgrading Global and Regional Managers from NNMi 9.0x/9.1x to NNMi 9.20

NNMi Versions Supported by Global Network Management

If a global manager is connected to a regional manager running NNMi 9.0x patch 2 or earlier, SNMP queries between the global and regional manager do not work. To remedy this, upgrade the regional manager to NNMi 9.0x patch 3 or later. To achieve the best results, the global manager must be the same version and NNMi patch level as the regional manager.

- ▶ HP does not support a regional manager running NNMi 9.0x or 9.1x connected to a global manager running NNMi 9.20. The global manager and regional managers must be running the same NNMi version.

Global Network Management Upgrade Steps

When upgrading NNMi management servers configured in a global network management environment to NNMi 9.20, the connections between the global manager and regional managers will drop until both the global and regional managers are upgraded to 9.20. For this reason, HP recommends you upgrade all of the servers at approximately the same time to minimize the total downtime.

For example, you might upgrade the NNMi management servers using the following steps:

- 1 Upgrade the regional managers to NNMi 9.20 and ensure proper operation. The global manager stays disconnected during the regional upgrades.
- 2 Upgrade the global manager to NNMi 9.20. The global manager performs a full resynchronization to obtain all events that occurred while the connection between the global manager and the regional managers was down. The effect is the same as if the administrator were to issue `nnmnodediscover.ovpl -all -fullsync` from the global manager. See the `nnmnodediscover.ovpl` reference page or the UNIX manpage for more information.

- ▶ Note the following:

- Updates to Status and Incidents could be delayed as NNMi resynchronizes following an upgrade.
- If you see the following message during this resynchronization, it does not indicate a problem:

The Causal Engine's large queue size is causing delayed updates to Status and Incidents. This could be due to resynchronization following an upgrade, application failover, restore from backup, or a manual resynchronization.

- Do not stop NNMi during this resynchronization. To ensure resynchronization has completed, keep NNMi running for several hours following the upgrade.

Global Network Management and NNM iSPIs or Third-Party Integrations

Each NNM iSPI or third-party integration has its own unique deployment guideline. For the examples in this chapter, you can deploy some NNM iSPIs on `regional1` only, `global1` only, or on both `regional1` and `global1`. For other NNM iSPIs or third-party integrations, you must have them installed on both `regional1` and `global1`. See the documentation for the NNM iSPI or third-party integration for more information.

HP Network Node Manager iSPI Performance for Metrics Software

If NNMi is deployed in a global network management environment, you must do the following:

- 1 Deploy one instance of Network Performance Server (NPS) for each NNMi management server. Every regional manager and the global manager must have separate instances of NPS installed and deployed.
- 2 Run the enablement script once on every regional manager and on the global manager.

Global Network Management and Address Translation Protocols

Each group of dynamic Network Address Translation (NAT) or dynamic Port Address Translation (PAT) or dynamic Network Address and Port Translation (NAPT) requires an NNMi regional manager, in addition to a tenant that is unique within the entire NNMi global network management configuration. See [Managing Overlapping IP Addresses in NAT Environments](#) on page 193. See also the NNMi help.

Configuring NNMi Advanced for IPv6

You must purchase and install an NNMi Advanced license to use the IPv6 management feature. References to NNMi in this chapter are for NNMi with an NNMi Advanced license installed.

IPv6 management in NNMi enables the discovery and monitoring of IPv6 addresses, including their interfaces, nodes and subnets. To provide a seamless integration, NNMi extends its IP Address model to include both IPv4 and IPv6 addresses. Whenever possible, NNMi treats all IP Addresses equally; most of the features associated with an IPv4 address are also available for IPv6 addresses. However, there are some exceptions. See the NNMi help for more information about IPv6 information displayed in the NNMi console.

This chapter contains the following topics:

- [Feature Description](#)
- [Prerequisites](#)
- [Licensing](#)
- [Supported Configuration](#)
- [Installing NNMi](#)
- [Activating IPv6 Features](#)
- [Deactivating IPv6 Features](#)

Feature Description

The NNMi IPv6 management feature provides the following:

- IPv6 inventory discovery for IPv6-only and dual-stacked devices
 - IPv6 addresses
 - IPv6 subnets
 - Associations between IPv6 Addresses, Subnets, Interfaces and Nodes

- Native IPv6 SNMP communication for the following:
 - Node discovery
 - Interface monitoring
 - Trap and inform reception and forwarding
- Automatic selection of IPv4 or IPv6 communication (management address) for dual-stacked devices. Use the NNMi console to set the SNMP management address preference to IPv4 or IPv6 using **Communication Configuration** located in the **Configuration** workspace.
- Native ICMPv6 communication for IPv6 Address fault monitoring.
- Seeded device discovery using an IPv6 address or hostname
- Automatic IPv6 device discovery using IPv6 Layer 3 neighbor discovery hints
- Automatic IPv6 device discovery using layer 2 neighbor discovery hints using LLDP (Link Layer Discovery Protocol) IPv6 neighbor information
- Consolidated presentation of IPv4 and IPv6 information
 - Inventory views for nodes, interfaces, addresses, subnets, and associations
 - Layer 2 Neighbor View and Topology Maps for IPv4 and IPv6 devices
 - Layer 3 Neighbor View and Topology Maps for IPv4 and IPv6 devices
 - Incidents, conclusions, root-cause analysis
- NNMi console actions: ping and traceroute for IPv6 addresses and nodes
- NNMi configuration using IPv6 addresses and address ranges
 - Communication configuration
 - Discovery configuration
 - Monitoring configuration
 - Node & Interface Groups
 - Incident configuration
- SDK Web-services support for IPv6 inventory and incidents
- NNM iSPI Performance for Metrics support for IPv6 interfaces

The NNMi IPv6 management feature excludes the following:

- Discovery of IPv6 subnet connections
- Use of IPv6 ping sweep for discovery
- IPv6 Network Path View (Smart Path)
- IPv6 Link Local Address fault monitoring
- Using IPv6 Link Local Addresses as discovery seeds

Prerequisites

Review the *NNMi Deployment Reference*, *NNMi Release Notes*, and *NNMi System and Device Support Matrix* for details on management server specifications and NNMi installation.

To use native IPv6 communication, the NNMi management server must be a dual-stacked system, meaning that it communicates using both IPv4 and IPv6.

Important: If you have IPv6 discovery configured on HP NNMi, and are using the HP Universal CMDB (HP UCMDB) integration, the UCMDB HP Discovery and Dependency mapping (DDM) import task fails. You need to disable IPv6 discovery to use the HP UCMDB integration with HP NNMi.

IPv6 is not supported on Windows operating systems. See the *NNMi System and Device Support Matrix* for information about the supported operating systems for IPv6. There are other requirements listed below:

- You must enable and configure IPv4 on at least one network interface.
- You must enable IPv6 and have a global unicast address or a unique local unicast address configured on at least one network interface that is connected to the IPv6 network you want to manage.
- You must configure IPv6 routes on the NNMi management server to enable NNMi to communicate with any devices you want NNMi to discover and monitor using IPv6.



You can use an IPv4-only NNMi management server, but doing so will limit NNMi from fully managing IPv4/IPv6 dual-stacked devices. For example, if you use an IPv4-only management server, NNMi cannot discover IPv6-only devices, cannot discover using IPv6 seeds and hints, and cannot monitor for faults on devices having IPv6 addresses.

The DNS server used by the NNMi management server must resolve hostnames to and from IPv6 addresses. For example, it must be able to resolve to and from an AAAA DNS record. That means the DNS server must map a hostname to a 128-bit IPv6 address. If an IPv6-capable DNS server is not available, NNMi will still function correctly; however NNMi does not determine nor display DNS hostnames for nodes using IPv6 addresses.

Licensing

As mentioned earlier, you must purchase and install an NNMi Advanced license to use the IPv6 management feature. For information about obtaining and installing your NNMi Advanced license, see [Licensing NNMi](#) on page 121.

The NNMi product includes a temporary Instant-On license password. This is a temporary, but valid NNMi Advanced license. You should obtain and install a permanent license password as soon as possible.

Supported Configuration

See the *NNMi System and Device Support Matrix* for additional information about the supported operating system configurations for NNMi.

Management Server

The following table shows the capabilities of both the IPv4-only and dual-stacked NNMi management server.

Table 25 Management Server Capabilities

Feature/Capability	IPv4-Only	Dual-Stack
IPv4 Communication (SNMP, ICMP)	Supported	Supported
IPv6 Communication (SNMP, ICMPv6)	Not Supported	Supported
Dual-Stack Managed Node	Supported	Supported
Discovery using IPv4 Seed	Supported	Supported
Discovery using IPv6 Seed	Not Supported	Supported
IPv4 Address and Subnet Inventory	Supported	Supported
IPv6 Address and Subnet Inventory	Supported	Supported
Interface Status and Performance using SNMP	Supported	Supported
IPv4 Address Status using ICMP	Supported	Supported
IPv6 Address Status using ICMPv6	Not Supported	Supported
IPv6-only Managed Node	Not Supported	Supported
Discovery using IPv6 Seed	Not Supported	Supported
IPv6 Address and Subnet Inventory	Not Supported	Supported

Table 25 Management Server Capabilities (cont'd)

Feature/Capability	IPv4-Only	Dual-Stack
Interface Status and Performance using SNMP	Not Supported	Supported
IPv6 Address Status using ICMPv6	Not Supported	Supported
IPv4-only Managed Node	Supported	Supported
Node Discovery using IPv4 Seed	Supported	Supported
Node Discovery using IPv4 Seed	Supported	Supported
Interface Status and Performance using SNMP	Supported	Supported
Interface Status and Performance using SNMP	Supported	Supported
IPv4 Address and Subnet Inventory	Supported	Supported

Supported SNMP MIBs for IPv6

NNMi supports the following SNMP MIBs for IPv6:

- RFC 4293 (current IETF standard)
- RFC 2465 (original IETF proposal)
- Cisco IP-MIB

Installing NNMi

During NNMi installation, the installation script includes IPv6 features; however, you must manually enable these IPv6 features. First, you must purchase and apply an NNMi Advanced license to enable the IPv6 features. Then you must manually configure IPv6 to work by editing the `nms-jboss.properties` file.

Activating IPv6 Features

Features requiring IPv6 communication, such as the discovery and of IPv6 only devices and the monitoring of IPv6 address status, require an NNMi management server to have an IPv6 global unicast address configured and operational.

The procedure shown below explains how to enable IPv6 features by doing the following:

- Install an NNMi Advanced license
- Enable the IPv6 Master Switch in the `nms-jboss.properties` file

► Review and verify all of the prerequisites described in the preceding section before continuing.

- 1 Use the temporary Instant-on license that comes with NNMi, or install an NNMi Advanced license. For information about obtaining and installing NNMi licenses, see [Licensing NNMi](#) on page 121. IPv6 features are not available with the basic NNMi license.
- 2 Edit the `nms-jboss.properties` file. Look in the following location:
 - *UNIX*: `$NNM_PROPS/nms-jboss.properties`
- 3 Locate the text that begins with `# Enable NNMi IPv6 Management`.

► NNMi provides a complete description of each property, showing them as comments in the `nms-jboss.properties` file.

- a To enable IPv6 communication in NNMi, un-comment the property:


```
java.net.preferIPv4Stack=false
```

► To un-comment a property, remove the `#!` characters from the beginning of a line.

- b To enable overall IPv6 management in NNMi, un-comment the property:


```
com.hp.nnm.enableIPv6Mgmt=true
```
- c Save and close the `nms-jboss.properties` file.

- 4 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.
- 5 Check the NNMi processes using the following command:

```
ovstatus -v ovjboss
```

Successful startup should look something like the following:

```
object manager name: ovjboss
state:                RUNNING
PID:                  <Process ID #>
last message:        Initialization complete.
exit status:          -
additional info:
```


SERVICE	STATUS
CommunicationModelService	Service is started
CommunicationParametersStatsService	Service is started
CustomPoller	Service is started
IslandSpotterService	Service is started
ManagedNodeLicenseManager	Service is started
MonitoringSettingsService	Service is started
NamedPoll	Service is started
msApa	Service is started
NmsCustomCorrelation	Service is started
NmsDisco	Service is started
NmsEvents	Service is started
NmsEventsConfiguration	Service is started
NmsExtensionNotificationService	Service is started
NnmTrapService	Service is started
PerformanceSpiAdapterTopologyChangeService	Service is started
PerformanceSpiConsumptionManager	Service is started
RbaManager	Service is started
RediscoverQueue	Service is started
SpmdjbossStart	Service is started
StagedIcmp	Service is started
StagedSnmp	Service is started
StatePoller	Service is started
TrapConfigurationService	Service is started
TrustManager	Service is started

- 6 After you enable IPv6, NNMi views immediately include the IPv6 inventory for newly discovered nodes. During the next discovery cycle, NNMi views show the IPv6 inventory associated with previously discovered nodes.
- 7 Optionally set the SNMP management address preference for dual-stacked managed nodes. Dual-stacked managed nodes are those nodes that can communicate using either IPv4 or IPv6. To do this, complete the following steps:
 - a From the NNMi console, click **Communication Configuration** located in the **Configuration** workspace.
 - b Locate the **Management Address Selection** section. Select IPv4, IPv6, or Any in the IP Version Preference field.
 - c Save your changes.

- d Restart the NNMi management server:

Run the **ovstop** command on the NNMi management server.

Run the **ovstart** command on the NNMi management server.

To speed things up, select nodes that you know are dual-stack nodes, and then use the **Actions > Configuration Poll** command located in the NNMi console. You can also use the `nmmnoderediscover.ovpl` script to add nodes to the NNMi discovery queue. See the `nmmnoderediscover.ovpl` reference page, or the UNIX manpage, for more information.

After you enable IPv6 communication on the NNMi management server, NNMi begins monitoring nodes for IPv6 address faults using ICMPv6.

Deactivating IPv6 Features

You can administratively disable IPv6 features using one of the following methods:

- 1 Turn off the IPv6 master switch in the `nms-jboss.properties` file, then restart NNMi.
- 2 Let the NNMi Advanced license expire, or replace it with a basic NNMi license.

For information about changing the NNMi license, see [Licensing NNMi](#) on page 121.

The following sections describe NNMi behavior and inventory cleanup after you disable IPv6.

IPv6 Monitoring Following Deactivation

If IPv6 management or IPv6 communication becomes completely disabled, the `StatePoller` service immediately stops monitoring IPv6 addresses with ICMPv6. NNMi sets the IP address state of these addresses to `Not Polled`. If you select an address, then use the **Actions > Monitoring Settings** for this address, NNMi displays `Fault ICMP Polling enabled: false` even though the associated `Monitoring Configuration` rule has the `IP Address Fault Polling enabled`.

IPv6 Inventory Following Deactivation

Once NNMi completely discovers your IPv6 inventory, you can enable NNMi to clean it up automatically in the following scenarios:

- You turned on the master IPv6 switch, then turned it off and restarted NNMi.
NNMi does not immediately remove the IPv6 inventory. NNMi removes the IPv6 inventory for SNMP nodes during the next discovery cycle. NNMi does not remove non-SNMP IPv6 nodes. You must manually delete IPv6 nodes from the NNMi inventory.
- Your NNMi Advanced license expired or someone removed the license. NNMi begins using the NNMi basic license, and the basic license has enough capacity to continue managing all of the discovered nodes.

NNMi immediately removes all of the non-SNMP IPv6 nodes from its inventory. NNMi rediscovers all of the SNMP nodes and removes all of the IPv6 data.

- Your NNMi Advanced license expired or someone removed the license. NNMi begins using the NNMi basic license, and the basic license does not have enough capacity to continue managing all of the discovered nodes. NNMi immediately removes all non-SNMP IPv6 nodes. The `Licensing` service marks the SNMP nodes that exceed the licensed inventory capacity with an `unmanaged` state. NNMi immediately removes IPv6 data from the managed SNMP nodes.

For the unmanaged SNMP nodes, complete these steps:

- a Install additional license capacity.
- b Use the **Actions > Management Mode > Manage** command located in the NNMi console to change the management mode for the nodes marked as `unmanaged` by the `Licensing` service. You can use the `nmmmanagementmode.ovpl` script to manage these nodes as well. See the *nmmmanagementmode.ovpl* reference page, or the UNIX manpage, for more information.
- c Use the **Actions > Configuration Poll** command located in the NNMi console to enable NNMi to discover them. You can use the `nmmnoderediscover.ovpl` script to discover these nodes as well. See the *nmmnoderediscover.ovpl* reference page, or the UNIX manpage, for more information.

- Your NNMi Advanced license expired or someone removed the license; you neglected to install an NNMi basic license.

NNMi immediately removes all non-SNMP IPv6 nodes and automatically unmanages the remaining nodes. To remedy this situation, complete these steps:

- a Install a valid license.
- b Use the **Actions > Management Mode > Manage** command located in the NNMi console to change the management mode for the nodes marked as unmanaged by the Licensing service. You can use the `nnmmanagementmode.ovpl` script to manage these nodes as well. See the `nnmmanagementmode.ovpl` reference page, or the UNIX manpage, for more information.
- c Use the **Actions > Configuration Poll** command located in the NNMi console to enable NNMi to discover the nodes you changed from unmanaged to managed. You can use the `nnmnoderediscover.ovpl` script to discover these nodes as well. See the `nnmnoderediscover.ovpl` reference page, or the UNIX manpage, for more information.
- d To create an IPv6 list, then remove the IPv6 inventory, use the **Actions > Configuration Poll** command to obtain configuration information from each managed node.

Known Issues When Cleaning Up IPv6 Inventory

You could experience leftover IPv6 inventory in the following situation: Suppose that NNMi successfully uses SNMP to manage an IPv6 node, then the node becomes inaccessible before the next discovery. Due to the design of the existing discovery system, the discovery process cannot update a node that loses its ability to communicate using SNMP. To remove these remaining nodes, you must fix the communication problem, then use the **Actions > Configuration Poll** command located in the NNMi console to obtain configuration information from these nodes. For native IPv6 nodes, delete the node directly from the NNMi console.

Running NNMi in a Solaris Zones Environment

For the supported versions of the Solaris operating system, HP Network Node Manager i Software (NNMi) runs without special configuration in a Solaris Zones environment.

This chapter contains the following topics:

- [Installing NNMi in a Solaris Zone](#)
- [Trap Forwarding in a Solaris Zone](#)
- [Running NNMi Application Failover in a Solaris Zones Environment](#)
- [Running NNMi under HA in a Solaris Zones Environment](#)

Installing NNMi in a Solaris Zone

If you plan to implement NNMi application failover in a Solaris Zones environment, see [Running NNMi Application Failover in a Solaris Zones Environment](#) on page 278.

If you plan to run the Solaris Zone under high availability (HA), see [Running NNMi under HA in a Solaris Zones Environment](#) on page 278.

For all other deployment models, install NNMi as described in the *HP Network Node Manager i Software Interactive Installation Guide*.

Trap Forwarding in a Solaris Zone

Suppose you want to forward the SNMP traps that NNMi receives from managed devices to another application. To do this, navigate to the **Trap Forward Configuration** in the **Configuration** workspace. See the NNMi help for more information.

Because the Solaris Zones environment does not support raw trap forwarding, do not select the **Original Trap** forwarding option. Choose one of the other forwarding options when running NNMi in a Solaris Zones environment.

Running NNMi Application Failover in a Solaris Zones Environment

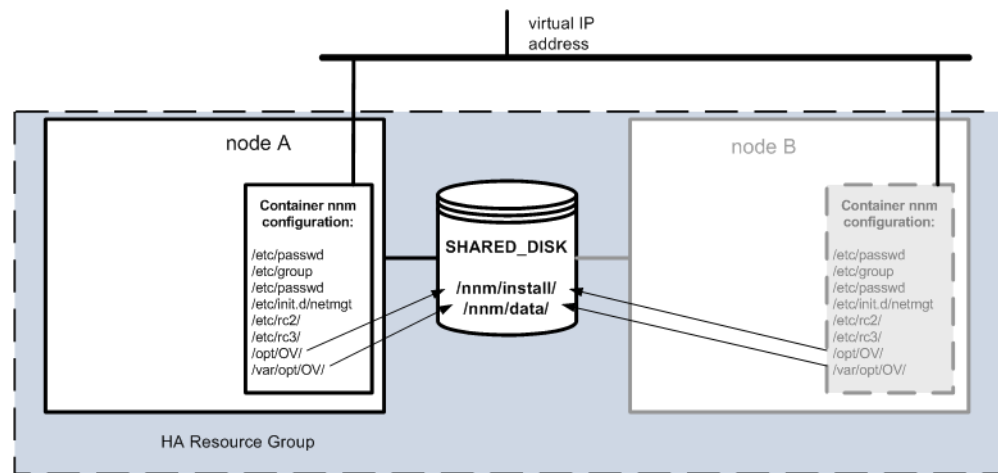
If you want to use the NNMi application failover feature in a Solaris Zones environment, install NNMi in its own zone on each of two physical systems.

Configure application failover as described in [Configuring NNMi for Application Failover](#) on page 283. Throughout the procedure, “server X” refers to one zone and “server Y” refers to the other zone.

Running NNMi under HA in a Solaris Zones Environment

In a Solaris Zones environment, you do not need to implement the NNMi-provided solution for running NNMi in an HA cluster. Because Veritas Cluster Server (VCS) is zone-aware, configure the HA resource group for the zone, as shown in [Figure 23](#).

Figure 23 NNMi in a Solaris Zone Running under HA



The configuration to run NNMi in this environment is minimal. The NNMi installation process creates the `nmsdbmgr` user in the `nmsdb` group and adds startup configuration to the host system. Replicate this setup to the second node in the HA cluster.

To install NNMi to run in a zone inside an HA resource group, follow these steps:

- 1 On the shared disk, create the NNMi installation folders:
 - /nmm/install
 - /nmm/data
- 2 On node A, create and prepare a new zone called **nmm**:
 - a Create zone **nmm**, as described in the Solaris zone documentation. Note all configuration parameters set during zone creation.
 - b Start zone **nmm**.

- c Log on to zone **nnm**, and then create the following symbolic links:
 - /opt/OV/ pointing to /nnm/install/ on the shared disk
 - /var/opt/OV/ pointing to /nnm/data/ on the shared disk
 - d Log off from and then shut down zone **nnm**.
 - 3 On node B, create an identical new zone called **nnm**, and then install NNMi:
 - a Create zone **nnm** with identical properties (including IP address) as zone **nnm** on node A.
 - b Start zone **nnm**.
 - c Log on to zone **nnm**, and then create the following symbolic links:
 - /opt/OV/ pointing to /nnm/install/ on the shared disk
 - /var/opt/OV/ pointing to /nnm/data/ on the shared disk
 - d Instruct the NNMi installer to follow the symbolic links by entering the following command:

```
PKG_NONABI_SYMLINKS=true
```
 - e Install NNMi inside the **nnm** zone.

NNMi installs into the /nnm/install/ and /nnm/data/ directories on the shared disk.
 - f Copy the following files to a temporary location (such as the shared disk) that is accessible from outside the **nnm** zone:
 - /etc/passwd
 - /etc/group
 - /etc/shadow
 - /etc/init.d/netmgt
 - g Log off from and then shut down zone **nnm**.
 - 4 On node A, copy the NNMi-modified system files, and then start NNMi:
 - a Start zone **nnm**.
 - b Log on to zone **nnm**, and then copy the files from the temporary location identified in [step 3](#) to the correct location in the zone:
 - /etc/passwd
 - /etc/group
 - /etc/shadow
 - /etc/init.d/netmgt

- c Create the following symbolic links (to duplicate the configuration created during NNMi installation on node B):
 - /etc/rc0.d/K01netmgt pointing to /etc/init.d/netmgt
 - /etc/rc1.d/K01netmgt pointing to /etc/init.d/netmgt
 - /etc/rc2.d/K01netmgt pointing to /etc/init.d/netmgt
 - /etc/rc3.d/S98netmgt pointing to /etc/init.d/netmgt
 - /etc/rcS.d/K01netmgt pointing to /etc/init.d/netmgt
- d Start NNMi by running the following command:

ovstart

- 5 Configure Veritas Cluster Server to create a resource group containing zone **nnm** on both node A and node B.

For more information, see the VCS documentation.

Resilience

HP Network Node Manager i Software (NNMi) supports two different approaches to protecting the NNMi data in case of hardware failure:

- NNMi application failover provides for disaster recovery by maintaining a copy of the embedded NNMi database transaction logs on an identically configured system. (If NNMi uses an Oracle database, the two systems connect to the same database at different times.)
- Running NNMi in a high availability (HA) cluster provides for nearly one hundred percent availability of the NNMi management server by maintaining the embedded NNMi database and configuration files on a shared disk. (If NNMi uses an Oracle database, the shared disk contains the NNMi configuration files, and the two systems connect to the same database at different times.)

In both approaches, if the current NNMi management server fails, the second system automatically becomes the NNMi management server.

Table 26 compares several aspects of these two approaches to NNMi data resilience.

Table 26 NNMi Data Resilience Comparison

Item for Comparison	NNMi Application Failover	NNMi Running in an HA Cluster
Required software products	NNMi or NNMi Advanced	<ul style="list-style-type: none"> • NNMi or NNMi Advanced • A separately purchased HA product
Time to fail over	<p>Embedded NNMi database: Time to process the transaction logs (under normal conditions, 10-60 minutes for NNMi without any NNM iSPIs).</p> <p>Oracle NNMi database: Almost instantaneous.</p>	Under normal conditions, 5-30 minutes for NNMi without any NNM iSPIs.
Transparency of failover	Partial. The IP address of the NNMi management server changes to the physical address of what was the standby server. Users must connect to the NNMi console using the new IP address. Some applications follow the movement of the NNMi management server, but most (including the NNM iSPIs) do not.	Complete. All connections use the virtual IP address of the HA cluster, which does not change on failover.
Relative proximity of active and standby servers	LAN or WAN	LAN or WAN (some HA products only)
Licenses purchased	<p>For each feature:</p> <ul style="list-style-type: none"> • A production license tied to the IP address of the initial active server. • A non-production license tied to the IP address of the initial standby server. 	<p>For each feature:</p> <p>A production or non-production license tied to the virtual IP address of the NNMi HA resource group.</p>

Table 26 NNMi Data Resilience Comparison


Item for Comparison	NNMi Application Failover	NNMi Running in an HA Cluster
Licenses installed	<ul style="list-style-type: none"> • Production license keys on the initial active server. • Non-production license keys on the initial standby server. 	<ul style="list-style-type: none"> • Non-production license keys on the initial active server and managed on the shared disk.
Support for NNM iSPIs	Support varies. See the documentation for each NNM iSPI.	
Interaction with Global Network Management	<ul style="list-style-type: none"> • Can configure each global manager for application failover or HA. • Can configure each regional manager for application failover or HA. • Each of these configurations requires two physical or virtual systems.^a • If a global manager or regional manager fails over, NNMi re-establishes the connections between the global managers and regional managers. 	
NNMi maintenance	NNMi must be taken out of the application failover cluster before applying a patch or upgrading.	NNMi can be patched and upgraded without unconfiguring HA.

a. Virtual machine support for HA is dependent on HA software vendors' support of virtual systems.

This section contains the following chapters:

- [Configuring NNMi for Application Failover](#)
- [Configuring NNMi in a High Availability Cluster](#)

Configuring NNMi for Application Failover



Many information technology professionals depend on HP Network Node Manager i Software (NNMi) to notify them when critical network equipment fails and to provide them with a root cause for the failure. They also need NNMi to continue to notify them of network equipment failures, even when the NNMi management server fails. **NNMi application failover** meets this need, transferring application control of NNMi processes from an active NNMi management server to a standby NNMi management server, providing continuance of NNMi functionality.

This chapter contains the following topics:

- [Application Failover Overview](#)
- [Application Failover Basic Setup](#)
- [Configuring NNMi for Application Failover](#)
- [Using the Application Failover Feature](#)
- [Returning to the Original Configuration Following a Failover](#)
- [NNM iSPIs and Application Failover](#)
- [Integrated Applications](#)
- [Disabling Application Failover](#)
- [Administrative Tasks and Application Failover](#)
- [Network Latency/Bandwidth Considerations](#)

Application Failover Overview

The application failover feature is available for NNMi installations that use either the embedded or Oracle databases. After configuring your systems to use the application failover feature, NNMi detects an NNMi management server failure and triggers a secondary server to assume NNMi functionality.

The following terms and definitions apply to configuring NNMi for application failover:

- **Active:** The server running the NNMi processes.
- **Standby:** The system in the NNMi cluster that is waiting for a failover event; this system is not running NNMi processes.
- **Cluster Member:** A Java process running on a system that is using JGroups technology to connect to a cluster; you can have multiple members on a single system.
- **Postgres:** The embedded database NNMi uses to store information such as topology, incidents, and configuration information.
- **Cluster Manager:** The `nmcluster` process and tool used to monitor and manage the servers for the application failover feature.

Application Failover Basic Setup

To deploy the application failover feature, install NNMi on two servers. This chapter refers to these two NNMi management servers as the **active** and **standby** servers. During normal operation, only the active server is running NNMi services.

The active and standby NNMi management servers are part of a cluster that monitors a heartbeat signal from both of the NNMi management servers. If the active server fails, resulting in the loss of its heartbeat, the standby server becomes the active server.

For application failover to work successfully, the NNMi management servers must meet the following requirements:

- Both NNMi management servers must be running the same type of operating system. For example, if the active server is running an HP-UX operating system, the standby server must also be running an HP-UX operating system.
- Both NNMi management servers must be running the same NNMi version. For example, if NNMi 9.20 is running on the active server, the identical NNMi version, NNMi 9.20, must be on the standby server. The NNMi patch levels must also be the same on both servers.
- The system password must be the same on both NNMi management servers.
- For NNMi installations on Windows operating systems, the `%NnmDataDir%` and `%NnmInstallDir%` system variables must be set to identical values on both servers.

- Both NNMi management servers must be running the same database. For example, both NNMi management servers must be running Oracle or both NNMi management servers must be running the embedded database. You cannot mix the two database types if you plan to use the application failover feature.
- Both NNMi management servers must have identical licensing attributes. For example, the node counts and licensed features must be identical.
- Do not enable application failover until NNMi is in an advanced stage of initial discovery. For more information see [Evaluate Discovery](#) on page 68.

For application failover to function correctly, the active and standby servers must have unrestricted network access to each other. After meeting this condition, complete the steps shown in [Configuring NNMi for Application Failover](#) on page 286. For more information see [NNMi 9.20 and Well-Known Ports](#) on page 429.



Any software that locks files or restricts network access can cause NNMi communication problems. Configure these applications to ignore the files and ports used by NNMi.



During an NNMi 9.20 installation or upgrade, the NNMi installation chooses a network interface for NNMi Cluster communications. The network interface chosen is generally the first non-loopback interface on the system. When the NNMi Cluster is configured, the configuration uses the chosen interface. If you have to adjust the interface, do the following:

- 1 Edit the following file:

- *Windows:*

```
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
```

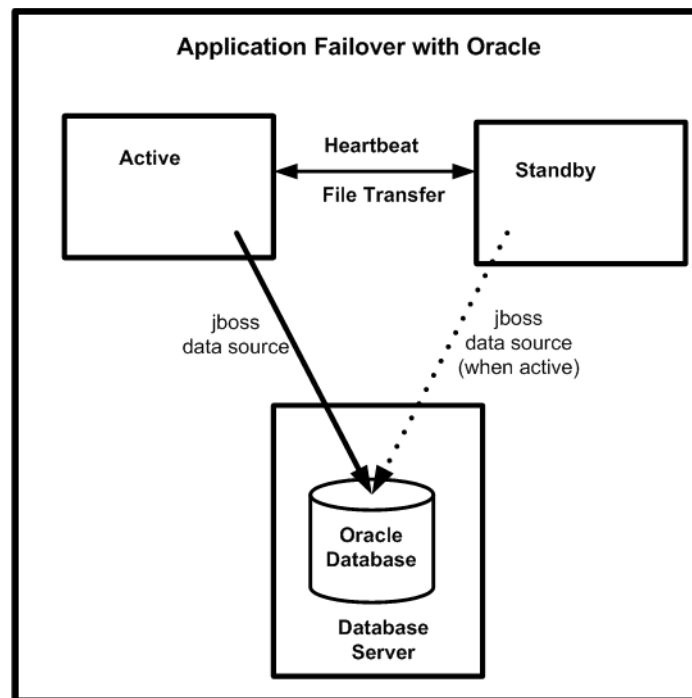
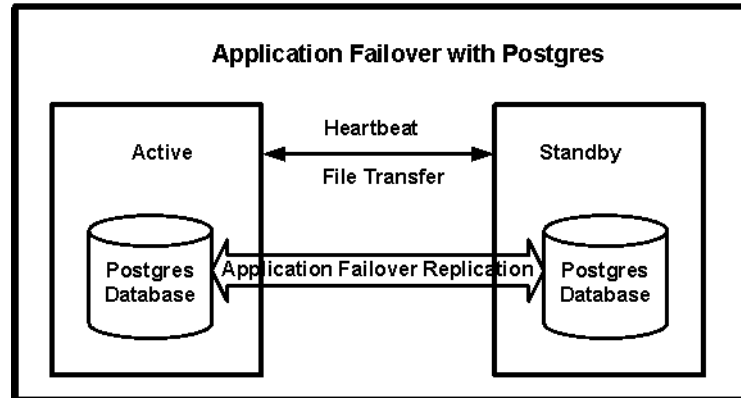
- *UNIX:* \$NnmDataDir/shared/nnm/conf/props/nms-cluster.properties

Parameters in the `nms-cluster.properties` file that have minimum and maximum values are documented, respectively, within the `nms-cluster.properties` file.

- 2 Adjust the `com.hp.ov.nms.cluster.interface` parameter to point to the desired interface.

Configuring NNMi for Application Failover

- 1 Install NNMi on the active server, server X, and the standby server, server Y, as described in the *HP Network Node Manager i Software Interactive Installation Guide*.



- 2 For each license on server X, obtain a similar non-production license for server Y and install it onto server Y as described in [Licensing NNMi](#) on page 121.
- 3 Run the `ovstop` command on each server to shut down NNMi.



If you are using application failover with Oracle as your database, your NNMi processes on the standby server should already be stopped.

- 4 If you are using application failover with Oracle as your database, follow the configuration steps in [Manually Configuring NNMi for Application Failover](#) on page 421.

Configuring your Cluster with the NNMi Cluster Setup Wizard (Embedded Database Users only)

The NNM Cluster Setup Wizard automates the process of configuring a cluster within NNMi for use with Application Failover. The wizard lets you:

- specify and validate cluster nodes
 - define cluster properties and ports
 - merge the `nnm.keystore` and `nnm.truststore` file content for both nodes into a single `nnm.keystore` and `nnm.truststore` file
- 1 Launch the Cluster Setup Wizard by entering the following into a supported Web browser:
 - **`http://<NNMIserv>:<port>/cluster`**
 - `<NNMIserv>` is the value of the NNMi host.
 - `<port>` is the value of the NNMi port.
 - 2 Enter your system **User Name** and **Password**, and then click the **Login** button to sign into NNMi.
 - 3 Enter **Local Hostname** and **Remote Cluster Node** values to define the cluster nodes, and then click **Next**.
 - 4 On the Communication Results page, review the communication verification results. If an error occurs, click **Previous** and fix the problem; otherwise, click **Next**.
 - 5 On the Define Cluster Properties page, enter the **Cluster Name**, define the **Backup Interval** (in hours), and specify whether to enable automatic failover. Click **Next**.
 - 6 On the Define Cluster Ports page, enter **Starting Cluster Port** and **File Transfer Port** values.



The NNMi Cluster uses 4 contiguous ports beginning with the **Starting Cluster Port**.

- 7 Click **Next**.
- 8 Review the summary information provided. Click **Previous** to go back and change configuration information; otherwise, click **Commit** to save the cluster configuration.
- 9 The final summary indicates whether the configuration was successful (indicated by “Success” messages for each of the items). If the configuration is unsuccessful, click **Previous** and fix the problem.

On successful cluster setup, click **Finish**.
- 10 Immediately stop NNMi on both nodes by running the **ovstop** command on both nodes.
- 11 Verify the two nodes are able to cluster by running the **nnmcluster** command on both nodes. If the nodes are not able to cluster, then see [Manually Configuring NNMi for Application Failover](#) on page 421.
- 12 Start NNMi on the desired active node using the **nnmcluster** command. Wait for NNMi to report ACTIVE (see [Manually Configuring NNMi for Application Failover](#) on page 421).
- 13 Start the standby node using the **ovstart** command.

Setting Cluster Communications (Optional)

During installation, NNMi queries all Network Interface Cards (NICs) on the system to find one to use for cluster communications (the first available NIC is chosen). If your system has multiple NICs, you can choose which NIC to use for `nnmcluster` operations by doing the following:

- 1 Run `nnmcluster -interfaces` to list all available interfaces. For more information, see the *nnmcluster* reference page, or the UNIX manpage.
- 2 Edit the following file:
 - Windows:
`%NnmDataDir%\conf\nnm\props\nms-cluster-local.properties`
 - UNIX:
`$(NnmDataDir)/conf/nnm/props/nms-cluster-local.properties`
- 3 Look for a line containing text similar to the following:
`com.hp.ov.nms.cluster.interface =<value>`
- 4 Change the value as desired.
- 5 Save the `nms-cluster-local.properties` file.

Using the Application Failover Feature

Now that you have both NNMi management servers running the cluster manager, with one active node and one standby node, you can use the cluster manager to view the cluster status. The cluster manager has three modes:

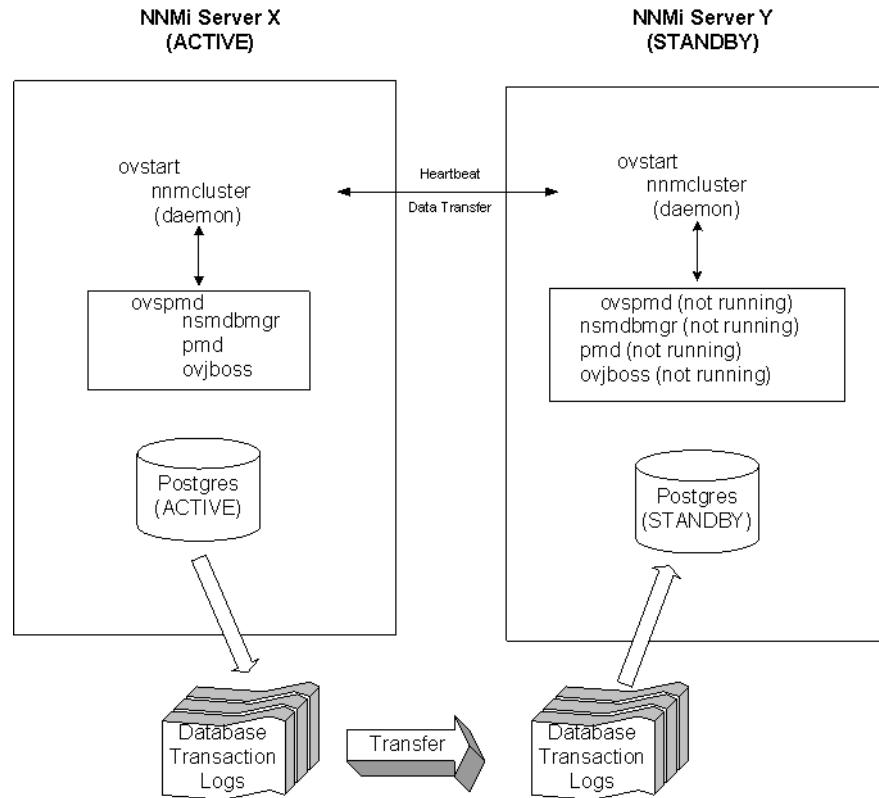
- **daemon mode:** The cluster manager process runs in the background, and uses the `ovstop` and `ovstart` commands to start and stop the NNMi services.
- **interactive mode:** The cluster manager runs an interactive session in which the NNMi administrator can view and change cluster attributes. For example, the NNMi administrator can use this session to enable or disable the application failover feature or shut down the daemon processes.
- **command line mode:** The NNMi administrator views and changes cluster attributes at the command prompt.

For more information, see the *nnmcluster* reference page, or the UNIX manpage.

Application Failover Behavior Using the Embedded Database

Figure 24 shows the application failover configuration for two NNMi management servers using the embedded database. Refer to this figure while reading the rest of this chapter.

Figure 24 Application Failover Configuration (embedded database)



NNMi 9.20 includes a streaming replication feature within application failover whereby database transactions are sent from the active server to the standby server, keeping the standby server in sync with the active server. This eliminates the need for database transaction logs to be imported on the standby server on failover (as was the case in earlier NNMi versions), thus greatly reducing the time needed for the standby server to take over as the active server. Another benefit of this feature is that database backup files are only sent from one node to another if and when needed, and given the regular transmission of database transaction files, the need for sending large database backup files should be infrequent.



For both the active and standby nodes, if you have a firewall enabled, ensure that the port you are using for the embedded database (port 5432 by default) is open. This port is set in the following file:

Windows: %NNM_CONF%\nrm\props\nms-local.properties

UNIX: \$NNM_CONF/nrm/props/nms-local.properties

After you start both the active and standby nodes, the standby node detects the active node, requests a database backup from the active node, but does not start NNMi services. This database backup is stored as a single Java-ZIP file. If the standby node already has a ZIP file from a previous cluster-connection, and NNMi finds that the file is already synchronized with the active server, the file is not retransmitted.

While both the active and standby nodes are running, the active node periodically sends database transaction logs to the standby node. You can modify the frequency of this data transfer by changing the value of the `com.hp.ov.nms.cluster.timeout.archive` parameter in the `nms-cluster.properties` file. These transaction logs accumulate on the standby node, and are available on the standby node any time it needs to become active.

When the standby node receives a full database backup from the active node, it places the information into its embedded database. It also creates a `recovery.conf` file to inform the embedded database that it should consume all received transaction logs before it becomes available to other services.

If the active node becomes unavailable for any reason, the standby node becomes active by running an `ovstart` command to start the NNMi services. The standby NNMi management server imports the transaction logs before starting the remaining NNMi services.

If the active NNMi system fails, the standby system begins discovery and polling activities. This transition keeps NNMi monitoring and polling your network while you diagnose and repair the failed system.



Note the following:

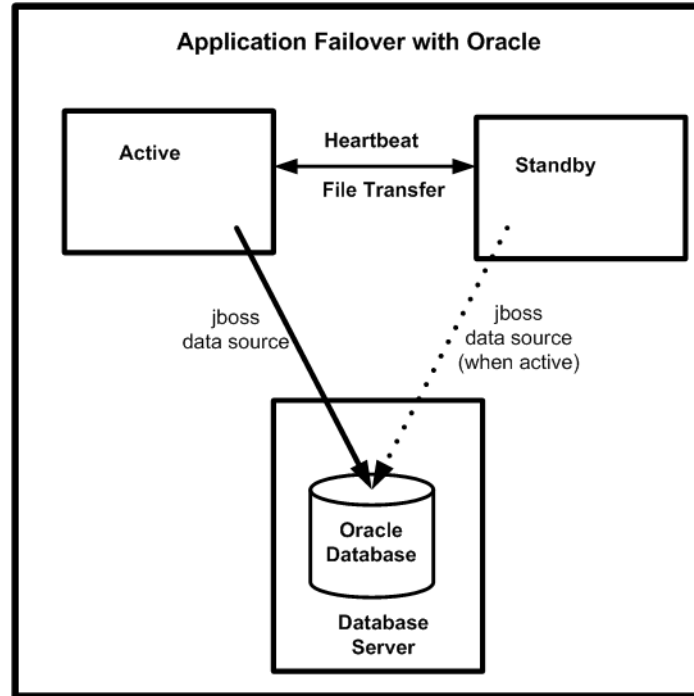
- Updates to Status and Incidents could be delayed as NNMi resynchronizes following an application failover.
- If you see the following message during this resynchronization, it does not indicate a problem:

The Causal Engine's large queue size is causing delayed updates to Status and Incidents. This could be due to resynchronization following an upgrade, application failover, restore from backup, or a manual resynchronization.
- Do not stop NNMi during this resynchronization. To ensure resynchronization has completed, keep NNMi running for several hours following the application failover.

Application Failover Behavior Using an Oracle Database

Figure 25 shows the application failover configuration for two NNMi management servers using an Oracle database. Refer to this figure while reading the rest of this chapter.

Figure 25 Application Failover Configuration (Oracle database)



If the active node becomes unavailable for any reason, the standby node becomes active by running an **ovstart** command to start the NNMi services.

If the active NNMi system fails, the standby system begins discovery and polling activities. This transition keeps NNMi monitoring and polling your network while you diagnose and repair the failed system.



Note the following:

- Updates to Status and Incidents could be delayed as NNMi resynchronizes following an application failover.
- If you see the following message during this resynchronization, it does not indicate a problem:

The Causal Engine's large queue size is causing delayed updates to Status and Incidents. This could be due to resynchronization following an upgrade, application failover, restore from backup, or a manual resynchronization.
- Do not stop NNMi during this resynchronization. To ensure resynchronization has completed, keep NNMi running for several hours following the application failover.

Application Failover Scenarios

There are several possible problems that can cause the active NNMi management server to stop sending heartbeats, and to initiate a failover:

- Scenario 1: The active NNMi management server fails.
- Scenario 2: The system administrator shuts down or reboots the active NNMi management server.
- Scenario 3: The NNMi administrator shuts down the cluster.
- Scenario 4: The network connection between the active and the standby NNMi management servers fails.

In scenario 4, both NNMi management servers run in the active state. When the network device comes back online, the two NNMi management servers automatically negotiate which node should become the new active node.

Additional `ovstart` and `ovstop` Options

When you use the `ovstop` and `ovstart` commands on NNMi management servers configured for application failover, NNMi runs the following commands:

- `ovstart`: `nnmcluster -daemon`
- `ovstop`: `nnmcluster -disable -shutdown`



If you run an `ovstop` command, NNMi does not failover to the standby node. HP designed the `ovstop` command to support temporary maintenance stoppages. To manually initiate a failover, use the `-failover` option with the `ovstop` command. For more information, see the `ovstop` reference page, or the UNIX manpage.

The following options to the `ovstop` command apply to NNMi management servers configured in an application failover cluster:

- `ovstop -failover`: This command stops the local daemon-mode cluster process and forces a failover to the standby NNMi management server. If the failover mode was previously disabled, it is re-enabled. This command is equivalent to: `nnmcluster -enable -shutdown`
- `ovstop -nofailover`: This command disables failover mode and then stops the local daemon-mode cluster process. No failover occurs. This command is equivalent to: `nnmcluster -disable -shutdown`
- `ovstop -cluster`: This command stops both the active and standby nodes, removing them both from the cluster. This command is equivalent to: `nnmcluster -halt`



If you run the `shutdown` command on NNMi management servers running UNIX operating systems, the `ovstop` command runs automatically and disables application failover. That might not be your desired result. To control application failover during maintenance windows, use the `nnmcluster -acquire` and `nnmcluster -relinquish` commands to set the active and standby nodes the way you want them before running the shutdown command. For more information see the `nnmcluster` reference page, or the UNIX manpage.

Application Failover Incidents

Any time the `nnmcluster` process or someone using the `nnmcluster` command starts a node as active, NNMi generates one of the following incidents:

- *NnmClusterStartup*: The NNMi cluster was started, and no active node was present. Therefore the node was started in the active state. This incident has a NORMAL severity.
- *NnmClusterFailover*: The NNMi cluster detected a failure of the active node. The standby node was then enabled and NNMi services started on the new active node. This incident has a MAJOR severity.

Returning to the Original Configuration Following a Failover

Suppose the active node fails and the standby node is functioning as the active node. After you fix the problem with the former active node, run the following command on the desired active node to return to the original configuration: `nnmcluster -acquire`. For more information, see the *nnmcluster* reference page, or the UNIX manpage.

NNM iSPIs and Application Failover

You can use the application failover feature for a Smart Plug-in (iSPI) that you deploy along with NNMi if the deployment meets the following requirements:

- The NNM iSPI runs on the NNMi management server.
- The NNM iSPI uses the same embedded database instance as NNMi.

The NNM iSPI Performance for Metrics and the NNM iSPI Performance for Traffic are exceptions to this description. If you plan to configure the NNMi application failover feature, you must install these iSPIs on dedicated servers. In this case, the iSPIs automatically connect to the new NNMi management server after failover occurs. As part of NNMi application failover configuration, run the enablement script for the NNM iSPI Performance for Metrics or the NNM iSPI Performance for Traffic on each NNMi management server in the cluster.

For more information, see *Support for Application Failover* in the NNM iSPI Performance for Metrics, the NNM iSPI Performance for QA, or the NNM iSPI Performance for Traffic help.

NNM iSPI Installation Information

To install an NNM iSPI on an NNMi management server that is already part of an application failover cluster, do the following:

- 1 As a precaution, run the **nnmconfigexport.ovpl** script on both the active and standby NNMi management servers before proceeding. For information, see [Best Practice: Save the Existing Configuration](#) on page 34.
- 2 As a precaution, back up the NNMi data on both the active and standby NNMi management servers before proceeding. For information, see [Backup Scope](#) on page 385.
- 3 Embedded database only: As a precaution, on the active NNMi management server run the **nnmcluster -dbsync** command and wait for the command to complete.
- 4 On the standby NNMi management server, run the following command:
nnmcluster -shutdown
- 5 Edit the following file on the standby NNMi management server.:
 - *Windows*:
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
 - *UNIX*: \$NnmDataDir/shared/nnm/conf/props/nms-cluster.properties
- 6 Comment out the com.hp.ov.nms.cluster.name option and save the file.
- 7 Run the **ovstart** command on the standby NNMi management server. This brings up NNMi services in the standalone (unclustered) state.
- 8 Install the NNM iSPI on the standby NNMi management server as described in the iSPI installation guide.
- 9 Run the **nnmcluster -halt** command on the active NNMi management server.

- 10 Edit the following file on the active NNMi management server:
 - *Windows*:
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
 - *UNIX*: \$NnmDataDir/shared/nnm/conf/props/nms-cluster.properties
- 11 Comment out the `com.hp.ov.nms.cluster.name` option and save the file.
- 12 Run the **ovstart** command on the active NNMi management server. This brings up NNMi services in the standalone (unclustered) state.
- 13 Install the NNM iSPI on the active NNMi management server as described in the iSPI installation guide.
- 14 Edit the following file on **both** the active and standby NNMi management servers:
 - *Windows*:
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
 - *UNIX*: \$NnmDataDir/shared/nnm/conf/props/nms-cluster.properties
- 15 Uncomment the `com.hp.ov.nms.cluster.name` option and save each file.
- 16 Run the **ovstart** command on the active NNMi management server.
- 17 Wait a few minutes for the active NNMi management server to become the first active node in the cluster. Run the **nnmcluster -display** command on the active NNMi management server and search the displayed results for the term `ACTIVE` as in `ACTIVE_NNM_STARTING` or `ACTIVE_SomeOtherState`. Do not continue with [step 18](#) until you know that the active NNMi management server is the active node.
- 18 Run the **ovstart** command on the standby NNMi management server.

Integrated Applications

When other HP Software or third-party products are integrated with NNMi, the affect of NNMi application failover on an integration depends on how a product communicates with NNMi. For more information, see the appropriate integration document.

If an integrated product must be configured with information about the NNMi management server, the following information applies:

- If long-term, you can update the NNMi management server information within the integrating product configuration. For more information, see the appropriate integration document.
- If the outage appears to be temporary, you can resume using the integrating product after server X returns to service. To return server X to service, follow these steps:
 - 1 On server X, run the following command:

```
nnmcluster -daemon
```

Server X joins the cluster and assumes a standby state.

- 2 On server X, run the following command:

```
nnmcluster -acquire
```

Server X changes to the active state.

If you anticipate that the original server X will be out of service for a longer time, you can update the NNMi management server IP address within the integrating product. For instructions on how to modify the IP address field, see the integrating product documentation.

Disabling Application Failover

Suppose you configure application failover, use it for a few days, then decide to completely disable it. The following information explains how to completely disable application failover. Complete the following instructions, including actions on both the active and standby NNMi management servers configured in the application failover cluster.

- 1 Run **nnmcluster -enable** command on the *active* NNMi management server.
- 2 Run the **nnmcluster -shutdown** command on the *active* NNMi management server.
- 3 Wait a few minutes for the old standby NNMi management server to become the new active NNMi management server.
- 4 Run the **nnmcluster -display** command on the new active (old standby) NNMi management server.
- 5 Search the displayed results for the ACTIVE_NNM_RUNNING status. Repeat [step 4](#) until you see the ACTIVE_NNM_RUNNING status.
- 6 Run the **nnmcluster -shutdown** command on the new active (old standby) NNMi management server.
- 7 Run the **nnmcluster -display** command repeatedly on the new active (old standby) until you no longer see a DAEMON process.
- 8 Edit the following file both NNMi management servers configured in the cluster:
 - *Windows:*
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
 - *UNIX:* \$NnmDataDir/shared/nnm/conf/props/nms-cluster.properties
- 9 Comment out the com.hp.ov.nms.cluster.name option on both NNMi management servers and save each file.
- 10 Edit the following file on both NNMi management servers:
 - *Windows:*
%NnmDataDir%\shared\nnm\databases\Postgres\postgresql.conf
 - *UNIX:* \$NnmDataDir/shared/nnm/databases/Postgres/postgresql.conf

- 11 Remove the following lines, which are automatically added by application failover. This is an example of what these lines could look like. These lines might look slightly different on your server.

```
# The following lines were added by the NNM cluster.
archive_command = ...
archive_timeout = 900
max_wal_senders = 4
archive_mode = 'on'
wal_level = 'hot_standby'
hot_standby = 'on'
wal_keep_segments = 500
listen_addresses = 'localhost,16.78.61.68'
```

Make sure to save your changes.

- 12 If these are Windows NNMi management servers, navigate to the Services (Local) console and do the following on each server:
- Set the Startup type for the HP NNM Cluster Manager to Disabled.
 - Set the Startup type for the HP OpenView Process Manager to Automatic.
- 13 Create the following trigger file, which tells Posgres to stop running in standby mode and to start fully running:

```
NnmDataDir/tmp/postgresTriggerFile
```

- 14 Run the **ovstart** command on the former active NNMi management server only. In the application failover configuration, this is the NNMi management server that has a permanent NNMi license.
- 15 If you were using a non-production license on the former standby server. Do not run the **ovstart** command on the former standby NNMi management server. In the application failover configuration, this is the NNMi management server that has a non-production license. To run this NNMi management server as a standalone server, you must purchase and install a permanent license. For more information, see [Licensing NNMi](#) on page 121.
- 16 If both NNMi management servers start successfully, then remove the following directory from both the standby and active NNMi management servers:
- Windows:** %NnmDataDir%\shared\nnm\databases\Postgres_standby
 - UNIX:** \$NnmDataDir/shared/nnm/databases/Postgres_standby



This directory is a default directory and is the value of the `com.hp.ov.nms.cluster.archivedir` parameter located in the `nms-cluster.properties` file. These instructions assume you did not change this value. If you changed the value of the `com.hp.ov.nms.cluster.archivedir` parameter in the `nms-cluster.properties` file, then remove the directory that equates to the new value.

- 17 Remove the following directory from both the standby and active NNMi management servers:
- Windows:** %NnmDataDir%\shared\nnm\databases\Postgres.OLD
 - UNIX:** \$NnmDataDir/shared/nnm/databases/Postgres.OLD

Administrative Tasks and Application Failover

The following information explains how to effectively manage application failover when doing administrative tasks such as patching and restarting NNMi management servers.

Application Failover and Upgrading to NNMi 9.20

If you plan to upgrade an earlier version of NNMi that is running in an NNMi application failover configuration, follow the steps in the appropriate section below based on the database you are using.

Embedded Database

To upgrade NNMi management servers configured for application failover and using the embedded database, follow these steps:

- 1 As a precaution, run the **nnmconfigexport.ovpl** script on both the active and standby NNMi management servers before proceeding. For information, see [Best Practice: Save the Existing Configuration](#) on page 34.

As a precaution, back up your NNMi data on both the active and standby NNMi management servers before proceeding. For information, see [Backup Scope](#) on page 385.

- 2 Complete the following steps on the active NNMi management server. Note that NNMi must be running for the following **nnmcluster** steps to work. Completing these steps will speed up the standby NNMi management server startup shown in [step 6](#) on page 299:
 - a Run the **nnmcluster** command.
 - b After NNMi prompts you, type **dbsync**, then press Enter. Review the displayed information to make sure it includes the following messages:

ACTIVE_DB_BACKUP: This means that the active NNMi management server is performing a new backup.

ACTIVE_NNM_RUNNING: This means that the active NNMi management server completed the backup referred to by the previous message.

STANDBY_RECV_DBZIP: This means that the standby NNMi management server is receiving a new backup from the active NNMi management server.

STANDBY_READY: This means that the standby NNMi management server is ready to perform if the active NNMi management server fails.
 - c Run **exit** or **quit** to stop the interactive **nnmcluster** process you started in [step a](#).

- 3 Run the **nnmcluster -shutdown** command on the standby NNMi management server. This shuts down all **nnmcluster** processes on the standby NNMi management server.
- 4 To verify there are no **nnmcluster** nodes running on the standby NNMi management server, *complete the following steps on the standby NNMi management server.*
 - a Run the **nnmcluster** command.

- b Verify that there are no (LOCAL) `nnmcluster` nodes present except the one marked (SELF). There might be one or more (REMOTE) nodes present.
 - c Run **exit** or **quit** to stop the interactive `nnmcluster` process you started in step a.
- 5 Complete the following steps on the standby NNMi management server to temporarily disable application failover:
- a Edit the following file:
 - Windows: `%NNM_SHARED_CONF%\props\nms-cluster.properties`
 - UNIX: `$NNM_SHARED_CONF/props/nms-cluster.properties`
 - b Comment out the `com.hp.ov.nms.cluster.name` parameter.
 - c Save your changes.
- 6 Start, then stop processes on the standby NNMi management server.
- a Run the **ovstart** command on the standby NNMi management server. Running the **ovstart** command causes the standby NNMi management server to import the transaction logs from the active NNMi management server.
 - b After the **ovstart** command completes, run the **ovstatus -v** command. All NNMi services should show the state `RUNNING`.
 - c Run the **ovstop** command on the standby NNMi management server.
- 7 Upgrade the standby NNMi management server to NNMi 9.20 using the instructions located in the *HP Network Node Manager i Software Interactive Installation Guide*.



You must upgrade all of the iSPIs that you have installed on the standby NNMi management server to iSPI versions that support NNMi 9.20.

You now have the former active NNMi management server running NNMi 9.0x or 9.1x and the former standby NNMi management server running NNMi 9.20. You have both of these NNMi management servers running independently with no database synchronization. That means you have both NNMi management servers monitoring the network in parallel. Do not leave these NNMi management servers in this configuration for more than a few hours, as this configuration is a violation of the non-production license installed on the former standby node.

To complete the upgrade, and remedy this situation, select a time to upgrade the former active node to NNMi 9.20. Have the operators temporarily use the former standby node to monitor the network while you complete the upgrade.

The remainder of this procedure assumes you plan to retain the database information from the former active node and discard the database information from the former standby node.

- 8 Run the **nnmcluster -halt** command on the former active NNMi management server.
- 9 To verify there are no `nnmcluster` nodes running on the former active NNMi management server, complete the following steps on the former active NNMi management server.
 - a Run the **nnmcluster** command.

- b Verify that there are no (LOCAL) nnmcluster nodes present except the one marked (SELF). There might be one or more (REMOTE) nodes present.
 - c Run **exit** or **quit** to stop the interactive nnmcluster process you started in step a.
- 10 Complete the following steps on the former active NNMi management server to temporarily disable application failover:

- a Edit the following file:
 - Windows: %NNM_SHARED_CONF%\props\nms-cluster.properties
 - UNIX: \$NNM_SHARED_CONF/props/nms-cluster.properties
- b Comment out the `com.hp.ov.nms.cluster.name` parameter.

Upgrade the former active NNMi management server to NNMi 9.20 using the instructions located in the *HP Network Node Manager i Software Interactive Installation Guide*.

▶ You must upgrade all of the iSPIs that you have installed on the former active NNMi management server to iSPI versions that support NNMi 9.20.

Now you have two servers running NNMi 9.20, but they are still independent since the databases are not synchronized.

- 11 Complete the following steps on the former active NNMi management server:
- a Run the **ovstop** command.
 - b Edit the following file:
 - Windows: %NNM_SHARED_CONF%\props\nms-cluster.properties
 - UNIX: \$NNM_SHARED_CONF/props/nms-cluster.properties
 - c Type in the value of the `com.hp.ov.nms.cluster.name` parameter.

▶ The NNMi upgrade procedure does not preserve commented-out properties. Therefore, you must retype the cluster name.

- d Uncomment the `com.hp.ov.nms.cluster.name` parameter.
- e Save your changes.

- 12 Run either the **ovstart** or **nnmcluster -daemon** command on the former active NNMi management server. It is now the active node.

- 13 Instruct the operators to begin using the active node to monitor the network.

▶ The former standby NNMi management server discards all of the database activity occurring during the maintenance window, from [step 8](#) on page 299 through [step 12](#) on page 300.

- 14 Complete the following steps on the former standby NNMi management server:
- a Run the **ovstop** command.
 - b Edit the following file:
 - Windows: %NNM_SHARED_CONF%\props\nms-cluster.properties
 - UNIX: \$NNM_SHARED_CONF/props/nms-cluster.properties
 - c Type in the value of the `com.hp.ov.nms.cluster.name` parameter.
 - d Uncomment the `com.hp.ov.nms.cluster.name` parameter.

- e Save your changes.
- 15 Run either the **ovstart** or **nnmcluster -daemon** command on the former standby NNMi management server.

This NNMi management server becomes the standby node, and receives a copy of the database from the active node.

- 16 If you installed the NNM iSPI Performance for QA, the NNM iSPI Performance for Metrics, or the NNM iSPI Performance for Traffic; are using the application failover feature; and completed the upgrade process shown above, run the NNM iSPI enablement script for each NNM iSPI on both the active and standby NNMi management servers. The path to the NNM iSPI enablement script is as follows:
- *Windows:* %NNMInstallDir%\bin\nmenableperfspi.ovpl
 - *UNIX:* /opt/OV/bin/nmenableperfspi.ovpl

Oracle Database



You must upgrade NNMi management servers separately because two NNMi management servers cannot be simultaneously connected to the same Oracle database.

To upgrade NNMi management servers configured for application failover and using an Oracle database, follow these steps:

- 1 As a precaution, run the **nnmconfigexport.ovpl** script on both the active and standby NNMi management servers before proceeding. For information, see [Best Practice: Save the Existing Configuration](#) on page 34.
- 2 As a precaution, back up your NNMi data on both the active and standby NNMi management servers before proceeding. For information, see [Backup Scope](#) on page 385.
- 3 Run the **nnmcluster -halt** command on the standby NNMi management server. This shuts down all **nnmcluster** processes on both the active and standby NNMi management server.
- 4 To verify there are no **nnmcluster** nodes running on either the active or standby NNMi management server, *complete the following steps on the standby NNMi management server.*
 - a Run the **nnmcluster** command.
 - b Verify that the only **nnmcluster** node present is one marked (SELF).
 - c Run **exit** or **quit** to stop the interactive **nnmcluster** process you started in step a.
- 5 *Complete the following steps on the standby NNMi management server to temporarily disable application failover:*
 - a Edit the following file:
 - *Windows:* %NNM_SHARED_CONF%\props\nms-cluster.properties
 - *UNIX:* \$NNM_SHARED_CONF/props/nms-cluster.properties
 - b Comment out the `com.hp.ov.nms.cluster.name` parameter.
 - c Save your changes.

- 6 Upgrade the standby NNMi management server to NNMi 9.20 using the instructions located in the *HP Network Node Manager i Software Interactive Installation Guide*.



You must upgrade all of the iSPIs that you have installed on the standby NNMi management server to iSPI versions that support NNMi 9.20.

You now have the former standby NNMi management server with NNMi 9.20 installed, and the former active NNMi management server with NNMi 9.0x or 9.1x installed.

- 7 Run the **ovstop** command on the former standby NNMi management server to disconnect the NNMi management server from the Oracle database.
- 8 *Complete the following steps on the former active NNMi management server to temporarily disable application failover:*
 - a Edit the following file:
 - *Windows:* %NNM_SHARED_CONF%\props\nms-cluster.properties
 - *UNIX:* \$NNM_SHARED_CONF/props/nms-cluster.properties
 - b Comment out the `com.hp.ov.nms.cluster.name` parameter.
- 9 Upgrade the former active NNMi management server to NNMi 9.20 using the instructions located in the *HP Network Node Manager i Software Interactive Installation Guide*.



You must upgrade all of the iSPIs that you have installed on the former active NNMi management server to iSPI versions that support NNMi 9.20.

Now you have two servers with NNMi 9.20 installed.

- 10 Complete the following steps on the former active NNMi management server:
 - a Run the **ovstop** command.
 - b Edit the following file:
 - *Windows:* %NNM_SHARED_CONF%\props\nms-cluster.properties
 - *UNIX:* \$NNM_SHARED_CONF/props/nms-cluster.properties
 - c Type in the value of the `com.hp.ov.nms.cluster.name` parameter.



The NNMi upgrade procedure does not preserve commented-out properties. Therefore, you must retype the cluster name.

- d Uncomment the `com.hp.ov.nms.cluster.name` parameter.
 - e Save your changes.
- 11 Run the **ovstart** or **nnmcluster -daemon** command on the former active NNMi management server. It is now the active node.
- 12 Complete the following steps on the former standby NNMi management server:
 - f Edit the following file:
 - *Windows:* %NNM_SHARED_CONF%\props\nms-cluster.properties
 - *UNIX:* \$NNM_SHARED_CONF/props/nms-cluster.properties
 - g Type in the value of the `com.hp.ov.nms.cluster.name` parameter.
 - h Uncomment the `com.hp.ov.nms.cluster.name` parameter.

- i Save your changes.
- 13 Run either the **ovstart** or **nnmcluster -daemon** command on the former standby NNMi management server.

This NNMi management server becomes the standby node.

- 14 If you installed the NNM iSPI Performance for QA, the NNM iSPI Performance for Metrics, or the NNM iSPI Performance for Traffic; are using the application failover feature; and completed the upgrade process shown above, run the NNM iSPI enablement script for each NNM iSPI on both the active and standby NNMi management servers. The path to the NNM iSPI enablement script is as follows:
 - *Windows:* %NNMInstallDir%\bin\nnmenableperfspi.ovpl
 - *UNIX:* /opt/OV/bin/nnmenableperfspi.ovpl

Application Failover and NNMi Patches

Both NNMi management servers must be running the same NNMi version and patch level. To add patches to the active and standby NNMi management servers, use one of the following procedures:

- [Applying Patches for Application Failover \(Shut Down Both Active and Standby\)](#)
Use this procedure when you are not concerned with an interruption in network monitoring.
- [Applying Patches for Application Failover \(Keep One Active NNMi Management Server\)](#)
Use this procedure when must avoid any interruptions in network monitoring.

Applying Patches for Application Failover (Shut Down Both Active and Standby)

This procedure results in both NNMi management servers being non-active for some period of time during the patch process. To apply patches to the NNMi management servers configured for application failover, follow these steps:

- 1 As a precaution, run the **nnmconfigexport.ovpl** script on both the active and standby NNMi management servers before proceeding. For information, see [Best Practice: Save the Existing Configuration](#) on page 34.
- 2 As a precaution, back up your NNMi data on both the active and standby NNMi management servers before proceeding. For information, see [Backup Scope](#) on page 385.
- 3 As a precaution, on the active NNMi management server, do the following steps:
 - a Run the **nnmcluster** command.
 - b Embedded database only: After NNMi prompts you, type **dbsync**, then press Enter. Review the displayed information to make sure it includes the following messages:
 - ACTIVE_DB_BACKUP: This means that the active NNMi management server is performing a new backup.
 - ACTIVE_NNM_RUNNING: This means that the active NNMi management server completed the backup referred to by the previous message.
 - STANDBY_READY: This shows the previous status of the standby NNMi management server.

`STANDBY_RECV_DBZIP`: This means that the standby NNMi management server is receiving a new backup from the active NNMi management server.
`STANDBY_READY`: This means that the standby NNMi management server is ready to perform if the active NNMi management server fails.

- 4 Run the **nnmcluster -halt** command on the active NNMi management server. This shuts down all `nnmcluster` processes on both the active and standby NNMi management servers.
- 5 To verify there are no `nnmcluster` nodes running on either server, *complete the following steps on both the active and standby NNMi management servers.*
 - a Run the **nnmcluster** command.
 - b Verify that there are no `nnmcluster` nodes present except the one marked `(SELF)`.
 - c Run **exit** or **quit** to stop the interactive `nnmcluster` process you started in step a.
- 6 On the active NNMi management server, comment out the `com.hp.ov.nms.cluster.name` parameter in the `nms-cluster.properties` file.
 - a Edit the following file:
 - *Windows*: `%NNM_SHARED_CONF%\props\nms-cluster.properties`
 - *UNIX*: `$NNM_SHARED_CONF/props/nms-cluster.properties`
 - b Comment out the `com.hp.ov.nms.cluster.name` parameter.
 - c Save your changes.
- 7 Apply the NNMi patch to the active NNMi management server using the instructions provided with the patch.
- 8 On the active NNMi management server, uncomment the `com.hp.ov.nms.cluster.name` parameter in the `nms-cluster.properties` file.
 - a Edit the following file:
 - *Windows*: `%NNM_SHARED_CONF%\props\nms-cluster.properties`
 - *UNIX*: `$NNM_SHARED_CONF/props/nms-cluster.properties`
 - b Uncomment the `com.hp.ov.nms.cluster.name` parameter.
 - c Save your changes.
- 9 Run the **ovstart** command on the active NNMi management server.
- 10 Verify that the patch installed correctly on the active NNMi management server by viewing information on the **Product** tab of the **Help > System Information** window in the NNMi console.
- 11 Run the **nnmcluster -dbsync** command to create a new backup.
- 12 On the standby NNMi management server, comment out the `com.hp.ov.nms.cluster.name` parameter in the `nms-cluster.properties` file as shown in [step a](#) on page 304 through [step c](#) on page 304.
- 13 Apply the NNMi patch to the standby NNMi management server.
- 14 On the standby NNMi management server, uncomment the `com.hp.ov.nms.cluster.name` parameter in the `nms-cluster.properties` file as shown in [step a](#) on page 304 through [step c](#) on page 304.
- 15 Run the **ovstart** command on the standby NNMi management server.

- 16 If you installed the NNM iSPI Performance for QA, the NNM iSPI Performance for Metrics, or the NNM iSPI Performance for Traffic; are using the application failover feature; and completed the patch process shown above, run the NNM iSPI enablement script for each NNM iSPI on both the active and standby NNMi management servers.

Applying Patches for Application Failover (Keep One Active NNMi Management Server)

This procedure results in one NNMi management server always being active during the patch process.



This process results in continuous monitoring of the network, however NNMi loses the transaction logs occurring during this patch process.

To apply NNMi patches to the NNMi management servers configured for application failover, follow these steps:

- 1 As a precaution, run the **nnmconfigexport.ovpl** script on both the active and standby NNMi management servers before proceeding. For information, see [Best Practice: Save the Existing Configuration](#) on page 34.
- 2 As a precaution, back up your NNMi data on both the active and standby NNMi management servers before proceeding. For information, see [Backup Scope](#) on page 385.
- 3 Run **nnmcluster** on one of the nodes.
- 4 Enter **dbsync** on the NNMi management server used in the previous step to synchronize the two databases.



The **dbsync** option works on an NNMi management server using the embedded database. Do not use the **dbsync** option on an NNMi management server configured to use an Oracle database.

- 5 Wait until the active NNMi management server reverts to `ACTIVE_NNM_RUNNING` and the standby NNMi management server reverts to `STANDBY_READY`. before continuing.
- 6 Exit or quit from the **nnmcluster** command.
- 7 Stop the cluster on the standby NNMi management server by running the following command on the standby NNMi management server:
nnmcluster -shutdown
- 8 Make sure the following processes and services terminate before continuing:
 - postgres
 - ovjboss
- 9 Make sure the **nnmcluster** process terminates before continuing. If the **nnmcluster** process will not terminate, manually kill the **nnmcluster** process only as a last resort.
- 10 Edit the following file on the standby NNMi management server:

Windows: %nnmDataDir%\shared\nnm\conf\props\nms-cluster.properties

UNIX: \$nnmDataDir/shared/nnm/conf/props/nms-cluster.properties

- 11 Comment out the cluster name by placing a # at the front of the line, then save your changes:
#com.hp.ov.nms.cluster.name = NNMiCluster
- 12 Install the NNMi patch on the standby NNMi management server.
- 13 At this point, the standby NNMi management server is patched but stopped, and the active NNMi management server is unpatched but running. Stop the active NNMi management server and immediately bring the standby NNMi management server online to monitor your network.
- 14 Shut down the cluster on the active NNMi management server by running the following command on the active NNMi management server:
nnmcluster -halt
- 15 Make sure the `nnmcluster` process terminates. If it does not terminate within a few minutes, manually kill the `nnmcluster` process.
- 16 On the standby NNMi management server, uncomment the cluster name from the `nms-cluster.properties` file.
- 17 Start the cluster on the standby NNMi management server by running the following command on the standby NNMi management server:
nnmcluster -daemon
- 18 Install the NNMi patch on the active NNMi management server.
- 19 At this point, the previous active NNMi management server is patched but offline. Bring it back into the cluster (as the standby NNMi management server) by performing the following:
 - a Uncomment the entry in the `nms-cluster.properties` file on the active NNMi management server.
 - b Start the active NNMi management server using the following command:
nnmcluster -daemon
- 20 To monitor the progress, run the following command on both the active and standby NNMi management servers:
nnmcluster

 Wait until the previous active NNMi management server finishes retrieving the database from the previous standby NNMi management server.
- 21 After the previous active NNMi management server displays `STANDBY_READY`, run the following command on the previous active NNMi management server:
nnmcluster -acquire
- 22 If you installed the NNM iSPI Performance for QA, the NNM iSPI Performance for Metrics, or the NNM iSPI Performance for Traffic; are using the application failover feature; and completed the patch process shown above, run the NNM iSPI enablement script for each NNM iSPI on both the active and standby NNMi management servers.

Application Failover and Restarting the NNMi Management Servers

You can restart the standby NNMi management server at any time with no special instructions. If you restart both the standby and active NNMi management servers, restart the active NNMi management server first.

To restart either the active or the standby NNMi management server, do the following.

- 1 Run the **nnmcluster -disable** command on the NNMi management server to disable the application failover feature.
- 2 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.
- 3 Run the **nnmcluster -enable** command on the NNMi management server to enable the application failover feature.

Application Failover Control after a Communication Failure

After a communication failure between two remote nodes is resolved, JGroups determines which member of the cluster becomes the controller based on the lowest IP address. The controller determines which node is the Active member (this node is always the node on which the controller is running). NNMi starts on the Active member. This functionality is subject to change in future releases.

Application Failover and Recovery from a Previous Database Backup (Embedded Database Only)

To restore your NNMi database from an original backup when active and standby NNMi management servers are configured for application failover, follow these steps:

- 1 Run the **nnmcluster -halt** command on the active NNMi management server.
- 2 Delete or move the following directory on both the active and standby NNMi management servers:
 - *Windows:* %NnmDataDir%\shared\nnm\databases\Postgres_standby
 - *UNIX:* \$NnmDataDir/shared/nnm/databases/Postgres_standby
- 3 Restore the database on the active NNMi management server:
 - a Modify the following file to comment out the cluster name:
 - *Windows:*
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
 - *UNIX:* \$NnmDataDir/shared/nnm/conf/props\nms-cluster.properties
 - b Restore the database as normal. See [Restoring NNMi Data](#) on page 388.
 - c Run the **ovstop** command on the active NNMi management server.

- d Modify the following file to uncomment the cluster name:
 - *Windows*:
%NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
 - *UNIX*: \$NnmDataDir/shared/nnm/conf/props/
nms-cluster.properties
- 4 Run the **ovstart** command on the active NNMi management server.
- 5 Wait until the active NNMi management server generates a new backup. To verify that this step is complete, run the **nnmcluster -display** command and look for an ACTIVE_NNM_RUNNING message.
- 6 Run the **ovstart** command on the standby NNMi management server. The standby NNMi management server copies and extracts the new backup. To verify that this step is complete, run the **nnmcluster -display** command and look for a STANDBY_READY message.

Network Latency/Bandwidth Considerations

NNMi application failover works by exchanging a continuous heartbeat signal between the nodes in the cluster. It uses this same network channel for exchanging other data files such as the NNMi embedded database, database transaction logs, and other NNMi configuration files. HP recommends using a high performance, low latency connection for NNMi application failover when implementing it over a WAN (wide area network).

The NNMi embedded database can become quite large, and can grow to 1GB or more even though this file is always compressed. Also, NNMi generates hundreds, or even thousands, of transaction logs during the built-in backup interval (a configuration parameter that defaults to six hours). Each transaction log can be several megabytes, up to a maximum size of 16 MB. (These files are also compressed). Example data collected from an HP test environments is shown here:

Number of nodes managed: 15,000

Number of interfaces: 100,000

Time to complete spiral discovery of all expected nodes: 12 hours

Size of database: 850MB (compressed)

During initial discovery: ~10 transaction logs per minute (peak of ~15/min)

10 TxLogs/minute X 12 hours = 7200 TxLogs @ ~10MB = ~72GB

This is a lot of data to send over the network. If the network between the two nodes is unable to keep up with the bandwidth demands of NNMi application failover, the standby node can fall behind in receiving these database files. This could result in a larger window of potential data loss if the active server fails.

Similarly, if the network between the two nodes has a high latency or poor reliability, this could result in a *false* loss-of-heartbeat between the nodes. For example, this can happen when the heartbeat signal does not respond in a timely manner, and the standby node assumes that the active node has failed. There are several factors involved in detecting loss-of-heartbeat. NNMi avoids false failover notification as long as the network keeps up with the application failover data transfer needs.

In HP's verification of multi-subnet NNMi application failover, the active and standby servers resided in the United States, one in Colorado and another in Houston. This provided acceptable bandwidth and latency, with no false failovers.

Application Failover and the NNMi Embedded Database

Application failover works with both the embedded and the Oracle database for NNMi 9.20. However, with Oracle, the database resides on a server that is separate from any NNMi management server. When you configure NNMi to work with an Oracle database, there is no database replication. This results in reduced network demands for application failover using an Oracle database. When using application failover with Oracle, the network uses less than 1% of the network demands as compared to using application failover with the embedded database. The information contained in this section explains NNMi traffic information related to application failover using the embedded database.

After you configure NNMi using the embedded database for application failover, NNMi does the following:

- 1 The active node performs a database backup, storing the data in a single ZIP file.
- 2 NNMi sends this ZIP file across the network to the standby node.
- 3 The standby node expands the ZIP file, and configures the embedded database to import transaction logs on the first startup.
- 4 The embedded database on the active node generates transaction logs, depending on database activity.
- 5 Application failover sends the transaction logs across the network to the standby node, where they accumulate on the disk.
- 6 When the standby node becomes active, NNMi starts, and the database imports all transaction logs across the network. The amount of time this takes depends on the number of files and complexity of the information stored within those files (some files take longer to import than other files of comparable size).
- 7 After the standby node imports all of the transaction logs, the database becomes available, and the standby node starts the remaining NNMi processes.
- 8 The original standby node is now active, and the procedure starts over at [step 1](#).

Network Traffic in and Application Failover Environment

NNMi transfers many items across the network from the active node to the standby node in an application failover environment:

- Database Activity: the database backup, as a single ZIP file.
- Transaction logs.
- A periodic *heartbeat* so that each application failover node verifies that the other node is still running.
- File comparison lists so that the standby node can verify that its files are in sync with those on the active node.
- Miscellaneous events, such as changes in parameters (enable/disable failover and others) and nodes joining or node leaving the cluster,

The first two items generate 99% of the network traffic used by application failover. This section explores these two items in more detail.

Database Activity: NNMi generates transaction logs for all database activity. Database activity includes everything in NNMi. This activity includes, but is not limited to, the following database activities:

- Discovering new nodes.
- Discovering attributes about nodes, interfaces, VLANs, and other managed objects.
- State polling and status changes.
- Incidents, events, and root cause analysis.
- Operator actions in the NNMi console.

Database activity is outside of your control. For example, an outage on the network results in NNMi generating many incidents and events. These incidents and events trigger state polling of devices on the network, resulting in updates to device status in NNMi. When the outage is restored, additional *node up* incidents result in further status changes. All of this activity updates entries in the database.

Although the embedded database itself grows with database activity, it reaches a stable size for your environment, with only moderate growth over time.

Database Transaction Logs: The embedded database works by creating an empty 16 MB file, then writing database transaction information to that file. NNMi closes this file, then makes it available to application failover after 15 minutes, or after writing 16 MB of data to the file, whichever comes first. That means that a completely idle database will generate one transaction log file every 15 minutes, and this file will be essentially *empty*. Application failover compresses all transaction logs, so an empty 16 MB file compresses down to under 1MB. A *full* 16MB file compresses to about 8 MB. Keep in mind that during periods of higher database activity, application failover generates more transaction logs in a shorter period of time, since each file gets full faster.

An Application Failover Traffic Test

The following test resulted in an average of about 2 transaction log files per minute, with an average file size of 7 MB per file. This is due to the database activity associated with discovery of the additional 5000 nodes added with each failover event. The database in this test case eventually stabilized at about 1.1GB (as measured by the size of the backup ZIP file), with 31,000 nodes and 960,000 interfaces.

Testing Method: During the first 4 hours, test personnel seeded NNMi with 5,000 nodes and waited until discovery stabilized. After 4 hours, test personnel induced failover (the standby node became active, and the previous-active node became standby). Immediately after failover, test personnel added approximately 5,000 more nodes, waited another 4 hours to let the NNMi discovery process stabilize, then induced another failover (failed back to the previous active node). Test personnel repeated this cycle several times with some variation in the time between failover (4 hours, then 6 hours, then 2 hours). After each failover event, test personnel measure the following:

- The size of the database backup ZIP file (created when the node first became active).
- The transaction logs: the total number of files and disk space utilization.
- The number of nodes and interfaces in the NNMi database immediately before inducing failover.
- Time to complete failover. This included the time from the initial `ovstop` command on the active node until the standby node became fully active with NNMi running.

Table 27 summarizes the results:

Table 27 Application Failover Test Results

Hours	DB.zip Size (MB)	No. of Tx Logs	Tx Logs (GB)	Nodes	Interfaces	Failover Time (Minutes)
4	6.5	50	.3	5,000	15,000	5
8	34	500	2.5	12,000	222,000	10
12	243	500	2.5	17,000	370,000	25
16	400	500	3.5	21,500	477,000	23
20	498	500	3.5	25,500	588,000	32
26	618	1100	7.5	30,600	776,000	30
28	840	400	2.2	30,600	791,000	31
30	887	500	2.5	30,700	800,000	16


Observations: When NNMi transferred files from the active node to the standby node, the transfer averaged about 5 GB every 4 hours, which is a continuous throughput of approximately 350KB/s (kilobytes per second) or 2.8 Mb/S (megabits per second).

- ▶ This data does not include any other application failover traffic, such as the heartbeat, file consistency checks, or other application failover communication. This data also excludes the overhead of network I/O, such as packet headers. This data only included the actual network payload of each file's contents moving across the network.
- ▶ The traffic generated by NNMi application failover environment is very bursty. Application failover identifies new transaction logs on the active node every five minutes and sends these logs to the standby node. Depending on network speed, the standby node should receive all of the new files in a short time, resulting in a relatively idle network for the remainder of that 5-minute interval.

Every time the active and standby nodes switch roles (the standby node becomes active and the active node becomes standby), the new active node will generate a complete database backup and send this across the network to the new standby node. This database backup also occurs periodically, backing up every 24 hours by default. Every time NNMi generates a new backup, it sends this backup to the standby node. Having this new backup available on the standby node reduces the failover time, as all of the transaction logs NNMi generated in that 24 hour interval are already in the database, and do not need to be imported at failover time.

The information provided in the above section will help you understand how the network might perform after a failover when using NNMi with application failover using the embedded database.

Configuring NNMi in a High Availability Cluster



High availability (HA) refers to a hardware and software configuration that provides for uninterrupted service should some aspect of the running configuration fail. An HA cluster defines a grouping of hardware and software that works together to ensure continuity in functionality and data when failover occurs.

NNMi provides support for configuring NNMi to run in an HA cluster under one of several separately purchased HA products. Most of the NNM Smart Plug-ins (iSPIs), but not the NNM iSPI NET diagnostics server, can also run under HA.

This chapter provides a template for configuring NNMi to run in an HA environment. This chapter does not provide end-to-end instructions for configuring your HA product. The HA configuration commands that NNMi provides are wrappers around the commands for the supported HA products. If you prefer, you can substitute the HA product-specific commands where these instructions specify NNMi-provided commands.

If you plan to install any NNM iSPIs on the NNMi management server, also see the documentation for those NNM iSPIs.


This chapter contains the following topics:

- [HA Concepts](#) on page 314
- [Verifying the Prerequisites to Configuring NNMi for HA](#) on page 319
- [Configuring High Availability](#) on page 321
- [Shared NNMi Data](#) on page 331
- [Licensing NNMi in an HA Cluster](#) on page 334
- [Maintaining the HA Configuration](#) on page 336
- [Unconfiguring NNMi from an HA Cluster](#) on page 340
- [Patching NNMi under HA](#) on page 344
- [Upgrading NNMi under HA from NNMi 9.0x/9.1x to NNMi 9.20](#) on page 345
- [Troubleshooting the HA Configuration](#) on page 350
- [HA Configuration Reference](#) on page 361

HA Concepts

Cluster architecture provides a single, globally coherent process and resource management view for the multiple nodes of a cluster. [Figure 26](#) shows an example of a cluster architecture.


Figure 26 Architecture of a High Availability Cluster



Each node in a cluster connects to one or more public networks and also connects to a private interconnect, representing a communication channel for transmitting data between cluster nodes.

In modern cluster environments such as HP Serviceguard, Veritas Cluster Server, Microsoft Failover Clustering, or Microsoft Cluster Services, applications are represented as compounds of resources, which are simple operations that enable applications to run in a cluster environment. The resources construct an **HA resource group**, which represents an application running in a cluster environment. [Figure 27](#) shows an example HA resource group.

Figure 27 Typical HA Resource Group Layout



This document uses the term *HA resource group* to designate a set of resources in any cluster environment. Each HA product uses a different name for the HA resource group. Table 28 lists the term for each supported HA product that equates to *HA resource group* for this document. (For the specific supported versions of each HA product, see the *NNMi System and Device Support Matrix*.)

Table 28 Terminology for HA Resource Group in the Supported HA Products

HA Product	Abbreviation	Equivalent Term for HA Resource Group
Microsoft Failover Clustering	MSFC	Resource Group
HP Serviceguard	SG	Package
Veritas Cluster Server	VCS	Service Group
Red Hat Cluster Suite	RHCS	Service

HA Terms

Table 29 lists and defines some common HA terms.

Table 29 Common HA Terms

Term	Description
HA resource group	An application running in a cluster environment (under an HA product). An HA resource group can simultaneously be a cluster object that represents an application in a cluster.
Volume group	One or more disk drives that are configured to form a single large storage area.
Logical volume	An arbitrary-size space in a volume group that can be used as a separate file system or as a device swap space.

Table 29 Common HA Terms (cont'd)

Term	Description
Primary cluster node	The first system on which the software product is installed, <i>and</i> the first system on which HA is configured. The shared disk is mounted on the primary cluster node for initial set up. The primary cluster node generally becomes the first active cluster node, but you do not need to maintain the primary designation after HA configuration is complete. The next time you update the HA configuration, another node might become the primary cluster node.
Secondary cluster node	Any system that is added to the HA configuration after the primary cluster node has been fully configured for HA.
Active cluster node	The system that is currently running the HA resource group.
Passive cluster node	Any system that is configured for HA but is not currently running the HA resource group. If the active cluster node fails, the HA resource group fails over to one of the available passive cluster nodes, which then becomes the active cluster node for that HA resource group.

NNMi HA Cluster Scenarios

- ▶ NNMi supports clusters where the application can run on more than two cluster nodes. See the *nms-ha* manpage and the *nnmdatareplicator.ovpl* reference page, or the UNIX manpage for more information.

For NNMi HA configuration, NNMi is installed on each system that will become part of an HA resource group. The NNMi database is located on a separate disk that is accessed by the NNMi programs running on each system. (Only one system, the active cluster node, accesses the shared disk at any given time.)

This approach is valid for the embedded and third-party database solutions.

- ▶ Run the NNMi database backup and restore scripts on the active cluster node only.


NNMi-only scenario

Figure 28 shows a graphical representation of the NNMi HA cluster scenario. In this figure the NNMi HA resource group is synonymous with the NNMi HA cluster.

Node A and node B are each a fully installed NNMi management server that contains the NNMi program and any NNM iSPIs that run on that system. The active cluster node accesses the shared disk for runtime data. Other products connect to NNMi by the virtual IP address of the HA resource group.

If the cluster contains more than two NNMi nodes, additional nodes are configured similarly to node B in Figure 28.

Figure 28 Basic Scenario for NNMi HA Cluster

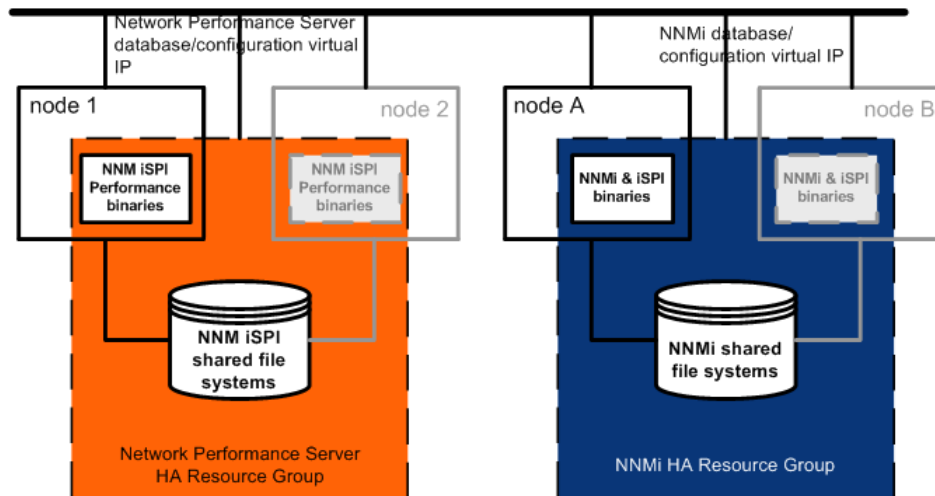


For information about how to implement this scenario, see [Configure NNMi for HA](#) on page 321 and [Configure NNM iSPIs for HA](#) on page 328.

NNMi and NNM Performance iSPIs on a standalone server scenario

If you are running any of the NNM Performance iSPIs on a standalone server, you can configure these NNM iSPIs to run as a separate HA resource group within the NNMi HA cluster, as shown in [Figure 29](#). The NNMi HA resource group is the same as that described for the NNMi-only scenario.

Figure 29 HA for NNMi and NNM Performance iSPIs on a Standalone Server



For information about how to implement this scenario, see [Configure NNMi for HA](#) on page 321 and [Configure NNM iSPIs for HA](#) on page 328.

Other options for the NNM Performance iSPIs on a standalone server are as follows:


- Run the NNM Performance iSPIs on a single system with no HA. Use this approach while evaluating the NNM iSPIs and for environments where it is not critical for performance data to be always available.
- Configure the NNM Performance iSPIs to run under a different HA cluster than that for NNMi. In this case, you must manage the NNM Performance iSPIs' dependency on NNMi manually.

NNMi with an Oracle database scenario

If your NNMi implementation uses Oracle for the main NNMi database, the Oracle database should be on a separate server, as shown in [Figure 30](#), for performance reasons. Therefore, you must configure two HA resource groups within the NNMi HA cluster:

- The NNMi HA resource group includes the NNMi nodes and a shared disk for NNMi data that is not stored in the Oracle database.
- The Oracle HA resource group contains the Oracle database server and the database disk.

Figure 30 HA for NNMi with an Oracle Database

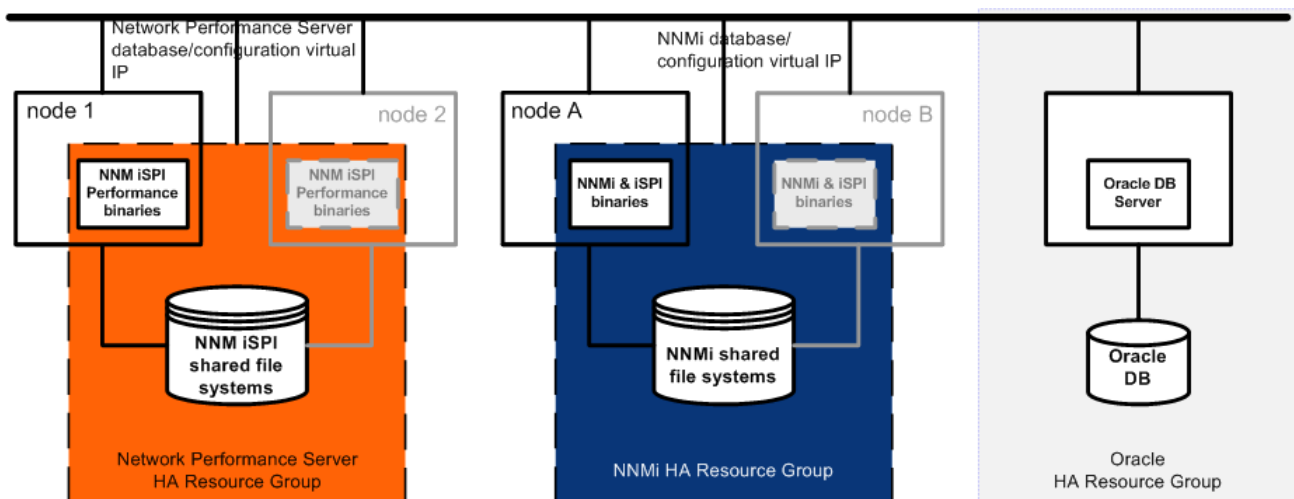


For information about how to implement this scenario, see [Configure NNMi for HA in an Oracle Environment](#) on page 329 and [Configure NNM iSPIs for HA](#) on page 328.

NNMi with an Oracle database and NNM Performance iSPIs on a standalone server scenario

If your NNMi implementation uses Oracle for the main NNMi database and you are running any of the NNM Performance iSPIs on a standalone server, you can configure three HA resource groups within the NNMi HA cluster, as shown in [Figure 31](#).

Figure 31 HA for NNMi with an Oracle Database and NNM Performance iSPIs on a Standalone Server



For information about how to implement this scenario, see [Configure NNMi for HA in an Oracle Environment](#) on page 329 and [Configure NNM iSPIs for HA](#) on page 328.

Manpages

With regard to HA configuration, the NNMi manpages contain the following topics:

- nnm-ha
- nnmhaconfigure.ovpl
- nnmhaunconfigure.ovpl
- nnmhadisk.ovpl
- nnmhaclusterinfo.ovpl
- nnmhastartrg.ovpl
- nnmhastoprg.ovpl

On the Windows operating system, these manpages are available as text files.

Verifying the Prerequisites to Configuring NNMi for HA

Successful configuration of NNMi for HA depends on a number of factors:

- Appropriate hardware
- Understanding of the HA product
- A methodical approach to configuration

Before you begin to configure NNMi for HA, complete the following preparation:

- 1 Verify that NNMi supports your HA product by checking the information in the *NNMi System and Device Support Matrix*.
- 2 Read the documentation for your HA product to familiarize yourself with the capabilities of that product and to make design decisions.



HA product documentation changes frequently. Be sure you have the most recent versions available.

- 3 Verify that each system to be included as a node in an NNMi HA cluster meets the following requirements:

- Meets all requirements described in the documentation for the HA product.
- Includes at least two network interface cards (NIC cards).



Review the HA product, operating system, and NIC card documentation to verify that these products can all work together.

- Supports the use of a virtual IP address for the HA resource group. This IP address is the IP address used for the NNMi license.



MSFC requires multiple virtual IP addresses, one for the HA cluster and one for each HA resource group. In this case, the virtual IP address of the NNMi HA resource group is the IP address used for the NNMi license.

- Supports the use of a shared disk or disk array
- ▶ Review the HA product, operating system, and disk manufacturer documentation to verify that these products, including the related SCSI cards, can all work together.
- Meets all requirements for NNMi as described in the *NNMi System and Device Support Matrix*.
- 4 If you plan to run any NNM iSPIs in the NNMi HA cluster, read the appropriate NNM iSPI documentation for additional HA configuration prerequisites.
 - 5 Allocate the following virtual IP addresses and hostnames:
 - One virtual IP address for the HA cluster (MSFC only)
 - One virtual IP address for each HA resource group to be configured
 - 6 From any system, use the `nslookup` command to validate correct DNS response for all of the IP addresses and hostnames you allocated in [step 5](#).
 - 7 Verify that operating system of each system is at the correct version and patch level for the HA product and NNMi.
 - 8 If necessary, install the HA product.

▶ In a Solaris Zones environment, install the HA product in the global zone.
 - 9 Prepare the shared disk as described in [Prepare the Shared Disk Manually](#) on page 332.
 - 10 Use the commands for your HA product to configure (if necessary) and test an HA cluster.

The HA cluster provides such functionality as checking the application heartbeat and initiating failover. The HA cluster configuration must, at a minimum, include the following items:

- (UNIX only) `ssh`, `remsh`, or both
 - (Windows only) Virtual IP address for the HA cluster that is DNS-resolvable
 - Virtual hostname for the HA cluster that is DNS-resolvable
 - A resource group that is unique and specific to NNMi.
- ▶ NNMi expects that the NNMi HA resource group includes all required resources. If this is not the case, use the HA product functionality to manage dependencies between the NNMi HA resource group and the other HA resource groups. For example, if Oracle is running in a separate HA resource group, configure the HA product to ensure that the Oracle HA resource group is fully started before the HA product starts the NNMi HA resource group.
- *MSFC*: Use the create cluster wizard of Failover Cluster Management for Windows Server 2008.
 - *ServiceGuard*:
 - Add `.rhosts` entries or `.ssh` entries for nodes.
 - Configure the HA product (`cmgetconf`, `cmcheckconf`, `cmapplyconf`). See the most recent documentation for your HA product on setting up a cluster.
 - *VCS*: Not necessary. Product installation created an HA cluster.

- *RHCS*: Add services (cman, rgmanager) as described in the RHCS documentation.

For information about testing the resources that you will place into the NNMi HA resource group, see [HA Resource Testing](#) on page 351.

Configuring High Availability

This section describes the procedures for configuring a new HA configuration for NNMi. It contains the following topics:

- [Configure NNMi Certificates for HA](#) on page 321
- [Configure NNMi for HA](#) on page 321
- [Configure NNM iSPIs for HA](#) on page 328
- [Configure NNMi for HA in an Oracle Environment](#) on page 329

➤ If you are running NNMi in a Solaris Zones environment, you do not need to follow the configuration process described in this chapter. See [Running NNMi under HA in a Solaris Zones Environment](#) on page 278.

➤ RHCS configuration requires a complete restart of the HA cluster daemons, including all applications, on each node in the HA cluster. Plan your configuration effort accordingly.

Configure NNMi Certificates for HA

The NNMi installation process configures a self-signed certificate for secure communications between the NNMi console and the NNMi database. The process for configuring NNMi for HA correctly shares the self-signed certificate among the primary and secondary cluster nodes. You do not need to take any extra steps to use the default certificate with NNMi running under HA.

If you want to use a different self-signed certificate or a Certificate Authority (CA)-signed certificate for NNMi communications, you must do some additional work. After obtaining the new certificate, complete the steps shown in [Configuring High Availability for a New Certificate](#) on page 134. You can complete this procedure before or after configuring NNMi for HA.

Configure NNMi for HA

The two distinct phases of configuring NNMi for HA are as follows:

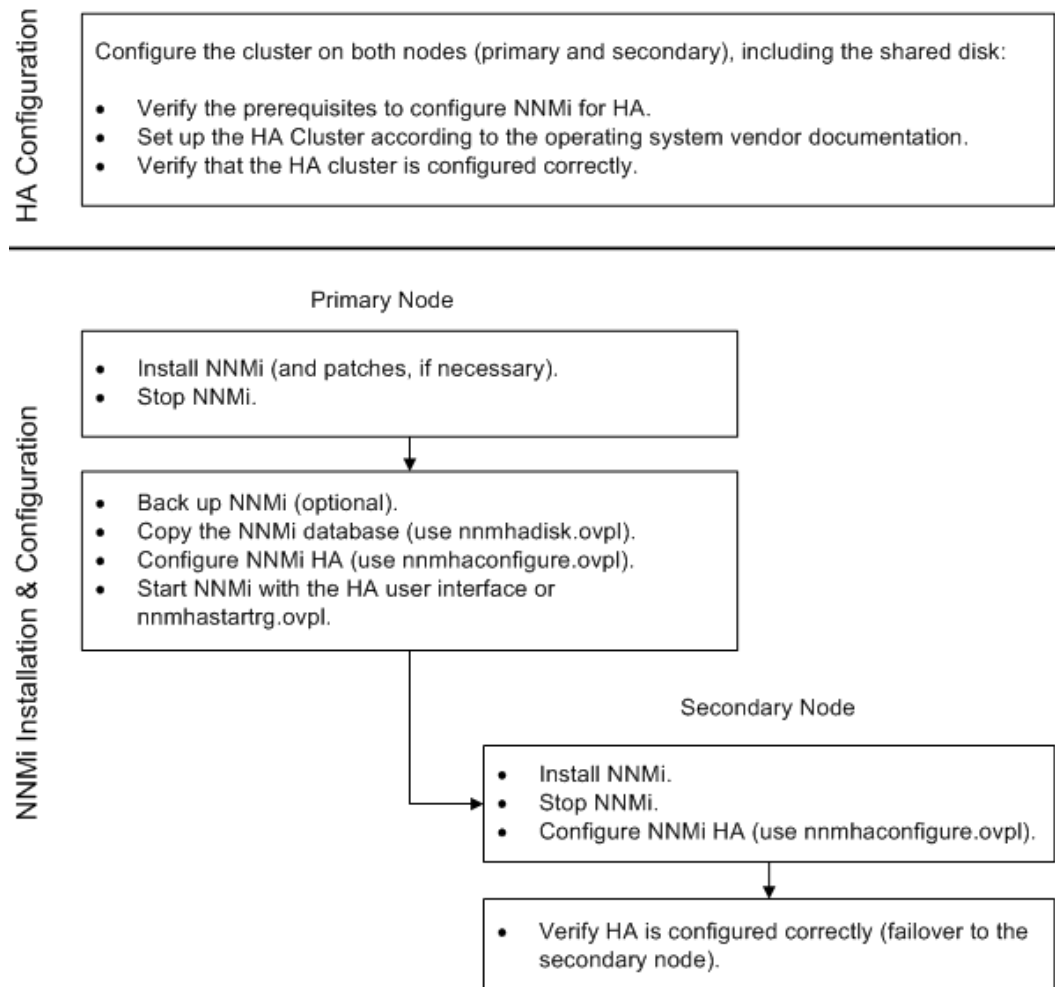
- 1 Copy the NNMi data files to the shared disk.
 - Do this task on the primary node, as described in [step 1](#) through [step 9](#) of [Configuring NNMi on the Primary Cluster Node](#) on page 325.
- 2 Configure NNMi to run under HA.
 - Do this task on the primary node, as described in [step 10](#) through [step 15](#) of [Configuring NNMi on the Primary Cluster Node](#) on page 325.
 - Also do this task on the secondary node, as described in [Configuring NNMi on the Secondary Cluster Nodes](#) on page 327.

Designate one HA cluster node as the primary NNMi management server. This is the node you expect to be active most of the time. Configure the primary node, and then configure all other nodes in the HA cluster as secondary nodes.

- ⚠ You *cannot* configure NNMi for HA simultaneously on multiple cluster nodes. After the HA configuration process is completed on one cluster node, proceed with the HA configuration on the next node, and so forth until NNMi is configured for HA on all nodes in the cluster environment.
- ➡ During failover, the NNMi console is unresponsive. After failover completes, NNMi users must log on to continue their NNMi console sessions.

Figure 32 provides an illustration of the NNMi HA configuration process.

Figure 32 NNMi HA Configuration Workflow





If you encounter errors during HA configuration, do the following:

- 1 Unconfigure NNMi from the HA environment by running the `nmhaunconfigure.ovpl` command.
- 2 Correct the condition indicated by the error message(s).
- 3 Reconfigure NNMi into the HA environment by running the `nmhaconfigure.ovpl` command.

For more information, see the `nmhaunconfigure.ovpl` and `nmhaconfigure.ovpl` reference pages, or the UNIX manpages.

NNMi HA Configuration Information

The HA configuration script collects information about the NNMi HA resource group. Please prepare the information listed in [Table 30](#) before you configure NNMi HA. This information is needed to execute the HA script (`nmhaconfigure.ovpl`) interactively, depending on your operating system or HA software.

Table 30 NNMi HA Primary Node Configuration Information

HA Configuration Item	Description
HA resource group	<p>The name of the resource group for the HA cluster that contains NNMi. This name must be unique, specific to NNMi, and not currently in use. See your HA system provider's reference material for information on valid names.</p> <p>Upon input of an HA resource group name, NNMi generates the following resources for UNIX and Windows systems:</p> <p style="padding-left: 40px;"><code><resource group name>-IP</code></p> <p style="padding-left: 40px;"><code><resource group name>-Mount</code></p> <p style="padding-left: 40px;"><code><resource group name>-App</code></p> <p>In addition, for Windows systems, the following resource is generated upon input of a virtual hostname:</p> <p style="padding-left: 40px;"><code><virtual hostname></code></p>
Virtual host short name	<p>The short name for the virtual host. This hostname must map to the virtual IP address for the HA resource group. The <code>nslookup</code> command must be able to resolve the virtual host short name and the virtual IP address.</p> <p>NOTE: If NNMi is unable to resolve the virtual host short name or the virtual host IP address, the HA configuration script could leave the system in an unstable state. Therefore, HP recommends that you implement a secondary naming strategy (such as entering the information in the <code>%SystemRoot%\system32\drivers\etc\hosts</code> file on the Windows operating system or <code>/etc/hosts</code> file on UNIX operating systems) in case DNS is not available during NNMi HA configuration.</p>
Virtual host netmask	<p>The subnet mask that is used with the virtual host IP address, which must be an IPv4 address.</p>

Table 30 NNMi HA Primary Node Configuration Information (cont'd)

HA Configuration Item	Description
Virtual host network interface	<p>The network interface on which the virtual host IP address is running. For example:</p> <ul style="list-style-type: none"> • <i>Windows</i>: Local Area Connection • <i>HP-UX</i>: lan0 • <i>Linux</i>: eth0 • <i>Solaris</i>: bge0
Shared file system type	<p>The type of shared disk configuration being used for the HA resource group. Possible values are:</p> <ul style="list-style-type: none"> • <code>disk</code>—The shared disk is a physically attached disk that uses a standard file system type. The HA configuration script can configure the shared disk. For more information, see the File system type entry in this table. • <code>none</code>—The shared disk uses a configuration other than that described for the <code>disk</code> option, such as NFS. After running the HA configuration script, configure the shared disk as described in Prepare the Shared Disk Manually on page 332.
File system type	<p>(UNIX only) The file system type of the shared disk (if the shared file system type is <code>disk</code>). The HA configuration scripts pass this value to the HA product so that it can determine how to validate the disk.</p> <p>HP has tested the following shared disk formats:</p> <ul style="list-style-type: none"> • <i>Windows</i>: Basic (see A Note about Shared Disk Configuration on Windows Server on page 334); SAN • <i>HP-UX</i>: vxfs • <i>Linux</i>: ext2, ext3, and vxfs for VCS and RHCS • <i>Solaris</i>: vxfs <p>NOTE: HA products support other file system types. If you use a shared disk format that HP has not tested, prepare the disk before configuring NNMi to run under HA, and then specify <code>none</code> for the shared file system type while running the NNMi HA configuration script.</p>
Disk information (disk group, volume group, and/or logical volume name, depending on the operating system used)	<p>The name associated with the disk information for the NNMi shared file system.</p> <p>NOTE: When you create/attach a disk on UNIX platforms, for example, with vxfs or lvm, you create different items, such as: disk group, volume group, logical volume. The names for these items are assigned by the system administrator at the time of creation. NNMi does not enforce any naming conventions. Contact your system administrator for your company's naming information.</p>
Mount point	<p>The directory location for mounting the NNMi shared disk. This mount point must be consistent between systems. (That is, each node must use the same name for the mount point.) For example:</p> <ul style="list-style-type: none"> • <i>Windows</i>: S:\ <ul style="list-style-type: none"> NOTE: Specify the drive completely. S and S: are unacceptable formats and do not provide access to the shared disk. • <i>UNIX</i>: /nnmmount

Configuring NNMi on the Primary Cluster Node

Complete the following procedure on the primary cluster node.

- ▶ If you are using Oracle for the main NNMi database, see [Configure NNMi for HA in an Oracle Environment](#) on page 329 first.
- ▶ If you are running NNMi in a Solaris Zones environment, you do not need to follow the configuration process described in this chapter. See [Running NNMi under HA in a Solaris Zones Environment](#) on page 278.

- 1 If you have not already done so, complete the procedure for [Verifying the Prerequisites to Configuring NNMi for HA](#) on page 319.
- 2 If you have not already done so, install NNMi (including the latest consolidated patch, if any), and then verify that NNMi is working correctly.
- 3 If you expect to run any NNM iSPiS on this NNMi management server, see [Configure NNM iSPiS for HA](#) on page 328 before continuing with this procedure.
- 4 Use the `nnmbackup.ovpl` command, or another database command, to back up all NNMi data. For example:

```
nnmbackup.ovpl -type offline -scope all -target nnmi_backups
```

For more information about this command, see [NNMi Backup and Restore Tools](#) on page 383.

- 5 Define the disk device group (and logical volume), consisting of at least one shared disk for the NNMi HA resource group. For example:
 - *MSFC*: Use Disk Management to configure the disk mount point and format the disk.
 - *Serviceguard*:
Use LVM commands such as `pvcreate`, `vgcreate`, and `lvcreate` to initialize the disk, create the volume group, and create the logical volume.
 - *VCS*:
Use VSF commands such as `vxdiskadm`, `vxassist`, and `mkfs` to add and initialize the disk, allocate disks by space, and create the logical volume.
 - *RHCS*:
Use LVM commands such as `pvcreate`, `vgcreate`, and `lvcreate` to initialize the disk, create the volume group, and create the logical volume.

- ▶ NNMi requires RHCS clusters be configured such that the `clusternode` names specified in the `/etc/cluster/cluster.conf` file must be fully qualified for NNMi to correctly start and stop.

For UNIX operating systems, a reference web site is:
<http://www.unixguide.net/unixguide.shtml>

- 6 Create the directory mount point (for example, `S:\` or `/nnmmount`), and then mount the shared disk:



After configuration, the HA product manages disk mounting. Do *not* update the file system table with this mount point.

- *Windows*: Use Windows Explorer and Disk Management.
- *UNIX*:
 - Use the `mkdir` and `mount` commands.
 - Verify that the shared disk directory mount point has been created with `root` as the user, `sys` as the group, and the permissions set to 555. For example:

```
ls -l /nnmmount
```

- 7 Stop NNMi:

```
ovstop -c
```



If NNMi is already installed on a node that you will include in this HA resource group, also run `ovstop -c` on that node at this time.

- 8 Copy the NNMi database to the shared disk:

- *Windows*:

```
%NnmInstallDir%\misc\nnm\ha\nnmhadisk.ovpl NNM \
-to <HA_mount_point>
```

- *UNIX*:

```
$NnmInstallDir/misc/nnm/ha/nnmhadisk.ovpl NNM \
-to <HA_mount_point>
```



To prevent database corruption, run this command (with the `-to` option) only one time. For information about alternatives, see [Re-Enable NNMi for HA after All Cluster Nodes are Unconfigured](#) on page 356.

- 9 (UNIX only) Unmount the shared disk and deactivate the disk group:

```
umount <HA_mount_point>
vgchange -a n <disk_group>
```

- 10 Verify that NNMi is not running:

```
ovstop -c
```

- 11 (RHCS only) Copy the NNMi custom script into place, and then restart the HA cluster daemons.

- a Copy the `/opt/OV/misc/nnm/ha/NNMscript.sh` file to the following location:

```
/usr/share/cluster/NNMscript.sh
```

- b Stop and then restart the `/sbin/ccsd` process.

12 Configure the NNMi HA resource group:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaconfigure.ovpl NNM
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaconfigure.ovpl NNM
```

[Table 30](#) on page 323 describes the information that this command requests.

13 (UNIX only) By default NNMi starts in the locale of the user who ran the `nmhaconfigure.ovpl` command. To change the NNMi locale, run the following command:

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \  
-config NNM -set HA_LOCALE <locale>
```

14 In [step 12](#), what value did you specify for the shared file system type (as described for [Shared file system type](#) and [File system type](#) in [Table 30](#) on page 323)?

- For type `disk`, the `nmhaconfigure.ovpl` command configured the shared disk. Continue with [step 15](#).
- For type `none`, prepare the shared disk as described in [Prepare the Shared Disk Manually](#) on page 332, and then continue with [step 15](#).

15 Start the NNMi HA resource group:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhastartrg.ovpl NNM \  
<resource_group>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhastartrg.ovpl NNM \  
<resource_group>
```

If NNMi does not start correctly, see [Troubleshooting the HA Configuration](#) on page 350.



Now that NNMi is running under HA, *do not* use the `ovstart` and `ovstop` commands for normal operation. Use these commands only when instructed to do so for HA maintenance purposes.

Configuring NNMi on the Secondary Cluster Nodes

Complete the following procedure on one secondary cluster node at a time.

- 1 If you have not already done so, complete the procedure for [Configuring NNMi on the Primary Cluster Node](#) on page 325.
- 2 If you have not already done so, complete the procedure for [Verifying the Prerequisites to Configuring NNMi for HA](#) on page 319.
- 3 If you have not already done so, install NNMi (including the latest consolidated patch, if any), and then verify that NNMi is working correctly.
- 4 Install the NNM iSPIs that you installed in [step 3](#) of [Configuring NNMi on the Primary Cluster Node](#) on page 325.
- 5 Stop NNMi:

```
ovstop -c
```

- 6 Create a mount point for the shared disk (for example, `S:\` or `/nnmmount`).



This mount point must use the same name as the mount point you created in [step 6](#) of the procedure [Configuring NNMi on the Primary Cluster Node](#).

- 7 (RHCS only) Copy the NNMi custom script into place, and then restart the HA cluster daemons.

- a Copy the `/opt/OV/misc/nnm/ha/NNMscript.sh` file to the following location:

```
/usr/share/cluster/NNMscript.sh
```

- b Stop and then restart the `/sbin/ccsd` process.

- 8 Configure the NNMi HA resource group:

- *Windows:* `%NnmInstallDir%\misc\nnm\ha\nnmhaconfigure.ovpl NNM`
- *UNIX:* `$NnmInstallDir/misc/nnm/ha/nnmhaconfigure.ovpl NNM`

Supply the HA resource group name when the command requests this information.

- 9 Verify that the configuration was successful:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \  
-group <resource_group> -nodes
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \  
-group <resource_group> -nodes
```

The command output lists all configured nodes for the specified HA resource group.

- 10 Optionally, test the configuration by taking the NNMi HA resource group on the primary node offline and then bringing the NNMi HA resource group on the secondary node online.

Configure NNM iSPIs for HA

If you expect to run any NNM iSPIs on the NNMi management server, read this section before configuring NNMi to run under HA.

NNM iSPI Performance for Metrics, NNM iSPI Performance for QA, and NNM iSPI Performance for Traffic

The NNM Performance iSPIs (NNM iSPI Performance for Metrics, NNM iSPI Performance for QA, and NNM iSPI Performance for Traffic) can be installed on the NNMi management server or on a standalone server, but not on a combination of these two options.

- If the NNM Performance iSPIs will be located on the NNMi management server, install the products before configuring NNMi to run under HA.
- If the NNM Performance iSPIs will be located on a standalone server, configure NNMi to run under HA before installing the products. During the NNM iSPI installation process, supply the NNMi HA resource group virtual hostname as the NNMi management server name.

NNM iSPI for MPLS, NNM iSPI for IP Multicast, and NNM iSPI for IP Telephony

The NNM iSPI for MPLS, NNM iSPI for IP Multicast, and NNM iSPI for IP Telephony can be installed on the NNMi management server only. Install these products before configuring NNMi to run under HA.

For information about configuring the NNM iSPIs to run under HA, see the documentation for the appropriate NNM iSPI.

NNM iSPI Network Engineering Toolset Software and NNMi Running under HA

The NNM iSPI Network Engineering Toolset Software SNMP trap analytics and Microsoft Visio export functionality are automatically installed with NNMi. No extra work is needed to run these tools under HA.

The NNM iSPI NET diagnostics server cannot be included in the NNMi HA resource group. Do not install this component on the NNMi management server. To run the NNM iSPI NET diagnostics server on a system that is outside the NNMi HA resource group, follow these steps:

- 1 Completely configure the NNMi HA resource group.
- 2 Install the NNM iSPI NET diagnostics server on a system that is outside the NNMi HA resource group. During the NNM iSPI NET diagnostics server installation process, supply the NNMi HA resource group virtual hostname as the NNM Server Hostname.

For more information, see the *NNM iSPI Network Engineering Toolset Software Planning and Installation Guide*.

If the NNM iSPI NET diagnostics server is already installed on an NNMi management server that will run under HA, uninstall the NNM iSPI NET diagnostics server before configuring NNMi to run under HA.



Uninstalling the NNM iSPI NET diagnostics server removes all existing reports.



It might be possible to save existing reports, as described here, but the following procedure is untested:

- 1 Use MySQL Workbench to perform a backup of the existing `nnminet` database. MySQL Workbench is available in the downloads area at dev.mysql.com.
- 2 Uninstall the NNM iSPI NET diagnostics server.
- 3 Configure NNMi to run under HA.
- 4 Install the NNM iSPI NET diagnostics server on a separate system.
- 5 Before running any flows, use MySQL Workbench to recover the `nnminet` database onto the new installation.

Configure NNMi for HA in an Oracle Environment

This sections presents a high-level overview of the process for configuring NNMi with an Oracle database to run under HA. The number of possible Oracle configurations is large, and the configuration process can vary according to the Oracle release. For the most accurate information about configuring Oracle to run under HA and creating an

NNMi dependency on the Oracle HA resource group, see the HA product documentation. You can also go to the Oracle web site (www.oracle.com) for information about the appropriate Oracle configuration for your HA product.

NNMi Dependency on Oracle

When Oracle and NNMi both run under HA, the NNMi HA resource group must include a shared disk for the NNMi data that is not stored in the Oracle database. Additionally, consider the following information:

- If the HA product supports dependencies, the recommended approach is to configure each product to run in a separate HA resource group. The Oracle HA resource group must be fully started before the NNMi HA resource group starts. If both HA resource groups are in the same HA cluster, you can modify the cluster configuration to set resource group ordering. If the HA resource groups are in different HA clusters, make sure that the NNMi HA resource group dependency on the Oracle HA resource group is met.
- If the HA product does not support dependencies, include the Oracle systems and the NNMi systems in the NNMi HA resource group.

Configuring NNMi for HA in an Oracle Environment

- 1 If you plan to run Oracle under HA, complete that configuration first.
- 2 Create an empty Oracle database instance for NNMi.
- 3 On the primary NNMi node, install NNMi (including the latest consolidated patch, if any). During installation, do the following:
 - a Select the **Oracle** database type, and then select **Primary Server Installation**.
 - b Specify the virtual IP address or hostname for the Oracle HA resource group (if applicable).
- 4 On the primary NNMi node, configure NNMi to run under HA as described in [Configuring NNMi on the Primary Cluster Node](#) on page 325.
- 5 Set up the NNMi dependency on the Oracle HA resource group.
For specific instructions, see the HA product documentation.
- 6 On the secondary NNMi node, install NNMi (including the latest consolidated patch, if any). During installation, do the following:
 - Select the **Oracle** database type, and then select **Secondary Server Installation**.
 - Specify the virtual IP address or hostname for the Oracle HA resource group (if applicable).
- 7 On the secondary NNMi node, configure NNMi to run under HA described in [Configuring NNMi on the Secondary Cluster Nodes](#) on page 327.
- 8 For each additional secondary NNMi node, repeat [step 6](#) and [step 7](#).

Shared NNMi Data

This implementation of NNMi running under HA requires the use of a separate disk for sharing files between all NNMi nodes in the HA cluster.



NNMi implementations that use Oracle as the primary database also require the use of a separate disk for shared data.

Data on the NNMi Shared Disk

This section lists the NNMi data files that are maintained on the shared disk when NNMi is running under HA.

The locations are mapped to the shared disk location as follows:

- *Windows:*
 - %NnmInstallDir% maps to %HA_MOUNT_POINT%\NNM\installDir
 - %NnmDataDir% maps to %HA_MOUNT_POINT%\NNM\dataDir
- *UNIX:*
 - \$NnmInstallDir maps to \$HA_MOUNT_POINT/NNM/installDir
 - \$NnmDataDir maps to \$HA_MOUNT_POINT/NNM/dataDir

The directories that are moved to the shared disk are as follows:

- *Windows:*
 - %NnmDataDir%\shared\nnm\databases\Postgres
The embedded database; not present when using an Oracle database.
 - %NnmDataDir%\log\nnm
The NNMi logging directory.
 - %NnmDataDir%\shared\nnm\databases\eventdb
The pmd events database.
 - %NnmDataDir%\nmsas\NNM\data
The transactional store used by ovjboss.
- *UNIX:*
 - \$NnmDataDir/shared/nnm/databases/Postgres
The embedded database; not present when using an Oracle database.
 - \$NnmDataDir/log/nnm
The NNMi logging directory.
 - \$NnmDataDir/shared/nnm/databases/eventdb
The pmd events database.
 - \$NnmDataDir/nmsas/NNM/data
The transactional store used by ovjboss.

The `nmmhadisk.ovpl` command copies these files to and from the shared disk. Run this command as the instructions in this chapter indicate. For a summary of the command syntax, see the *nmm-ha* manpage.

Replication of Configuration Files

The NNMi HA implementation uses file replication to maintain copies of the NNMi configuration files on all NNMi nodes in the HA cluster. By default, NNMi manages file replication, copying NNMi configuration files from the active node to a passive node during the failover process. The `nmdatareplicator.conf` file specifies the NNMi folders and files included in data replication.

Disabling Data Replication

You can disable data replication as follows:

- 1 Edit the following file:
 - *Windows*: %NnmDataDir%\shared\nnm\conf\ov.conf
 - *UNIX*: \$NnmDataDir/shared/nnm/conf/ov.conf
- 2 Include the following line:


```
DISABLE_REPLICATION=DoNotReplicate
```
- 3 Save your changes.
- 4 From a command prompt, run `ovstop`.
- 5 From a command prompt, run `ovstart`.



When you change files (for example, configuration files) on the Active node, these files are automatically replicated to the Standby node on failover.

Prepare the Shared Disk Manually

If the shared disk is of a format that HP has tested (as listed in [Table 30](#) on page 323), the HA configuration script prepares the shared disk, and you can ignore this section.

If the shared disk uses a non-tested configuration, such as disk formats supported by the HA product, you must prepare the disk manually. Enter the value `none` for the file system type during HA configuration, and then configure the shared disk and the NNMi HA resource group's use of the shared disk.



You can configure the disk before or after configuring the NNMi HA resource group.


To prepare the shared disk manually, follow these steps:

- 1 Configure the shared disk as described in [Configuring a SAN or a Physically Connected Disk](#) on page 333.
- 2 Configure the NNMi HA resource group to recognize the disk by completing both of the following procedures:
 - [Setting the HA Variables in the ov.conf File](#) on page 333
 - [Moving the Shared Disk into the NNMi HA Resource Group](#) on page 334

Configuring a SAN or a Physically Connected Disk

Connecting and formatting a disk that disk into a vxfs or ext3 file system. To configure a SAN or a physically-connected disk, follow these steps:


- 1 Verify that the shared disk is *not* configured to be mounted at system boot time. The resource group is responsible for mounting the shared disk.
- 2 Connect the device:
 - For a SAN disk, add the SAN device to the network. The logical volume on the SAN disk should be in exclusive mode, if that mode is available.
 - For a physically-connected disk, attach the disk using a Y cable.
- 3 Add operating system entries to all cluster nodes (disk group, logical volume, volume group, and disk):
 - For a SAN disk, the entries reference the SAN.
 - For a physically-connected disk, the entries reference the disk hardware.
- 4 Format the disk using a disk format listed in [Table 30](#) on page 323.
- 5 Ensure that the SAN mounts.

 For UNIX systems, a reference web site is: <http://www.unixguide.net/unixguide.shtml>
- 6 Unmount and deport the disk.
- 7 To test the configuration, add the disk to a resource group and initiate failover.

Setting the HA Variables in the ov.conf File

The NNMi HA resource group uses the following variables to access the shared disk:

- HA_POSTGRES_DIR=<HA_mount_point>/NNM/dataDir/shared/nnm/databases/Postgres
- HA_EVENTDB_DIR=<HA_mount_point>/NNM/dataDir/shared/nnm/eventdb
- HA_NNM_LOG_DIR=<HA_mount_point>/NNM/dataDir/log
- HA_JBOSS_DATA_DIR=<HA_mount_point>/NNM/dataDir/nmsas/NNM/data
- HA_MOUNT_POINT=<HA_mount_point>
- HA_CUSTOMPOLLER_DIR=<HA_mount_point>/NNM/dataDir/shared/nnm/databases/custompoller

 If you plan to run any NNM iSPIs in the NNMi HA resource group, also set the ov.conf variables for each of those NNM iSPIs. For more information, see the documentation for the appropriate NNM iSPI.

To set the product variables for accessing the shared disk in the ov.conf file, run the following command for each of the preceding variables:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \  
-config NNM -set <variable> <value>
```

- *UNIX*:

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \  
-config NNM -set <variable> <value>
```

Moving the Shared Disk into the NNMi HA Resource Group

Modify the disk configuration file according to the product documentation to move the shared disk into the NNMi HA resource group. For example:



You can also use this process to add other resources, such as a NIC card or a backup disk to the NNMi HA resource group.

- *MSFC*: Use Failover Management to add resources to the resource group.
- *ServiceGuard*:

```
/etc/cmcluster/<resource_group>/<resource_group>.cntl
```
- *VCS*: Add disk entries and links to the HA configuration file by using the `/opt/VRTSvcs/bin/hares` command. For example:
- *RHCS*:

```
/etc/cluster/cluster.conf
```

A Note about Shared Disk Configuration on Windows Server

According to Microsoft Knowledge Base article 237853, dynamic disks are not supported for clustering with Windows Server 2008. To ensure the correct disk configuration, review the information located on the following web sites:

- <http://support.microsoft.com/kb/237853>
- http://www.petri.co.il/difference_between_basic_and_dynamic_disks_in_windows_xp_2000_2003.htm

Licensing NNMi in an HA Cluster

NNMi requires two licenses to run NNMi in an HA cluster:

- one production license tied to the IP address of one of the physical cluster nodes
- one non-production license tied to the virtual IP address of the NNMi HA resource group

The NNMi license keys are managed on the shared disk. Therefore, each NNMi HA resource group requires only the non-production license keys for each separately licensed product.

When licensing NNMi in an HA cluster, you must update the `licenses.txt` file on the shared disk with the new information from the license file on the active node. Complete the following procedure to correctly license NNMi in an HA cluster.

To correctly license NNMi in an HA cluster, perform these steps on the active NNMi cluster node:

- 1 Obtain and install a permanent non-production license key for each of your ordered products as described in [Licensing NNMi](#) on page 121. When prompted for the IP address of the NNMi management server, provide the virtual IP address of the NNMi HA resource group.
- 2 Update the `licenses.txt` file on the shared disk with the new information from the `LicFile.txt` file on the active node. Do one of the following:
 - If the `licenses.txt` file exists in the `NNM` directory on the shared disk, append the new license keys in `LicFile.txt` on the active node to `licenses.txt` on the shared disk.
 - If the `licenses.txt` file does not exist on the shared disk, copy `LicFile.txt` from the active node to `licenses.txt` in the `NNM` directory on the shared disk.

On the active node, the `LicFile.txt` file is in the following location:

- *Windows:* `<drive>:\Program Files (x86)\HP\HP BTO Software\data\shared\nnm\conf\licensing\LicFile.txt`
- *UNIX:* `/var/opt/OV/shared/nnm/conf/licensing/LicFile.txt`

On the shared disk, example locations of the `licenses.txt` file are as follows:

- *Windows:* `S:\NNM\licenses.txt`
- *UNIX:* `/nnmount/NNM/licenses.txt`

Maintaining the HA Configuration

Maintenance Mode

When you need to apply NNMi patches or update to a newer version of NNMi, put the NNMi HA resource group into maintenance mode to prevent failover during the process. When the NNMi HA resource group is in maintenance mode, you (or an installation script) can run the `ovstop` and `ovstart` commands as needed on the primary (active) cluster node.



Never run the `ovstart` or `ovstop` commands on a secondary (backup) cluster node.

Putting an HA Resource Group into Maintenance Mode

Putting an HA resource group into maintenance mode disables HA resource group monitoring. When an HA resource group is in maintenance mode, stopping and starting the products in that HA resource group do not cause failover.

To put an HA resource group into maintenance mode, on the active cluster node, create the following file:

- *Windows*: %NnmDataDir%\hacluster*<resource_group>*\maintenance
- *UNIX*: \$NnmDataDir/hacluster/*<resource_group>*/maintenance



The maintenance file contents are as follows:

- To disable monitoring of the HA resource group, create the maintenance file. The file can be empty or can contain the keyword `NORESTART`.
- To prevent NNMi from starting during a configuration procedure, the first line of the maintenance file must contain only the single word:
`NORESTART`

Removing an HA Resource Group from Maintenance Mode

Taking an HA resource group out of maintenance mode re-enables HA resource group monitoring. Stopping the products in that HA resource group causes the HA resource group to fail over to a passive cluster node.

To remove an HA resource group from maintenance mode, follow these steps:

- 1 Verify that NNMi is running correctly:

```
ovstatus -c
```

All NNMi services should show the state `RUNNING`.

- 2 Delete the `maintenance` file from the node that was the active cluster node before maintenance was initiated. This file is described in [Putting an HA Resource Group into Maintenance Mode](#).

Maintaining NNMi in an HA Cluster

Starting and Stopping NNMi

While NNMi is running under HA, *do not* use the `ovstart` and `ovstop` commands unless instructed to do so for HA maintenance purposes. For normal operation, use the NNMi-provided HA commands or the appropriate HA product commands for starting and stopping HA resource groups.

Changing NNMi Hostnames and IP Addresses in a Cluster Environment

A node in a cluster environment can have more than one IP address and hostname. If a node becomes a member of another subnet, you might need to change its IP addresses. As a result, the IP address or fully-qualified domain name might change.

For example, on UNIX systems, the IP address and the related hostname are generally configured in one of the following:

- `/etc/hosts`
- Domain Name Service (DNS)
- Network Information Service (NIS on HP-UX or Linux, NIS+ on Solaris)

NNMi also configures the hostname and IP address of the management server for the managed node in the NNMi database.

If you are moving from a non-name-server environment to a name-server environment (that is, DNS or BIND), make sure that the name server can resolve the new IP address.

Hostnames work within IP networks to identify a managed node. While a node might have many IP addresses, the hostname is used to pinpoint a specific node. The system hostname is the string returned when you use the `hostname` command.

When changing the virtual hostname or IP address of the NNMi HA resource group, you must update the `licenses.txt` file on the shared disk with the new information from the license file on the active node. Complete the following procedure to correctly update the HA configuration.

To change the virtual hostname or IP address of the NNMi HA resource group, perform these steps on the active NNMi cluster node:

- 1 Convert and the permanent non-production license keys for the prior virtual IP address of the NNMi HA resource group to the new virtual IP address of the NNMi HA resource group.



Do not install the new license keys at this time.

- 2 Put the NNMi HA resource group into maintenance mode as described in [Putting an HA Resource Group into Maintenance Mode](#) on page 336.

3 Stop the NNMi HA resource group:

• *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhastoprg.ovpl NNM \
<resource_group>
```

• *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhastoprg.ovpl NNM \
<resource_group>
```

4 Change the IP address or node name of the NNMi HA resource group:

- a In the `ov.conf` file, edit the `NNM_INTERFACE` entry to be the new hostname or IP address.
- b In the `ovspmd.auth` file, edit any lines containing the old hostname to contain the new hostname.

The `ov.conf` and `ovspmd.auth` files are available in the following location:

- *Windows:* %NnmDataDir%\shared\nnm\conf
- *UNIX:* \$NnmDataDir/shared/nnm/conf

5 If you changed the node name of the NNMi HA resource group, set NNMi to use the new fully-qualified domain name of the NNMi HA resource group with the `nmmsetofficialfqdn.ovpl` command. For example:

```
nmmsetofficialfqdn.ovpl newnnmi.servers.example.com
```

For more information, see the `nmmsetofficialfqdn.ovpl` reference page, or the UNIX manpage.

6 Change the cluster configuration to use the new IP address:

• *MSFC:*

In Failover Cluster Management, open `<resource_group>`.

Double-click `<resource_group>-ip`, select **Parameters**, and then enter the new IP address.

• *Serviceguard:*

On the active HA cluster node, edit the `/etc/cmcluster/<resource_group>/<resource_group>.cntl` file to replace `IP[0]=<old_IP_address>` with `IP[0]=<new_IP_address>`. (If you moved the NNMi HA resource group to a different subnet, also replace `SUBNET[0]=<old_subnet_mask>` with `SUBNET[0]=<new_subnet_mask>`.) Then use `cmapplyconf` to update all other systems.

• *VCS:*

```
$NnmInstallDir/misc/nnm/ha/nnmhargconfigure.ovpl NNM \
<resource_group> -set_value <resource_group>-ip \
Address <new_IP_address>
```

• *RHCS:*

On the active HA cluster node, edit the `/etc/cluster/cluster.conf` file to replace `ip address="<old_IP_address>"` with `ip address="<new_IP_address>"`. Then run `ccs_tool update /etc/cluster/cluster.conf` to update all other systems.

7 Install the permanent non-production license keys for the new virtual IP address of the NNMi HA resource group as described in [Licensing NNMi](#) on page 121.

- 8 Update the `licenses.txt` file on the shared disk with the new information from the `LicFile.txt` file on the active node. Do one of the following:
 - If the `licenses.txt` file exists in the NNM directory on the shared disk, append the new license keys in `LicFile.txt` on the active node to `licenses.txt` on the shared disk.
 - If the `licenses.txt` file does not exist on the shared disk, copy `LicFile.txt` from the active node to `licenses.txt` in the NNM directory on the shared disk.

On the active node, the `LicFile.txt` file is in the following location:

- *Windows:* `<drive>:\Program Files (x86)\HP\HP BTO Software\data\shared\nnm\conf\licensing\LicFile.txt`
- *UNIX:* `/var/opt/OV/shared/nnm/conf/licensing/LicFile.txt`

On the shared disk, example locations of the `licenses.txt` file are as follows:

- *Windows:* `S:\NNM\licenses.txt`
- *UNIX:* `/nnmount/NNM/licenses.txt`

- 9 Start the NNMi HA resource group:
 - *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhastartrg.ovpl NNM \  
<resource_group>
```
 - *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhastartrg.ovpl NNM \  
<resource_group>
```

- 10 Verify that NNMi started correctly:

```
ovstatus -c
```

All NNMi services should show the state RUNNING.

- 11 Take the NNMi HA resource group out of maintenance mode as described in [Removing an HA Resource Group from Maintenance Mode](#) on page 336.

Stopping NNMi Without Causing Failover

When you need to perform NNMi maintenance, you can stop NNMi on the active cluster node without causing failover to a currently passive node. Follow these steps on the active cluster node:

- 1 Put the NNMi HA resource group into maintenance mode as described in [Putting an HA Resource Group into Maintenance Mode](#) on page 336.
- 2 Stop NNMi:

```
ovstop -c
```

Restarting NNMi after Maintenance

If you have stopped NNMi in the manner that prevents failover, follow these steps to restart NNMi and HA monitoring:

- 1 Start NNMi:

```
ovstart -c
```

- 2 Verify that NNMi started correctly:

```
ovstatus -c
```

All NNMi services should show the state `RUNNING`.

- 3 Take the NNMi HA resource group out of maintenance mode as described in [Removing an HA Resource Group from Maintenance Mode](#) on page 336.

Maintaining Add-on NNM iSPIs in an NNMi HA Cluster

The NNM iSPIs are closely linked to NNMi. When add-on NNM iSPIs are installed on the nodes in the NNMi HA cluster, use the NNMi HA cluster maintenance procedures as written.

Unconfiguring NNMi from an HA Cluster

The process of removing an NNMi node from an HA cluster involves undoing the HA configuration for that instance of NNMi. You can then run that instance of NNMi as a standalone management server, or you can uninstall NNMi from that node.

If you want to keep NNMi configured for high availability, the HA cluster must contain one node that is actively running NNMi and at least one passive NNMi node. If you want to completely remove NNMi from the HA cluster, unconfigure the HA functionality on all nodes in the cluster.

To completely unconfigure NNMi from an HA cluster, follow these steps:

- 1 Determine which node in the HA cluster is active. On any node, run the following command:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \  
-group <resource_group> -activeNode
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \  
-group <resource_group> -activeNode
```

- 2 On each passive node, unconfigure any add-on NNM iSPIs from the HA cluster.

For information, see the documentation for each NNM iSPI.

- 3 On any node in the HA cluster, verify that the add-on NNM iSPIs on all passive nodes have been unconfigured from the HA cluster:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \  
-config NNM -get NNM_ADD_ON_PRODUCTS
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \  
-config NNM -get NNM_ADD_ON_PRODUCTS
```

The command output lists the add-on iSPI configurations in the format `<iSPI_PM_Name>[hostname_list]`. For example:

```
PerfSPIHA[hostname1, hostname2]
```

At this time, only the active node hostname should appear in the output. If a passive node hostname appears in the output, repeat [step 2](#) until this command output includes only the active node hostname.

- 4 On each passive node, unconfigure NNMi from the HA cluster:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaunconfigure.ovpl NNM \  
<resource_group>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaunconfigure.ovpl NNM \  
<resource_group>
```

This command removes access to the shared disk but does not unconfigure the disk group or the volume group.

- 5 On each passive node, move the NNMi HA resource group-specific files to a separate location for safe-keeping:



If you do not plan to reconfigure the NNMi HA resource group, you do not need to save a copy of these files, and you can delete them at this time.

- *MSFC:* In Windows Explorer, delete the `%NnmDataDir%\hacluster\<resource_group>` folder.

- *Serviceguard:*

```
rm -rf /var/opt/OV/hacluster/<resource_group>  
rm -rf /etc/cmcluster/<resource_group>
```

- *VCS:*

```
rm -rf /var/opt/OV/hacluster/<resource_group>
```

- *RHCS:*

```
rm -rf /var/opt/OV/hacluster/<resource_group>
```

- 6 On the active node, unconfigure any add-on NNM iSPIs from the HA cluster.

For information, see the documentation for each NNM iSPI. On any node in the HA cluster, verify that the add-on NNM iSPIs on all nodes have been unconfigured from the HA cluster:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \  
-config NNM -get NNM_ADD_ON_PRODUCTS
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \  
-config NNM -get NNM_ADD_ON_PRODUCTS
```

If any hostname appears in the output, repeat [step 6](#) until this command output indicates that no iSPIs are configured.

- 7 On the active node, stop the NNMi HA resource group:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhastoprg.ovpl NNM \  
<resource_group>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhastoprg.ovpl NNM \  
<resource_group>
```

This command does not remove access to the shared disk. Nor does it unconfigure the disk group or the volume group.

- 8 On the active node, unconfigure NNMi from the HA cluster:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaunconfigure.ovpl NNM \  
<resource_group>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaunconfigure.ovpl NNM \  
<resource_group>
```

This command removes access to the shared disk but does not unconfigure the disk group or the volume group.

- 9 On the active node, move the NNMi HA resource group-specific files to a separate location for safe-keeping:



If you do not plan to reconfigure the NNMi HA resource group, you do not need to save a copy of these files, and you can delete them at this time.

- *MSFC:* In Windows Explorer, delete the
%NnmDataDir%\hacluster*<resource_group>*\ folder.

- *Serviceguard:*

```
rm -rf /var/opt/OV/hacluster/<resource_group>  
rm -rf /etc/cmcluster/<resource_group>
```

- *VCS:*

```
rm -rf /var/opt/OV/hacluster/<resource_group>
```

- *RHCS:*

```
rm -rf /var/opt/OV/hacluster/<resource_group>
```

- 10 Unmount the shared disk.

- If you want to reconfigure the NNMi HA cluster at some point, you can keep the disk in its current state.
- If you want to use the shared disk for another purpose, copy all data that you want to keep (as described in [Running NNMi Outside HA with the Existing Database](#) on page 343), and then use the HA product commands to unconfigure the disk group and volume group.

Running NNMi Outside HA with the Existing Database

If you want to run NNMi outside HA on any node with the existing database, follow these steps:

- 1 On the active node (if one still exists), ensure that NNMi is not running:

```
ovstop
```

Alternatively, check the status of the `ovspmd` process by using Task Manager (Windows) or the `ps` command (UNIX).

- 2 On the current node (where you want to run NNMi outside HA), verify that NNMi is not running:

```
ovstop
```



To prevent data corruption, make sure that no instance of NNMi is running and accessing the shared disk.

- 3 (UNIX only) Activate the disk group, for example, on HP-UX Serviceguard:

```
vgchange -a e <disk_group>
```

- 4 Use the appropriate operating system commands to mount the shared disk. For example:

- *Windows:* Use Server Manager—>Disk Management.
- *UNIX:* `mount /dev/vgnrm/lvnm /nnmmount`

- 5 Copy the NNMi files from the shared disk to the local disk:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhadisk.ovpl NNM \  
-from <HA_mount_point>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhadisk.ovpl NNM \  
-from <HA_mount_point>
```

- 6 Use the appropriate operating system commands to unmount the shared disk. For example:

- *Windows:* Use Windows Explorer.
- *UNIX:* `umount /nnmmount`

- 7 (UNIX only) Deactivate the disk group, for example:

```
vgchange -a n <disk_group>
```

- 8 Obtain and install the permanent production license keys for the physical IP address of this NNMi management server as described in [Licensing NNMi](#) on page 121.

- 9 Start NNMi:

```
ovstart -c
```

NNMi is now running with a copy of the database that was formerly used by the NNMi HA resource group. Manually remove from the NNMi configuration any nodes that you do not want to manage from this NNMi management server.

Patching NNMi under HA

To apply a patch for NNMi, work in HA maintenance mode. Follow these steps:

- 1 Determine which node in the HA cluster is active:
 - *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \
-group <resource_group> -activeNode
```
 - *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \
-group <resource_group> -activeNode
```
- 2 On the active node, put the NNMi HA resource group into maintenance mode as described in [Putting an HA Resource Group into Maintenance Mode](#) on page 336. Include the NORESTART keyword.
- 3 On all passive nodes, put the NNMi HA resource group into maintenance mode as described in [Putting an HA Resource Group into Maintenance Mode](#) on page 336. Include the NORESTART keyword.
- 4 On the active node, follow these steps:
 - a Stop NNMi:


```
ovstop -c
```
 - b Back up the shared disk by performing a disk copy.
 - c *Optional.* Use the nnmbackup.ovpl command, or another database command, to back up all NNMi data. For example:



```
nnmbackup.ovpl -type offline -scope all -target nnm_backups
```

 For more information about this command, see [NNMi Backup and Restore Tools](#) on page 383.
 - d Apply the appropriate NNMi and NNM iSPI patches to the system.
 - e Start NNMi:


```
ovstart -c
```
 - f Verify that NNMi started correctly:


```
ovstatus -c
```

 All NNMi services should show the state RUNNING.
- 5 On each passive node, apply the appropriate patches to the system.



Never run the ovstart or ovstop commands on a secondary (backup) cluster node.
- 6 On all passive nodes, take the NNMi HA resource group out of maintenance mode as described in [Removing an HA Resource Group from Maintenance Mode](#) on page 336.
- 7 On the active node, take the NNMi HA resource group out of maintenance mode as described in [Removing an HA Resource Group from Maintenance Mode](#) on page 336.

Upgrading NNMi under HA from NNMi 9.0x/9.1x to NNMi 9.20

Follow the appropriate procedure for your environment:

- [Upgrade NNMi with the Embedded Database on all Supported Operating Systems](#) on page 345
- [Upgrade NNMi with Oracle on all Supported Operating Systems](#) on page 349

Upgrade NNMi with the Embedded Database on all Supported Operating Systems



As of NNMi 9.10, Serviceguard is no longer supported on the Linux operating system. If NNMi is currently running under Serviceguard HA, you cannot follow the procedure in this section. Instead, unconfigure NNMi from HA as described in [Unconfiguring NNMi from an HA Cluster](#) on page 309, upgrade NNMi on all nodes, and then configure NNMi to run under a supported HA product as described in [Configure NNMi for HA](#) on page 321. Alternatively, you can configure NNMi for NNMi application failover as described in [Configuring NNMi for Application Failover](#) on page 283.

Upgrading NNMi includes upgrading the Postgres database software to a newer version. For this reason, NNMi must be taken out of operation for the duration of the upgrade process.



NNMi will be unavailable for approximately 30 to 60 minutes during this upgrade procedure.

To upgrade from NNMi 9.0x or 9.1x under HA to NNMi 9.20 under HA, upgrade the active node to update the embedded database, and then upgrade the passive node while NNMi is still in maintenance mode. Follow these steps:

- 1 Ensure that the NNMi 9.0x or 9.1x configuration is consistent across all HA nodes by forcing a failover, in turn, to each of the passive nodes.
- 2 For NNMi 9.0x, ensure that all nodes are running NNMi 9.0x Patch 5 or a higher version. For NNMi 9.1x, use patch 3 or higher.

If necessary, upgrade each system to the appropriate consolidated patch.

- 3 Check the `ov.conf` files on both systems to ensure that they have the correct values. The `ov.conf` file is available in the following location:

- *Windows*: `%NnmDataDir%\shared\nnm\conf`
- *UNIX*: `$NnmDataDir/shared/nnm/conf`

- 4 Determine which node in the NNMi 9.0x or 9.1x HA cluster is active:

- *Windows*:

```
%NnmInstallDir%\misc\nnm\ha\nnmhaclusterinfo.ovpl \
-group <resource_group> -activeNode
```
- *UNIX*:

```
$NnmInstallDir/misc/nnm/ha/nnmhaclusterinfo.ovpl \
-group <resource_group> -activeNode
```

The remainder of this procedure refers to the currently active node as server X and the currently passive node as server Y.

- 5 For HP-UX systems, on server Y, edit the `/etc/cmcluster/<resource_group>/<resource_group>.mon` file as follows:
 - a Locate the following line:


```
if [ ! -f /var/opt/OV/hacluster/$SHA_RESOURCE_GROUP/maint_NNM -a
! -f /var/opt/OV/hacluster/$SHA_RESOURCE_GROUP/maint_NNM ]
```
 - b Change the second "maint_NNM" to "maintenance".
 - c Failover the application and repeat [step a](#) and [step b](#) on the node that is no longer running the resource group.
- 6 For Windows systems, perform the following:
 - a On server X, stop the `<resource_group>-app` resource.
 - b Check the Access Control Lists (ACLs) on the file `%NnmDataDir%\hacluster\<resource_group>\hamscs.vbs` (be sure to remember these).
 - c Save the `hamscs.vbs` file.
 - d Copy the `%NnmInstallDir%\misc\nnm\ha\nnmhamscs.vbs` script to a temporary directory where you can edit the file.
 - e Open the copy of the `nmhamscs.vbs` file and change all references for `product_name` to be **NNM**. You can reference the original script for the value. Save the `nmhamscs.vbs` file.
 - f As Administrator, copy the updated `nmhamscs.vbs` script to `%NnmDataDir%\hacluster\<resource_group>\hamscs.vbs`.
 - g Check the ACLs again to ensure that they are the same as before.
 - h Start the `<resource_group>-app` resource.
 - i Verify that the resource comes online. If not, check the cluster logs to see if there are any syntax errors. (You can use the following command to generate a cluster log: `cluster log /gen`. If you must specify a folder, you can do so using the following syntax: `cluster log /gen /copy:<my folder>`.)
 - j Run `ovstop`.
- 7 On server X, disable HA resource group monitoring by creating the following maintenance file:
 - *Windows:*

```
%NnmDataDir%\hacluster\<resource_group>\maintenance
```



Ensure that the maintenance file does not have a `.txt` extension, which can occur if the file has been edited with a text editor, such as Notepad.

- *UNIX:*

```
$NnmDataDir/hacluster/<resource_group>/maintenance
```

The file can be empty.

- 8 On server X, upgrade NNMi:
 - a Upgrade NNMi to the current version as described in the *HP Network Node Manager i Software Upgrade Reference*, available on the HP manuals web site.

The database upgrade occurs during this step.

- b To verify that the upgrade completed correctly, enter the following command:

```
ovstart
```

All NNMi services should show the state RUNNING.

- c Upgrade all add-on NNM iSPIs to version 9.20.

For information, see the documentation for each NNM iSPI.



If your environment includes standalone NNM iSPIs, you must also upgrade those products to version 9.20 for correct functionality. You can complete those upgrades after completing this procedure.

- 9 For Windows systems, do the following:

- a Copy the updated `nnmhamscs.vbs` script (see [step f](#) within [step 6](#)) from Server X to `%NnmDataDir%\hacluster\<resource group>\hamscs.vbs` on Server Y.
- b Check the ACLs to ensure that they are the same as before.

- 10 On server X, run the following command: `nnmhadisk.ovpl NNM -replicate`.

- 11 On server Y, disable HA resource group monitoring by creating the following maintenance file:

- *Windows:*

```
%NnmDataDir%\hacluster\<resource_group>\maintenance
```



Ensure that the maintenance file does not have a `.txt` extension, which can occur if the file has been edited with a text editor, such as Notepad.

- *UNIX:*

```
$NnmDataDir/hacluster/<resource_group>/maintenance
```

The file can be empty.

- 12 On server Y, upgrade NNMi:

- a Upgrade NNMi to the current version as described in the *HP Network Node Manager i Software Upgrade Reference*, available on the HP manuals web site.
- b Verify that the upgrade completed without error.
- c Upgrade all add-on NNM iSPIs to version 9.20.

For information, see the documentation for each NNM iSPI.

- 13 If the HA cluster includes multiple passive nodes, repeat [step 12](#) for each passive node.

- 14 For HP-UX systems, on the node not running the resource group, run the following commands:

```
cd /etc/cmcluster/<resource_group>
```

```
cp <resource_group>.mon <resource_group>.mon.save
```

```
cp /opt/OV/misc/nnm/ha/mcsg/NNM/rg.mon <resource_group>.mon
```

- 15 On server X, delete the maintenance file:

- *Windows:*

```
%NnmDataDir%\hacluster\<resource_group>\maintenance
```

- *UNIX*:
`$NnmDataDir/hacluster/<resource_group>/maintenance`

16 Perform the following post-installation steps:

- a Verify that the following variables are set:

```

NNM_INTERFACE
HA_MOUNT_POINT
NNM_ADD_ON_PRODUCTS
HA_LOCALE (not required if running in C)

```

These variables are defined in the following locations:

HP-UX Serviceguard:

```
/etc/cmcluster/<resource_group>/<resource_group>.public.env
```

Veritas:

```
/opt/VRTSvcs/bin/hagrp -display | grep UserStrGlobal
```

Windows: Using regedit, the values are in the following location:

```
HKEY_LOCAL_MACHINE\Cluster\Groups\<group>\Parameters
```

- b If the variables are not set, you can run the following commands for each missing value:

```
/opt/OV/misc/nnm/ha/nnmhaclusterinfo.ovpl -config NNM -set
NNM_INTERFACE <value for NNM_INTERFACE>
```

```
/opt/OV/misc/nnm/ha/nnmhaclusterinfo.ovpl -config NNM -set
HA_MOUNT_POINT <value for HA_MOUNT_POINT>
```

```
/opt/OV/misc/nnm/ha/nnmhaclusterinfo.ovpl -config NNM -set
NNM_ADD_ON_PRODUCTS <value for NNM_ADD_ON_PRODUCTS>
```

```
/opt/OV/misc/nnm/ha/nnmhaclusterinfo.ovpl -config NNM -set
HA_LOCALE <value for HA_LOCALE>
```



HA_LOCALE is only needed if you are attempting to use a localized language.

17 For all Linux HA upgrades, run the following sets of commands, as applicable for your system:

— *RHEL*:

```
rm /etc/rc.d/rc*.d/S98netmgt
```

```
rm /etc/rc.d/rc*.d/K01netmgt
```

— *SuSE*:

```
rm /etc/init.d/rc*.d/S98netmgt
```

```
rm /etc/init.d/rc*.d/K01netmgt
```



When using Windows Server 2008 R2, the Network Name resource may have the name "Network Name". This name should be the short name for the virtual IP address. If applicable, change the name as follows:

- 1 Using Failover Cluster Management, select the Network Name resource.
- 2 Right-click and select **Properties**.
- 3 Change the name.

Upgrade NNMi with Oracle on all Supported Operating Systems

To upgrade NNMi for HA in an Oracle environment, follow the procedure described in [Upgrade NNMi with the Embedded Database on all Supported Operating Systems](#) on page 345.

Troubleshooting the HA Configuration

This section includes the following topics:

- [Common Configuration Mistakes](#) on page 350
- [Configuration Issues with RHCS 6](#) on page 351
- [HA Resource Testing](#) on page 351
- [General HA Troubleshooting](#) on page 352
- [NNMi-Specific HA Troubleshooting](#) on page 356
- [NNM iSPI-Specific HA Troubleshooting](#) on page 360

Common Configuration Mistakes

Some common HA configuration mistakes are listed here:

- Incorrect disk configuration
 - VCS: If a resource cannot be probed, the configuration is somehow wrong. If a disk cannot be probed, the disk might no longer be accessible by the operating system.
 - Test the disk configuration manually and confirm against HA documentation that the configuration is appropriate.
- The disk is in use and cannot be started for the HA resource group.
Always check that the disk is not activated before starting the HA resource group.
- MSFC: Bad network configuration
If network traffic is flowing across multiple NIC cards, RDP sessions fail when activating programs that consume a large amount of network bandwidth, such as the NNMi ovjboss process.
- Some HA products do not automatically restart at boot time.
Review the HA product documentation for information about how to configure automatic restart on boot up.
- Adding NFS or other access to the OS directly (resource group configuration should be managing this).
- Being in the shared disk mount point during a failover or offlining of the HA resource group.
HA kills any processes that prevent the shared disk from being unmounted.
- Reusing the HA cluster virtual IP address as the HA resource virtual IP address (works on one system and not the other)
- Timeouts are too short. If the products are misbehaving, HA product might time out the HA resource and cause a failover.
MSFC: In Failover Cluster Management, check the value of the **Time to wait for resource to start** setting. NNMi sets this value to 15 minutes. You can increase the value.

- Not using maintenance mode
Maintenance mode was created for debugging HA failures. If you attempt to bring a resource group online on a system, and it fails over shortly afterwards, use the maintenance mode to keep the resource group online to see what is failing.
- Not reviewing cluster logs (cluster logs can show many common mistakes).

Configuration Issues with RHCS 6

It is possible for the `/etc/cluster/cluster.conf` file versions to differ between the two systems in an HA environment if the `ricci` service is down or has been intentionally disabled. Therefore, monitor the `cluster.conf` file regularly to ensure that the file versions are synchronized.

If the `cluster.conf` file versions are not synchronized, you may experience problems when you attempt to do any of the following:

- apply changes to `cluster.conf`
- unconfigure a resource group
- start the cluster
- use the `clustat` command

HA Resource Testing

This section describes the general approach for testing the resources that you will place into the NNMi HA resource group. This testing identifies hardware configuration problems. It is recommended to perform this testing *before* configuring NNMi to run under HA. Note the configuration values that generate positive results, and use these value when performing the complete configuration of the NNMi HA resource group.

For specific details regarding any of the commands listed here, see the most recent documentation for your HA product.

To test HA resources, follow these steps:

- 1 If necessary, start the HA cluster.
- 2 (Windows only) Verify that the following virtual IP addresses have been defined for the HA cluster:
 - A virtual IP address for the HA cluster
 - A virtual IP address for each HA resource groupEach of these IP addresses should not be used elsewhere.
- 3 Add an HA resource group to the HA cluster.
Use a non-production name, such as `test`, for this HA resource group.
- 4 Test the connection to the HA resource group:
 - a Add the virtual IP address and corresponding virtual hostname for the resource group as a resource to the HA resource group.
Use the values that you will later associate with the NNMi HA resource group.

- b Fail over from the active cluster node to the passive cluster node to verify that the HA cluster correctly fails over.
 - c Fail over from the new active cluster node to the new passive cluster node to verify failback.
 - d If the resource group does not fail over correctly, log on to the active node, and then verify that the IP address is properly configured and accessible. Also verify that no firewall blocks the IP address.v
- 5 Configure the shared disk as described in [Configuring a SAN or a Physically Connected Disk](#) on page 333.
- 6 Test the connection to the shared disk:
 - a Add the shared disk as a resource to the HA resource group as described in [Moving the Shared Disk into the NNMi HA Resource Group](#) on page 334.
 - b Fail over from the active cluster node to the passive cluster node to verify that the HA cluster correctly fails over.
 - c Fail over from the new active cluster node to the new passive cluster node to verify failback.
 - d If the resource group does not fail over correctly, log on to the active node, and then verify that the disk is mounted and available.
- 7 Keep a record of the commands and inputs that you used to configure the shared disk. You might need this information when configuring the NNMi HA resource group.
- 8 Remove the resource group from each node:
 - a Remove the IP address entry.
 - b Offline the resource group, and then remove resource group from the node.

At this point, you can use the NNMi-provided tools to configure NNMi to run under HA.

General HA Troubleshooting

The topics in this section apply to HA configuration for NNMi and the NNM iSPIs. They include:

- [Error: Wrong Number of Arguments](#)
- [Resource Hosting Subsystem Process Stops Unexpectedly \(Windows Server 2008 R2\)](#)
- [Product Startup Times Out \(Solaris\)](#)
- [Log Files on the Active Cluster Node Are Not Updating](#)
- [Cannot Start the NNMi HA Resource Group on a Particular Cluster Node](#)

Error: Wrong Number of Arguments

The name of the product Perl module is a required parameter to most of the NNMi HA configuration commands.

- For NNMi, use the value `NNM`.

- To determine what value to use for an NNM iSPI, see the documentation for that NNM iSPI.

Resource Hosting Subsystem Process Stops Unexpectedly (Windows Server 2008 R2)

Starting an HA cluster resource on a computer running the Windows Server 2008 R2 operating system stops the Resource Hosting Subsystem (Rhs.exe) process unexpectedly.

For information about this known problem, see the Microsoft Support web site article *The Resource Hosting Subsystem (Rhs.exe) process stops unexpectedly when you start a cluster resource in Windows Server 2008 R2*, which is available from **<http://support.microsoft.com/kb/978527>**.



Always run the NNMi resource in a separate resource monitor (rhs.exe) specific to the resource group.

Product Startup Times Out (Solaris)

One or more of the `/var/adm/messages*` files contains a message similar to the following example:

```
VCS ERROR V-16-1-13012 Thread(...) Resource(<resource_group>-app):
online procedure did not complete within the expected time.
```

This message indicates that the product did not start completely within the Veritas timeout value. The NNMi-provided HA configuration scripts define this timeout to be 15 minutes.

You can adjust the Veritas timeout value. For example, to change the Veritas timeout value to 30 minutes (1800 seconds), you can run the following commands in order:

```
/opt/VRTSvcs/bin/haconf -makerw
/opt/VRTSvcs/bin/hares -modify <resource_group>-app OnlineTimeout 1800
/opt/VRTSvcs/bin/haconf -dump -makero
```

Product Startup Times Out (Windows MSCS 2008)

After upgrading from NNMi 9.0x, if the app resource (`<resource>-app`) in the Failover Cluster Manager changes from "Pending" to "Failed", there might be a timeout issue. If this situation occurs, do the following:

- 1 Use the `cluster log /gen` command to generate the `cluster.log` file.
- 2 Open the log located in the following directory:
C:\Windows\cluster\reports\cluster.log
- 3 If you see an error in the `cluster.log` file similar to the following, you have a `DeadlockTimeout` issue:

```
ERR [RHS] Resource <resource-name>-APP handling deadlock. Cleaning
current operation.
```

- 4 The `DeadlockTimeout` is the total time for failover when the agent might be blocked. The `PendingTimeout` represents either the online or offline operation. The `DeadlockTimeout` default value is 45 minutes (2,700,000 milliseconds), and the `PendingTimeout` default value is 30 minutes (1,800,000 milliseconds). You can change the `DeadlockTimeout` and the `PendingTimeout` values. For example, to set a `DeadlockTimeout` of 75 minutes and a `PendingTimeout` of 60 minutes, you can run the following commands:

```
cluster res "<resource_group>-APP" /prop DeadlockTimeout=4500000
cluster res "<resource_group>-APP" /prop PendingTimeout=3600000
```

See your HA vendor documentation for more information.

Log Files on the Active Cluster Node Are Not Updating

This situation is normal. It occurs because the log files have been redirected to the shared disk.

For NNMi, review the log files in the location specified by `HA_NNM_LOG_DIR` in the `ov.conf` file.

Cannot Start the NNMi HA Resource Group on a Particular Cluster Node

If the `nmhastartrg.ovpl` or `nmhastartrg.ovpl` command does not correctly start, stop, or switch the NNMi HA resource group, review the following information:

- *MSFC*:
 - In Failover Cluster Management, review the state of the NNMi HA resource group and underlying resources.
 - Review the Event Viewer log for any errors.
- *Serviceguard*:

Review the `<resource_group>.cntl.log` file and the syslog files for errors. The most common problems are leaving the system in a state where a resource cannot be added, for example, having a disk group misconfigured such that it cannot be activated.

```
/etc/cmcluster/<resource_group>/<resource_group>.cntl.log
```
- *VCS*:
 - Run `/opt/VRTSvcs/bin/hares -state` to review the resource state.
 - For failed resources, review the `/var/VRTSvcs/log/<resource>.log` file for the resource that is failing. Resources are referenced by the agent type, for example: `IP*.log`, `Mount*.log`, and `Volume*.log`.
- *RHCS*:

Review the `<resource_group>.cntl.log` file and the syslog files for errors. The most common problems are leaving the system in a state where a resource cannot be added, for example, having a disk group misconfigured such that it cannot be activated.

```
/etc/cmcluster/<resource_group>/<resource_group>.cntl.log
```

If you cannot locate the source of the problem, you can manually start the NNMi HA resource group by using the HA product commands:

- 1 Mount the shared disk.
- 2 Assign the virtual host to the network interface:
 - *MSF*:
 - Start Failover Cluster Management.
 - Expand the resource group.
 - Right-click `<resource_group>-ip`, and then click **Bring Online**.
 - *Serviceguard*: Run `/usr/sbin/cmmodnet` to add the IP address.
 - *VCS*: `/opt/VRTSvcs/bin/hares -online <resource_group>-ip \ -sys <local_hostname>`
 - *RHCS*: Run `/usr/sbin/cmmodnet` to add the IP address.
- 3 Start the NNMi HA resource group. For example:
 - *Windows*:


```
%NnmInstallDir%\misc\nnm\ha\nnmhastartrg.ovpl NNM \
          -start <resource_group>
```
 - *UNIX*:


```
$NnmInstallDir/misc/nnm/ha/nnmhastartrg.ovpl NNM \
          -start <resource_group>
```

The return code 0 indicates that NNMi started successfully.

The return code 1 indicates that NNMi did not start correctly.

NNMi-Specific HA Troubleshooting

The topics in this section apply to HA configuration for NNMi only. They include:

- [Re-Enable NNMi for HA after All Cluster Nodes are Unconfigured](#)
- [NNMi Does Not Start Correctly Under HA](#)
- [Changes to NNMi Data are Not Seen after Failover](#)
- [nmsdbmgr Does Not Start after HA Configuration](#)
- [pmd Does Not Start after HA Configuration](#)
- [NNMi Runs Correctly on Only One HA Cluster Node \(Windows\)](#)
- [Disk Failover Does Not Occur](#)
- [Shared Disk is Not Accessible \(Windows\)](#)
- [Shared Disk Does Not Contain Current Data](#)
- [Shared Disk Files Are Not Found by the Secondary Node after Failover](#)

Re-Enable NNMi for HA after All Cluster Nodes are Unconfigured

When all NNMi HA cluster nodes have been unconfigured, the `ov.conf` file no longer contains any mount point references to the NNMi shared disk. To re-create the mount point reference without overwriting the data on the shared disk, follow these steps on the primary node:

- 1 If NNMi is running, stop it:

```
ovstop -c
```

- 2 Reset the reference to the shared disk:

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhadisk.ovpl NNM \
-setmount <HA_mount_point>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhadisk.ovpl NNM \
-setmount <HA_mount_point>
```

- 3 In the `ov.conf` file, verify the entries related to HA mount points.

For the location of the `ov.conf` file, see [NNMi HA Configuration Files](#) on page 361.

NNMi Does Not Start Correctly Under HA

When NNMi does not start correctly, it is necessary to debug whether the issue is a hardware issue with the virtual IP address or the disk, or whether the issue is some form of application failure. During this debug process, put the system in maintenance mode *without* the `NORESTART` keyword.

- 1 On the active node in the HA cluster, disable HA resource group monitoring by creating the following maintenance file:

- *Windows:* `%NnmDataDir%\hacluster\<resource_group>\maintenance`
- *UNIX:* `$NnmDataDir/hacluster/<resource_group>/maintenance`

2 Start NNMI:

```
ovstart
```

3 Verify that NNMI started correctly:

```
ovstatus -c
```

All NNMI services should show the state RUNNING. If this is not the case, troubleshoot the process that does not start correctly.

4 After completing your troubleshooting, delete the maintenance file:

- *Windows:* %NnmDataDir%\hacluster*<resource_group>*\maintenance
- *UNIX:* \$NnmDataDir/hacluster/*<resource_group>*/maintenance

Changes to NNMI Data are Not Seen after Failover

The NNMI configuration points to a different system than where NNMI is running. To fix the problem, verify that the `ov.conf` file has appropriate entries for the following items:

- `NNM_INTERFACE=<virtual_hostname>`
- `HA_RESOURCE_GROUP=<resource_group>`
- `HA_MOUNT_POINT=<HA_mount_point>`
- `NNM_HA_CONFIGURED=YES`
- `HA_POSTGRES_DIR=<HA_mount_point>/NNM/dataDir/shared/nnm/databases/Postgres`
- `HA_EVENTDB_DIR=<HA_mount_point>/NNM/dataDir/shared/nnm/eventdb`
- `HA_CUSTOMPOLLER_DIR=<HA_mount_point>/NNM/dataDir/shared/nnm/databases/custompoller`
- `HA_NNM_LOG_DIR=<HA_mount_point>/NNM/dataDir/log`
- `HA_JBOSS_DATA_DIR=<HA_mount_point>/NNM/dataDir/nmsas/NNM/data`
- `HA_LOCALE=C`

For the location of the `ov.conf` file, see [NNMI HA Configuration Files](#) on page 361.

nmsdbmgr Does Not Start after HA Configuration

This situation usually occurs as a result of starting NNMI after running the `nmhaconfigure.ovpl` command but without the `nmhadisk.ovpl` command with the `-to` option having been run. In this case, the `HA_POSTGRES_DIR` entry in the `ov.conf` file specifies the location of the embedded database on the shared disk, but this location is not available to NNMI.

To fix this problem, follow these steps:

- 1 On the active node in the HA cluster, disable HA resource group monitoring by creating the following maintenance file:
 - *Windows:* %NnmDataDir%\hacluster*<resource_group>*\maintenance
 - *UNIX:* \$NnmDataDir/hacluster/*<resource_group>*/maintenance
- 2 Copy the NNMI database to the shared disk:
 - *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhadisk.ovpl NNM \
-to <HA_mount_point>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhadisk.ovpl NNM \
-to <HA_mount_point>Start the NNMi HA resource group:
```



To prevent database corruption, run this command (with the `-to` option) only one time. For information about alternatives, see [Re-Enable NNMi for HA after All Cluster Nodes are Unconfigured](#) on page 356.

- *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nnmhastartrg.ovpl NNM \
<resource_group>
```

- *UNIX:*

```
$NnmInstallDir/misc/nnm/ha/nnmhastartrg.ovpl NNM \
<resource_group>
```

- 3 Start NNMi:

```
ovstart
```

- 4 Verify that NNMi started correctly:

```
ovstatus -c
```

All NNMi services should show the state RUNNING.

- 5 After completing your troubleshooting, delete the maintenance file:

- *Windows:* %NnmDataDir%\hacluster*<resource_group>*\maintenance
- *UNIX:* \$NnmDataDir/hacluster/*<resource_group>*/maintenance

pmd Does Not Start after HA Configuration

This situation usually occurs after a configuration error such as not setting up the shared disk correctly. The failure of the `pmd` process occurs when the `ovjboss` process does not fully start.

Review the following log file:

- *Windows:* %HA_MOUNT_POINT%\NNM\dataDir\log\nnm\jbossServer.log
- *UNIX:* \$HA_MOUNT_POINT/NNM/dataDir/log/nnm/jbossServer.log

NNMi Runs Correctly on Only One HA Cluster Node (Windows)

The Windows operating system requires two different virtual IP addresses, one for the HA cluster and one for the HA resource group. If the virtual IP address of the HA cluster is the same as that of the NNMi HA resource group, NNMi only runs correctly on the node associated with the HA cluster IP address.

To correct this problem, change the virtual IP address of the HA cluster to a unique value for the network.

Disk Failover Does Not Occur

This situation can happen when the operating system does not support the shared disk. Review the HA product, operating system, and disk manufacturer documentation to determine whether these products can all work together.

If disk failure occurs, NNMI does not start on failover. Most likely, `nmsdbmgr` fails because the `HA_POSTGRES_DIR` directory does not exist. Verify that the shared disk is mounted and that the appropriate files are accessible.

Shared Disk is Not Accessible (Windows)

The command `nmhaclusterinfo.ovpl -config NNM -get HA_MOUNT_POINT` returns nothing.

The drive of the shared disk mount point must be fully specified (for example, `S:\`) during HA configuration.

To correct this problem, run the `nmhaconfigure.ovpl` command on each node in the HA cluster. Fully specify the drive of the shared disk mount point.

Shared Disk Does Not Contain Current Data

Responding to the `nmhaconfigure.ovpl` command question about disk type with the text `none` bypasses the code for setting the disk-related variables in the `ov.conf` file. To fix this situation, follow the procedure in [Prepare the Shared Disk Manually](#) on page 332.

Shared Disk Files Are Not Found by the Secondary Node after Failover

The most common cause of this situation is that the `nmhadisk.ovpl` command was run with the `-to` option when the shared disk was not mounted. In this case, the data files are copied to the local disk, so the files are not available on the shared disk.

To fix this problem, follow these steps:

- 1 On the active node in the HA cluster, disable HA resource group monitoring by creating the following maintenance file:
 - *Windows:* `%NnmDataDir%\hacluster\<resource_group>\maintenance`
 - *UNIX:* `$NnmDataDir/hacluster/<resource_group>/maintenance`
- 2 Log on to the active node, and then verify that the disk is mounted and available.
- 3 Stop NNMI:


```
ovstop
```
- 4 Copy the NNMI database to the shared disk:
 - *Windows:*

```
%NnmInstallDir%\misc\nnm\ha\nmhadisk.ovpl NNM \  
-to <HA_mount_point>
```

- *UNIX*:

```
$NnmInstallDir/misc/nnm/ha/nnmhadisk.ovpl NNM \  
-to <HA_mount_point>
```



To prevent database corruption, run this command (with the `-to` option) only one time. For information about alternatives, see [Re-Enable NNMi for HA after All Cluster Nodes are Unconfigured](#) on page 356.

5 Start the NNMi HA resource group:

- *Windows*:

```
%NnmInstallDir%\misc\nnm\ha\nnmhastartrg.ovpl NNM \  
<resource_group>
```
- *UNIX*:

```
$NnmInstallDir/misc/nnm/ha/nnmhastartrg.ovpl NNM \  
<resource_group>
```

6 Start NNMi:

```
ovstart
```

7 Verify that NNMi started correctly:

```
ovstatus -c
```

All NNMi services should show the state `RUNNING`.

8 After completing your troubleshooting, delete the maintenance file:

- *Windows*: %NnmDataDir%\hacluster*<resource_group>*\maintenance
- *UNIX*: \$NnmDataDir/hacluster/*<resource_group>*/maintenance

NNM iSPI-Specific HA Troubleshooting

For information about troubleshooting an NNM iSPI running under HA, see the documentation for that NNM iSPI.

HA Configuration Reference

NNMi HA Configuration Files

Table 31 lists the NNMi HA configuration files. These files apply to NNMi and add-on NNM iSPIs on the NNMi management server. These files are installed to the following location:

- *Windows*: %NnmDataDir%\shared\nnm\conf
- *UNIX*: \$NnmDataDir/shared/nnm/conf

Table 31 NNMi HA Configuration Files

File Name	Description
ov.conf	Updated by the <code>nmhaclusterinfo.ovpl</code> command to describe the NNMi HA implementation. NNMi processes read this file to determine the HA configuration.
nmdatareplicator.conf	Used by the <code>nmdatareplicator.ovpl</code> command to determine which NNMi folders and files are included in data replication from the active node to the passive nodes. If you implement a different method of replicating the NNMi configuration, see this file for a list of the data to include. For more information, see the comments in the file.

NNMi-Provided HA Configuration Scripts

Table 32 and Table 33 list the HA configuration scripts that are included with NNMi. The NNMi-provided scripts listed in Table 32 are convenience scripts that can be used to configure HA for any product that has a customer Perl module. If you prefer, you can use the HA product-provided commands to configure HA for NNMi.

On the NNMi management server, the NNMi-provided HA configuration scripts are installed to the following location:

- *Windows*: %NnmInstallDir%\misc\nnm\ha
- *UNIX*: \$NnmInstallDir/misc/nnm/ha

Table 32 NNMi HA Configuration Scripts

Script Name	Description
nmhaconfigure.ovpl	Configures NNMi or an NNM iSPI for an HA cluster. Run this script on all nodes in the HA cluster.
nmhaunconfigure.ovpl	Unconfigures NNMi or an NNM iSPI from an HA cluster. Optionally, run this script on one or more nodes in the HA cluster.
nmhaclusterinfo.ovpl	Retrieves cluster information regarding NNMi. Run this script as needed on any node in the HA cluster.

Table 32 NNMi HA Configuration Scripts (cont'd)

Script Name	Description
nmhadisk.ovpl	Copies NNMi and NNM iSPI data files to and from the shared disk. During HA configuration, run this script on the primary node. At other times, run this script per the instructions in this chapter.
nmhastartrg.ovpl	Starts the NNMi HA resource group in an HA cluster. During HA configuration, run this script on the primary node.
nmhastoprg.ovpl	Stops the NNMi HA resource group in an HA cluster. During HA unconfiguration, run this script on the primary node.

The NNMi-provided scripts listed in [Table 33](#) are used by the scripts listed in [Table 32](#) on page 361. Do not run the scripts listed in [Table 33](#) directly.

Table 33 NNMi HA Support Scripts

Script Name	Description
nmdatareplicator.ovpl	Checks the nmdatareplicator.conf configuration file for changes and copies files to remote systems.
nmharg.ovpl	Starts, stops, and monitors NNMi in an HA cluster. For Serviceguard configurations, used by <resource_group>.cntl. For VCS configurations, used by the VCS start, stop, and monitor scripts. (nmhargconfigure.ovpl configures this usage.) Also used by nmhastartrg.ovpl to enable and disable tracing.
nmhargconfigure.ovpl	Configures HA resources and resource groups. Used by nmhaconfigure.ovpl and nmhaunconfigure.ovpl.
nmhastart.ovpl	Starts NNMi in an HA cluster. Used by nmharg.ovpl.
nmhastop.ovpl	Stops NNMi in an HA cluster. Used by nmharg.ovpl.
nmhamonitor.ovpl	Monitors NNMi processes in an HA cluster. Used by nmharg.ovpl.
nmhamscs.vbs	Is a template for creating a script to start, stop, and monitor NNMi processes in a MSFC HA cluster. The generated script is used by MSFC and is stored in the following location: %NmDataDir%\hacluster\ <resource_group>\hamscs.vbs</resource_group>

NNMi HA Configuration Log Files

The following log files apply to the HA configuration for NNMi and add-on NNM iSPIs on the NNMi management server:

- **Windows configuration:**
 - %NnmDataDir%\tmp\HA_nnmhaserver.log
 - %NnmDataDir%\log\haconfigure.log
- **UNIX configuration:**
 - \$NnmDataDir/tmp/HA_nnmhaserver.log
 - \$NnmDataDir/log/haconfigure.log
- **Windows runtime:**
 - **Event Viewer log**
 - %HA_MOUNT_POINT%\NNM\dataDir\log\nnm\ovspmd.log
 - %HA_MOUNT_POINT%\NNM\dataDir\log\nnm\public\postgres.log
 - %HA_MOUNT_POINT%\NNM\dataDir\log\nnm\public\nmsdbmgr.log
 - %HA_MOUNT_POINT%\NNM\dataDir\log\nnm\jbossServer.log
 - %SystemRoot%\Cluster\cluster.log
This is the log file for cluster runtime issues including: adding and removing resources and resource groups; other configuration issues; starting and stopping issues.
- **HP-UX runtime:**
 - /etc/cmcluster/<resource_group>/<resource_group>.cntl.log
This is the log file for the resource group.
 - /var/adm/syslog/syslog.log
 - /var/adm/syslog/OLDSyslog.log
 - \$HA_MOUNT_POINT/NNM/dataDir/log/nnm/ovspmd.log
 - \$HA_MOUNT_POINT/NNM/dataDir/log/nnm/public/postgres.log
 - \$HA_MOUNT_POINT/NNM/dataDir/log/nnm/public/nmsdbmgr.log
 - \$HA_MOUNT_POINT/NNM/dataDir/log/nnm/jbossServer.log

- *Linux or Solaris runtime for VCS:*


Table 34 Linux or Solaris Runtime for VCS

Resource	Log File
<code><resource_group>-app</code>	<ul style="list-style-type: none"> • <code>/var/VRTSvcs/log/Application_A.log</code> • <code>\$HA_MOUNT_POINT/NNM/dataDir/log/nnm/ovspmd.log</code> • <code>\$HA_MOUNT_POINT/NNM/dataDir/log/nnm/public/postgres.log</code> • <code>\$HA_MOUNT_POINT/NNM/dataDir/log/nnm/public/nmsdbmgr.log</code> • <code>\$HA_MOUNT_POINT/NNM/dataDir/log/nnm/jbossServer.log</code> • <code>/var/adm/messages*</code>
<code><resource_group>-dg</code> <code><resource_group>-volume</code> <code><resource_group>-mount</code>	<ul style="list-style-type: none"> • <code>/var/VRTSvcs/log/DiskGroup_A.log</code> • <code>/var/VRTSvcs/log/Volume_A.log</code> • <code>/var/VRTSvcs/log/Mount_A.log</code> • <code>/var/adm/messages*</code>
<code><resource_group>-ip</code>	<ul style="list-style-type: none"> • <code>/var/VRTSvcs/log/IP_A.log</code> • <code>/var/adm/messages*</code>

For operating system-specific issues related to the HA resources, review the `/var/adm/messages*` files. For `<resource_group>-app`, look for messages regarding unable to start process.

- *Linux runtime for RHCS:*
 - `/var/adm/syslog/syslog.log`
 - `$HA_MOUNT_POINT/NNM/dataDir/log/nnm/ovspmd.log`
 - `$HA_MOUNT_POINT/NNM/dataDir/log/nnm/public/postgres.log`
 - `$HA_MOUNT_POINT/NNM/dataDir/log/nnm/public/nmsdbmgr.log`
 - `$HA_MOUNT_POINT/NNM/dataDir/log/nnm/jbossServer.log`

NNMi Northbound Interface



HP Network Node Manager i Software (NNMi) provides the NNMi northbound interface for forwarding NNMi incidents to any application that can receive SNMPv2c traps. For each NNMi management server, you can implement the NNMi northbound interface to multiple northbound applications, each configured separately.

NNMi includes support for using the NNMi northbound interface to integrate with the following products:

- The Operations Management functionality of the HP Business Service Management (BSM) platform.
- The HP Operations Manager (HPOM) active messages browser.
- IBM Tivoli Netcool/OMNIbus.
- HP ArcSight Logger

To integrate with a different northbound application, follow the instructions in this chapter.

This chapter contains the following topics:

- [NNMi Northbound Interface](#)
- [Enabling the NNMi Northbound Interface](#)
- [Using the NNMi Northbound Interface](#)
- [Changing the NNMi Northbound Interface](#)
- [Disabling the NNMi Northbound Interface](#)
- [Troubleshooting the NNMi Northbound Interface](#)
- [Application Failover and the NNMi Northbound Interface](#)
- [NNMi Northbound Interface Destination Form Reference](#)

NNMi Northbound Interface

The NNMi northbound interface forwards NNMi management events as SNMPv2c traps to a northbound application. The northbound application might filter, act on, and show the NNMi traps. The northbound application might also provide tools for accessing the NNMi console in the context of an NNMi trap.

The NNMi northbound interface can send incident lifecycle state change notifications, incident correlation notifications, and incident deletion notifications to the northbound application. In this way, the northbound application can replicate the results of NNMi causal analysis.

The NNMi northbound interface can also forward the SNMP traps that NNMi receives to the northbound application. The NNMi northbound interface does not forward events generated by NNM 6.x or 7.x management stations to the northbound application.

Value

The NNMi northbound interface enables event consolidation in a third-party or custom event consolidator. The NNMi northbound interface enriches events with information that can be used to integrate other applications with NNMi.

Supported Versions

The information in this chapter applies to NNMi version 9.00 or higher.

For the most recent information about supported hardware platforms and operating systems, see the *NNMi System and Device Support Matrix*.

Terminology

This chapter uses the following terms:

- Northbound application—Any application that can receive and process SNMPv2c traps.
- Trap-receiving component—The portion of a northbound application that receives SNMP traps.
 - Some applications include a separately installable component that receives SNMP traps and forwards them to another component for processing.
 - For any northbound application that does not include such a component, “trap-receiving component” is synonymous with “northbound application.”
- NNMi northbound interface—The NNMi functionality that forwards NNMi incidents as SNMPv2c traps to a northbound application.
- Northbound destination—One configuration of the NNMi northbound interface that defines the connection to the trap-receiving component of a northbound application and specifies the types of traps that NNMi will send to that northbound application.

Documentation

This chapter describes how to configure NNMi to forward NNMi incidents to any northbound application. For information about a particular northbound application, see that application's documentation.

Enabling the NNMi Northbound Interface



NNMi does not limit the amount of information sent in an SNMP trap using UDP. If any network hardware in the transmission path cannot handle the size of the trap data, or if network traffic is heavy, the trap might be lost. Therefore, it is recommended that the trap-receiving component of the northbound application be installed on the NNMi management server. The northbound application is responsible for ensuring reliable information transfer.

To enable the NNMi northbound interface, follow these steps:

- 1 If necessary, configure the northbound application to understand the NNMi trap definitions.
- 2 On the NNMi management server, configure NNMi incident forwarding:
 - a In the NNMi console, open the **HP NNMi–Northbound Interface Destinations** form (**Integration Module Configuration > Northbound Interface**), and then click **New**.
(If you have selected an available destination, click **Reset** to make the **New** button available.)
 - b Select the **Enabled** check box to make the remaining fields on the form available.
 - c Enter the information for connecting to the northbound application.
For information about these fields, see [Northbound Application Connection Parameters](#) on page 375.
 - d Specify the sending options and incident filter for the content to send to the northbound application.
For information about these fields, see [NNMi Northbound Interface Integration Content](#) on page 376.
 - e Click **Submit** at the bottom of the form.
A new window opens, showing a status message. If the message indicates a problem with the settings, click **Return**, and then adjust the values as suggested by the text of the error message.
- 3 *Optional.* Create contextual interaction with NNMi by creating URLs that provide access to NNMi views from the northbound application.
For information, in the NNMi console, click **Help > NNMi Documentation Library > Integrate NNMi Elsewhere with URLs**.

Using the NNMi Northbound Interface

When the NNMi northbound interface is enabled, the northbound destination determines the information that NNMi sends to a northbound application. Configure the northbound application to show and interpret the forwarded traps, as appropriate in your network environment. For complete information about the contents and format of the traps that NNMi sends to a northbound application, see the `hp-nnmi-nbi.mib` and `hp-nnmi-registration.mib` files.

NNMi sends only one copy of each management event, SNMP trap, or notification trap to a northbound destination. NNMi does not queue traps. If the trap-receiving component of a northbound application is unavailable when NNMi forwards a trap, the trap is lost.

This section describes the types of traps the integration can send. For information about setting the content configuration, see [NNMi Northbound Interface Integration Content](#) on page 376.

Incident Forwarding

Management events

When the northbound destination includes management events, NNMi forwards each management event incident to the northbound application when that incident changes to the REGISTERED lifecycle state.

The OID of the forwarded management event is the SNMP Object ID on the **Management Event Configuration** form in the NNMi console. NNMi forwards all custom management events with the OID 1.3.6.1.4.1.11.2.17.19.2.0.9999.

Third-party SNMP traps

When the northbound destination includes third-party SNMP traps, NNMi forwards each incoming SNMPv1, v2c, or v3 format trap to the northbound application when the associated incident changes to the REGISTERED lifecycle state. NNMi preserves the original trap varbinds in order (as defined in the MIB) and appends the NNMi-specific varbinds to the message payload. If the original trap does not contain all of the defined varbinds, NNMi pads NULL values for the missing varbinds. If the MIB is not loaded in NNMi, NNMi cannot correctly reconstruct the trap and append the NNMi incident data; therefore, NNMi does not forward this trap.

For third-party SNMP traps, note the following:

- Because NNMi reconstructs a trap from its SNMP trap incident, the forwarded trap is in SNMPv2c format regardless of the format the original trap was in when NNMi received it.
- The forwarded SNMP trap shows the NNMi management server as the source object. To determine the original source object, examine the values of the (n+21)th varbind, IncidentNodeHostname (1.3.6.1.4.1.11.2.17.19.2.2.21) and the (n+24)th varbind, IncidentNodeMgmtAddr (1.3.6.1.4.1.11.2.17.19.2.2.24), where n is the number of varbinds defined for the trap in the MIB.

If any of the devices that NNMi manages also send traps to the northbound application, the northbound application must manage the duplicate device traps.

For a comparison of trap forwarding mechanisms, see *Trap and Incident Forwarding* in the *NNMi Deployment Reference*.

Incident Lifecycle State Change Notifications

The information in this section varies with the selections made to the **Sending Options** in the **HP NNMi–Northbound Interface Destination** page.

Enhanced closed traps

When the northbound destination includes enhanced closed notifications, NNMi sends an `EventLifecycleStateClosed` (1.3.6.1.4.1.11.2.17.19.2.0.1000) trap to the northbound application when the lifecycle state of an incident changes to `CLOSED` in NNMi. The `EventLifecycleStateClosed` trap includes much of the data from the original incident. The previous lifecycle state value is not included. The `EventLifecycleStateClosed` trap identifies the original incident in the sixth varbind, `IncidentUuid` (1.3.6.1.4.1.11.2.17.19.2.2.6).

State change traps

When the northbound destination includes lifecycle state changed notifications, NNMi sends a `LifecycleStateChangeEvent` (1.3.6.1.4.1.11.2.17.19.2.0.1001) trap to the northbound application when the lifecycle state of an incident changes to the `IN PROGRESS`, `COMPLETED`, or `CLOSED` lifecycle state in NNMi. The northbound application can associate the `LifecycleStateChangeEvent` with the original incident.

The `LifecycleStateChangeEvent` trap identifies the original incident and the lifecycle state change in the following varbinds:

- `IncidentUuid`, the sixth varbind (1.3.6.1.4.1.11.2.17.19.2.2.6)
This value matches the value of the sixth varbind in a management event or the (n+6)th varbind in a third-party SNMP trap varbind.
- `IncidentLifecycleStatePreviousValue`, the seventh varbind (1.3.6.1.4.1.11.2.17.19.2.2.200)
- `IncidentLifecycleStateCurrentValue`, the eighth varbind (1.3.6.1.4.1.11.2.17.19.2.2.201)

The following table lists the possible integer values for lifecycle state.

Name	Integer Value
registered	1
inprogress	2
completed	3
closed	4
dampened	5

Incident Correlation Notifications

When the northbound destination includes incident correlation notifications, NNMi sends incident correlation traps to the northbound application as NNMi causal analysis correlates incidents. The northbound application can use the information in the traps to replicate the correlation changes.

Single correlation traps

For the single correlation trap option, the integration sends the following correlation traps:

- `EventDedupCorrelation` (1.3.6.1.4.1.11.2.17.19.2.0.1100)
- `EventImpactCorrelation` (1.3.6.1.4.1.11.2.17.19.2.0.1101)

- EventPairwiseCorrelation (1.3.6.1.4.1.11.2.17.19.2.0.1102)
- EventRateCorrelation (1.3.6.1.4.1.11.2.17.19.2.0.1103)
- EventApaCorrelation (1.3.6.1.4.1.11.2.17.19.2.0.1104)
- EventCustomCorrelation (1.3.6.1.4.1.11.2.17.19.2.0.1105)

Each trap identifies one parent-child incident correlation relationship in the following varbinds:

- IncidentCorrelationIndicatorParentUuid, the sixth varbind (1.3.6.1.4.1.11.2.17.19.2.2.6)
- IncidentCorrelationIndicatorChildUuid, the seventh varbind (1.3.6.1.4.1.11.2.17.19.2.2.300)

Group correlation traps

For the group correlation option, the integration sends the following correlation traps:

- EventDedupCorrelationGroup (1.3.6.1.4.1.11.2.17.19.2.0.2100)
- EventImpactCorrelationGroup (1.3.6.1.4.1.11.2.17.19.2.0.2101)
- EventPairwiseCorrelationGroup (1.3.6.1.4.1.11.2.17.19.2.0.2102)
- EventRateCorrelationGroup (1.3.6.1.4.1.11.2.17.19.2.0.2103)
- EventApaCorrelationGroup (1.3.6.1.4.1.11.2.17.19.2.0.2104)
- EventCustomCorrelationGroup (1.3.6.1.4.1.11.2.17.19.2.0.2105)

Each trap identifies the parent-child incident correlation relationships in the following varbinds:

- IncidentCorrelationIndicatorParentUuid, the sixth varbind (1.3.6.1.4.1.11.2.17.19.2.2.6)
- IncidentCorrelationIndicatorChildCount, the seventh varbind (1.3.6.1.4.1.11.2.17.19.2.2.301)
- IncidentCorrelationIndicatorChildUuidCsv, the eighth varbind (1.3.6.1.4.1.11.2.17.19.2.2.302)

This value is a comma-separated-value list of child incident UUIDs.

Incident Deletion Notifications

When the northbound destination includes incident deletion notifications, NNMi sends an EventDeleted (1.3.6.1.4.1.11.2.17.19.2.0.3000) trap to the northbound application when an incident is deleted in NNMi. The EventDeleted trap identifies the original incident in the sixth varbind, IncidentUuid (1.3.6.1.4.1.11.2.17.19.2.2.6).

Event Forwarding Filter

When the northbound destination includes an incident filter, the object identifiers (OIDs) in the filter include or exclude (depending on the selected configuration option) the following event types:

- NNMi management event incidents
- Third-party SNMP traps
- EventLifecycleStateClosed traps
- LifecycleStateChangeEvent traps

- EventDeleted traps
- Correlation notification traps

The following notes apply to correlation notification traps:

- If the incident filter prevents the forwarding of the parent incident for a correlation, NNMi does not send a correlation notification trap to the northbound application.
- If the incident filter prevents the forwarding of a child incident for a correlation, the forwarded correlation notification trap does not include that child incident's UUID. (If the correlation notification trap would not contain any child incident UUIDs, NNMi does not send that trap to the northbound application.)
- The DuplicateCorrelation management event is forwarded independently of the EventDedupCorrelation or EventDedupCorrelationGroup correlation notification traps. Likewise, the RateCorrelation management event is forwarded independently of the EventRateCorrelation or EventRateCorrelationGroup correlation notification traps. If the incident filter prevents the forwarding of one of these correlation notification traps, NNMi might still forward the associated management events.

Changing the NNMi Northbound Interface

To change the NNMi northbound interface configuration parameters, follow these steps:

- 1 In the NNMi console, open the **HP NNMi–Northbound Interface Destinations** form (**Integration Module Configuration > Northbound**).
- 2 Select a destination, and then click **Edit**.
- 3 Modify the values as appropriate.

For information about the fields on this form, see [NNMi Northbound Interface Destination Form Reference](#) on page 374.

- 4 Verify that the **Enabled** check box at the top of the form is selected, and then click **Submit** at the bottom of the form.

The changes take effect immediately.

Disabling the NNMi Northbound Interface

No SNMP trap queuing occurs while a northbound destination is disabled.

To discontinue the forwarding of NNMi incidents to a northbound application, follow these steps:

- 1 In the NNMi console, open the **HP NNMi–Northbound Interface Destinations** form (**Integration Module Configuration > Northbound**).
- 2 Select a destination, and then click **Edit**.



Alternatively, click **Delete** to entirely remove the configuration for the selected destination.

- 3 Clear the **Enabled** check box at the top of the form, and then click **Submit** at the bottom of the form.

The changes take effect immediately.

Troubleshooting the NNMi Northbound Interface

If the NNMi northbound interface is not working as expected, follow these steps until you have resolved the problem:

- 1 Verify that the trap destination port is not blocked by a firewall.
Ensure that the NNMi management server can directly address the northbound application by host and port.
- 2 Verify that the integration is running correctly:
 - a In the NNMi console, open the **HP NNMi–Northbound Interface Destinations** form (**Integration Module Configuration > Northbound**).
 - b Select a destination, and then click **Edit**.
 - c Verify that the **Enabled** check box is selected.
- 3 If the northbound destination includes management events, verify this functionality:
 - a In the **Closed Key Incidents** view of the NNMi console, open any incident.
 - b Set the incident lifecycle state to **Registered**, and then click  **Save**.
 - c Set the incident lifecycle state to **Closed**, and then click  **Save and Close**.
 - d After 30 seconds, determine whether the northbound application received an EventLifecycleStateClosed trap (or a LifecycleStateChangeEvent trap) for this incident should.
 - If the northbound application received the trap, continue with [step 4](#).
 - If the northbound application did not receive the trap, configure a new northbound destination to connect with a different northbound application, and then repeat this test from [step a](#).

If the repeated test succeeds, the problem is with the first northbound application. Consult that application’s documentation for troubleshooting information.

If the repeated test fails, contact HP Support for assistance.
- 4 If the northbound destination includes SNMP traps, verify this functionality:
 - a Generate an SNMP trap against a node in the NNMi topology by entering the following command on the NNMi management server:

```
nnmsnmpnotify.ovpl -u username -p password -a \  
discovered_node NNMi_node linkDown
```

Where *discovered_node* is the hostname or IP address of a node in the NNMi topology and *NNMi_node* is the hostname or IP address of the NNMi management server.

- b After 30 seconds, determine whether the northbound application received the forwarded trap.
 - If the northbound application received the trap, the NNMi northbound interface is working correctly.
 - If the northbound application did not receive the trap, configure a new northbound destination to connect with a different northbound application, and then repeat this test from [step a](#).

If the repeated test succeeds, the problem is with the first northbound application. Consult that application's documentation for troubleshooting information.

If the repeated test fails, contact HP Support for assistance.

Application Failover and the NNMi Northbound Interface

If the NNMi management server will participate in NNMi application failover, the information in this topic applies to any integration that implements the NNMi northbound interface for sending traps to a northbound application.

The traps that NNMi sends to a northbound application include the NNMi URL in the NmsUrl varbind (1.3.6.1.4.1.11.2.17.19.2.2.2). Traps received before application failover reference what is now the standby NNMi management server. When the URL points to the standby NNMi management server, any actions that use the URL value (for example, launching the NNMi console) will fail.

Local Northbound Application

If the trap-receiving component of the northbound application is located on the NNMi management server, the following considerations apply to the configuration of the NNMi northbound interface:

- The trap-receiving component of the northbound application must be installed and configured identically on the active and standby NNMi management servers. Configure SNMP trap reception on the same port on both NNMi management servers.
- Configure the NNMi northbound interface on the primary NNMi management server only.

On the **HP NNMi–Northbound Interface Destination** form, select either the **NNMi FQDN** or the **Use Loopback** option for **Host** identification.

At startup, the NNMi northbound interface determines the correct name or IP address of the current NNMi management server. In this way, the northbound interface sends traps to the trap-receiving component of the northbound application on the active NNMi management server.

Remote Northbound Application

If the trap-receiving component of the northbound application is not located on the NNMi management server, configure the NNMi northbound interface on the primary NNMi management server only. On the **HP NNMi–Northbound Interface Destination** form, select the **Other** option for **Host** identification.

NNMi Northbound Interface Destination Form Reference

The **HP NNMi–Northbound Interface Destination** form contains the parameters for configuring communications between NNMi and a northbound application. This form is available from the **Integration Module Configuration** workspace. (On the **HP NNMi–Northbound Interface Destinations** form, click **New**; or select a destination, and then click **Edit**.)



Only NNMi users with the Administrator role can access the **HP NNMi–Northbound Interface Destination** form.

The **HP NNMi–Northbound Interface Destination** form contains information for the following areas:

- [Northbound Application Connection Parameters](#) on page 375
- [NNMi Northbound Interface Integration Content](#) on page 376
- [NNMi Northbound Interface Destination Status Information](#) on page 378

To apply changes to the integration configuration, update the values on the **HP NNMi–Northbound Interface Destination** form, and then click **Submit**.

Northbound Application Connection Parameters

Table 35 lists the parameters for configuring the connection to the northbound application.

Table 35 Northbound Application Connection Information

Field	Description
Host	<p>The fully-qualified domain name (preferred) or the IP address of the server that contains the trap-receiving component of the northbound application.</p> <p>The integration supports the following methods for identifying the server:</p> <ul style="list-style-type: none"> • NNMi FQDN NNMi manages the connection to the northbound application on the NNMi management server and the Host field becomes read-only. This is the recommended configuration for northbound applications on the NNMi management server. • Use Loopback NNMi manages the connection to the northbound application on the NNMi management server and the Host field becomes read-only. • Other Enter a hostname or IP address for identifying the northbound application server in the Host field. NNMi validates that the hostname or IP address in the Host field is not configured as a loopback adapter. This is the default configuration. <p>NOTE: If the NNMi management server participates in NNMi application failover, see Application Failover and the NNMi Northbound Interface on page 373 for information about the impact of application failover on the integration.</p>
Port	<p>The UDP port where the northbound application receives SNMP traps.</p> <p>Enter the port number specific to the northbound application.</p> <p>NOTE: If the trap-receiving component of the northbound application is on the NNMi management server, this port number must be different from the port NNMi uses to receive SNMP traps, as set in the SNMP Port field on the Communication Configuration form in the NNMi console.</p>
Community String	<p>A read-only community string for the northbound application to receive traps.</p> <p>If the northbound application configuration requires a community string in the received SNMP traps, enter that value.</p> <p>If the northbound application configuration does not require a specific community string, use the default value, which is <code>public</code>.</p>

NNMi Northbound Interface Integration Content

Table 36 lists the parameters for configuring the content the NNMi northbound interface sends to the northbound application.

Table 36 NNMi Northbound Interface Content Configuration Information

Field	Description
Incidents	<p>The incident forwarding specification.</p> <ul style="list-style-type: none"> • Management NNMi forwards only NNMi-generated management events to the northbound application. • 3rd Party SNMP Trap NNMi forwards only SNMP traps that NNMi receives from managed devices to the northbound application. • Syslog NNMi forwards only ArcSight Syslog messages that NNMi receives from managed devices to the northbound application using the NorthBound Integration module. <p>NNMi begins forwarding incidents as soon as you enable the northbound destination.</p> <p>For more information, see Incident Forwarding on page 368.</p>
Lifecycle State Changes	<p>The incident change notification specification.</p> <ul style="list-style-type: none"> • Enhanced Closed NNMi sends an incident closed trap to the northbound application for each incident that changes to the CLOSED lifecycle state. This is the default configuration. • State Changed NNMi sends an incident lifecycle state changed trap to the northbound application for each incident that changes to the IN PROGRESS, COMPLETED, or CLOSED lifecycle state. • Both NNMi sends an incident closed trap to the northbound application for each incident that changes to the CLOSED lifecycle state. Additionally, the integration sends an incident lifecycle state changed trap to the northbound application for each incident that changes to the IN PROGRESS, COMPLETED, or CLOSED lifecycle state. <p>NOTE: In this case, each time an incident changes to the CLOSED lifecycle state, the integration sends two notification traps: an incident closed trap and an incident lifecycle state changed trap.</p> <p>For more information, see Incident Lifecycle State Change Notifications on page 369.</p>

Table 36 NNMi Northbound Interface Content Configuration Information (cont'd)

Field	Description
Correlations	<p>The incident correlation notification specification.</p> <ul style="list-style-type: none"> • None NNMi does not notify the northbound application of incident correlations resulting from NNMi causal analysis. This is the default configuration. • Single NNMi sends a trap for each parent-child incident correlation relationship resulting from NNMi causal analysis. • Group NNMi sends one trap per correlation that lists all child incidents correlated to a parent incident. <p>For more information, see Incident Correlation Notifications on page 369.</p>
Deletions	<p>The incident deletion specification. This selection configures whether to send a deletion trap to the northbound application for the selections made in the Incidents field.</p> <ul style="list-style-type: none"> • Don't Send NNMi does not notify the northbound application when incidents are deleted in NNMi. This is the default configuration. • Send NNMi sends a deletion trap to the northbound application for each incident that is deleted in NNMi. <p>For more information, see Incident Deletion Notifications on page 370.</p>
NNMi Console Access	<p>The connection protocol specification in the URL for browsing to the NNMi console from the northbound application. The traps that NNMi sends to the northbound application include the NNMi URL in the NmsUrl varbind (1.3.6.1.4.1.11.2.17.19.2.2.2).</p> <p>The configuration page defaults to the setting that matches the NNMi configuration. If the NNMi console is configured to accept both HTTP and HTTPS connections, you can change the HTTP connection protocol specification in the NNMi URL. For example, if all users of the northbound application are on the intranet, you can set NNMi console access from the northbound application to be over HTTP. To change the protocol for connecting to the NNMi console from the northbound application, select the HTTP option or the HTTPS option as appropriate.</p>

Table 36 NNMi Northbound Interface Content Configuration Information (cont'd)

Field	Description
Incident Filters	<p>A list of object identifiers (OIDs) the integration uses to filter the events sent to the northbound application. Each filter entry can be a valid numeric OID (for example, .1.3.6.1.6.3.1.1.5.4.1.3.6.1.4.1.9) or OID prefix (for example, .1.3.6.1.6.3.1.1.5.*).</p> <p>Select one of the following options:</p> <ul style="list-style-type: none"> • None NNMi sends all events to the northbound application. This is the default configuration. • Include NNMi sends only the specific events that match the OIDs identified in the filter. • Exclude NNMi sends all events except for the specific events that match the OIDs identified in the filter. <p>Specify the incident filter:</p> <ul style="list-style-type: none"> • To add a filter entry, enter the text in the lower text box, and then click Add. • To delete a filter entry, select that entry from the list in the upper box, and then click Remove. <p>For more information, see Event Forwarding Filter on page 370.</p>

NNMi Northbound Interface Destination Status Information

Table 37 lists the read-only status information for the northbound destination. This information is useful for verifying that the integration is working correctly.

Table 37 NNMi Northbound Interface Destination Status Information

Field	Description
Trap Destination IP Address	The IP address the destination host name resolves to. This value is unique to this northbound destination.
Uptime (seconds)	<p>The time (in seconds) since the northbound component was last started. The traps that NNMi sends to a northbound application include this value in the sysUptime field (1.3.6.1.2.1.1.3.0).</p> <p>This value is the same for all integrations that use the NNMi northbound interface. To see the latest value, either refresh or close and re-open the form.</p>
NNMi URL	<p>The URL for connecting to the NNMi console. The traps that NNMi sends to a northbound application include this value in the NmsUrl varbind (1.3.6.1.4.1.11.2.17.19.2.2.2).</p> <p>This value is unique to this northbound destination.</p>

MIB Information used by the NNMi Northbound Interface

Complete the following steps to load specific MIBs into NNMi, then view the management information used for incident notifications sent by the NNMi northbound integration.

- 1 From a command prompt, run the **nnmloadmib.ovpl -load hp-nnmi.mib** command to load the `hp-nnmi.mib` file.
- 2 From a command prompt, run the **nnmloadmib.ovpl -load p-nnmi-registration.mib** command to load the `hp-nnmi-registration.mib` file.
- 3 From a command prompt, run the **nnmloadmib.ovpl -load hp-nnmi-nbi.mib** command to load the `hp-nnmi-nbi.mib` file.
- 4 Optional Step: From a command prompt, run the **nnmloadmib.ovpl -load hp-nnmi-isp-perf-nbi.mib** command to load the `hp-nnmi-isp-perf-nbi.mib` file.
- 5 From the NNMi console, open the **Configuration** workspace.
- 6 Click **MIBs->Loaded MIBs**.
- 7 Double-click each of the MIBs you just loaded; then click **MIB Variables** to view the MIB information.

Maintaining NNMi

This section contains the following chapters:

- [NNMi Backup and Restore Tools](#)
- [Maintaining NNMi](#)
- [NNMi Logging](#)
- [Running NNMi in a Xen Virtualization Environment](#)

NNMi Backup and Restore Tools

A good backup and restore strategy is key to ensuring the uninterrupted operations of any business. HP Network Node Manager i Software (NNMi) is an important asset for network operations and should be backed up regularly.

The two types of critical data related to an NNMi installation are as follows:

- Files in the file system
- Data in the relational database (embedded or external)

This chapter explains the tools that NNMi provides for backing up and restoring important NNMi files and data.

This chapter contains the following topics:

- [Backup and Restore Commands](#)
- [Backing up NNMi Data](#)
- [Restoring NNMi Data](#)
- [Backup and Restore Strategies](#)
- [Backing up and Restoring the Embedded Database Only](#)

Backup and Restore Commands

NNMi provides the following scripts for backing up and restoring NNMi data:

- `nmbackup.ovpl`—Backs up all necessary file system data (including configuration information) and any data stored in the NNMi embedded database.
- `nmrestore.ovpl`—Restores a backup that was created by using the `nmbackup.ovpl` script.
- `nmbackupembdb.ovpl`—Creates a complete backup of the NNMi embedded database (but not the file system data) while NNMi is running.
- `nmrestoreembdb.ovpl`—Restores a backup that was created by using the `nmbackupembdb.ovpl` script.
- `nmresetembdb.ovpl`—Drops the NNMi embedded database tables. Run the `ovstart` command to recreated the tables.

For command syntax, see the appropriate reference page, or the UNIX manpage.

Backing up NNMi Data

The NNMi backup command (`nmbackup.ovpl`) copies key NNMi file system data and some or all of the tables in the NNMi Postgres database to the specified target directory. The NNMi backup command can create a tar archive of the backup data, or you can compress the backup files using your own tools. You can then use any appropriate tool to save a copy of the backup.



If your NNMi implementation uses Oracle for the main NNMi database, the NNMi backup and restore commands work with the NNMi file system data only. External database maintenance should be handled as part of the existing database backup and restore procedures.

The back up and restore data might or might not include data from any NNM iSPIs installed in your network environment. Check the documentation that came with each NNM iSPI for details.



Any software that locks files (for example, anti-virus or system backup software), can interrupt NNMi access to the NNMi database. This can cause problems such as an inability to read from or write to a file that is being used by another process, such as an anti-virus application. For the NNMi Postgres database, configure these applications to exclude the NNMi database directory (`%NNM_DB%` on Windows, and `$NNM_DB` on UNIX). Use `nmbackup.ovpl` to back up the NNMi database regularly.

Backup Type

The NNMi backup command supports two types of backups:

- Online backups occur while NNMi is running. NNMi ensures that the database tables are synchronized in the backed up data. Operators can be actively using the NNMi console and other processes can be interacting with the NNMi database during an online backup. With an online backup, you can back up all NNMi data or only some of the data according to function, as described in [Backup Scope](#). For

the embedded NNMi database, the `nmsdbmgr` service must be running. For an external database, the backup includes NNMi file system data. NNMi processes do not have to be running to back up an external database.

- Offline backups occur while NNMi is completely stopped. With an offline backup, the backup scope applies to the file system files only. An offline backup always includes the complete NNMi database regardless of the backup scope. For the embedded NNMi database, the backup copies the Postgres database files. For an external database, the backup includes NNMi file system data only.

Backup Scope

The NNMi backup command provides several scopes that define how much NNMi is backed up.

Configuration scope

The configuration scope (`-scope config`) loosely aligns to the information in the **Configuration** workspace of the NNMi console.

The configuration scope includes the following data:

- For online backups, only those embedded database tables that store NNMi configuration information.
- For offline backups, the entire embedded database.
- For all backups, the NNMi configuration information in the file system as listed in [Table 38](#).

Topology scope

The topology scope (`-scope topology`) loosely aligns to the information in the **Inventory** workspace of the NNMi console. Because the network topology is dependent on the configuration that was used for discovering that topology, the topology scope includes the configuration scope.

The topology scope includes the following data:

- For online backups, only those embedded database tables that store NNMi configuration and network topology information.
- For offline backups, the entire embedded database.
- For all backups, the NNMi configuration information in the file system as listed in [Table 38](#). Currently, there are no file system files associated with the topology scope.

Event scope

The event scope (`-scope event`) loosely aligns to the information in the **Incident Browsing** workspace of the NNMi console. Because events are dependent on the network topology related to those events, the event scope includes the configuration and topology scopes.

The event scope includes the following data:

- For online backups, only those embedded database tables that store NNMi configuration, network topology, and event information.
- For offline backups, the entire embedded database.
- For all backups, the NNMi configuration information in the file system as listed in [Table 38](#) and the NNMi event information as listed in [Table 39](#).

All scope

The complete backup (`-scope all`) includes all important NNMi files and the complete embedded database.

Table 38 Configuration Scope Files and Directories

Directory or File name	Description
%NnmInstallDir%/conf (Windows only)	Configuration information
%NnmInstallDir%\misc\nms\lic \$NnmInstallDir/misc/nms/lic	Miscellaneous license information
%NnmInstallDir%\nmsas\server\nms\conf \$NnmInstallDir/nmsas/server/nms/conf	jboss configuration
%NnmDataDir%\conf \$NnmDataDir/conf	Configuration that might be shared by other HP products
%NnmDataDir%\conf\nnm\props \$NnmDataDir/conf/nnm/props	Local NNMi configuration properties files
<ul style="list-style-type: none"> • <i>Windows Server 2008:</i> <drive>:\Program Files (x86)\HP\HP BTO Software\data\shared\nnm\conf\licensing\LicFile.txt • <i>UNIX:</i> /var/opt/OV/shared/nnm/conf/licensing/LicFile.txt 	License information
%NnmDataDir%\NNMVersionInfo \$NnmDataDir/NNMVersionInfo	NNMi version information file
%NnmDataDir%\shared\nnm\user-snmp-mibs \$NnmDataDir/shared/nnm/user-snmp-mibs	Shared user-added SNMP MIB information
%NnmDataDir%\shared\nnm\actions \$NnmDataDir/shared/nnm/actions	Shared lifecycle transition actions
%NnmDataDir%\shared\nnm\certificates \$NnmDataDir/shared/nnm/certificates	Shared NNMi SSL certificates
%NnmDataDir%\shared\nnm\conf \$NnmDataDir/shared/nnm/conf	Shared NNMi configuration information
%NnmDataDir%\shared\nnm\conf\licensing \$NnmDataDir/shared/nnm/conf/licensing	Shared NNMi license configuration information
%NnmDataDir%\shared\nnm\lrf \$NnmDataDir/shared/nnm/lrf	Shared NNMi component registration files
%NnmDataDir%\shared\nnm\conf\props \$NnmDataDir/shared/nnm/conf/props	Shared NNMi configuration properties files
%NnmDataDir%\shared\nnm\www\htdocs\images \$NnmDataDir/shared/nnm/www\htdocs/images	Shared background images for NNMi node group maps

In this context, files in the shared directories are those shared with another NNMi management server in an NNMi application failover or high availability environment.

Table 39 Event Scope Files and Directories

Directory or File name	Description
\$NnmDataDir/log/nnm/signin.0.0.log	NNMi console sign-in log

Restoring NNMi Data

The NNMi restore script (`nmrestore.ovpl`) places the backup data on the NNMi management server. The type and scope of the backup determines what NNMi can restore.

- ▶ If you use the `nmrestore.ovpl` script to place database records on a second NNMi management server, both NNMi management servers must have the same type of operating system and NNMi version and patch level.

Placing the backup data from one NNMi management server onto a second NNMi management server means that both servers have the same database UUID. After you restore NNMi on the second NNMi management server, uninstall NNMi from the original NNMi management server.

- To restore an online backup, NNMi copies the file system data to the correct locations and overwrites the contents of the database tables that were included in the backup. Objects that have been deleted since the backup are restored, and objects that have been created since the backup are deleted. Additionally, any objects that were changed after the backup was taken revert to their state at the time of the backup. For the embedded NNMi database, the `nmsdbmgr` service must be running. For an external database, the restore includes NNMi file system data only and no NNMi processes must be running.
- To restore an offline backup, NNMi overwrites the Postgres files in the file system, completely replacing the database files with the contents of the backup. For an external database, the backup includes NNMi file system data only.

With the `-force` option, the `nmrestore.ovpl` command stops all NNMi processes, starts the `nmsdbmgr` service (if restoring from an online backup of the NNMi embedded database), restores the data, and then restarts all NNMi processes.

If the provided source is a tar file, the NNMi restore command extracts the tar file to a temporary folder in the current working directory. In this case, either ensure that the current working directory has adequate storage to support the temporary folder, or extract the archive before running the restore command.

- ▶ Because the database schema might change from one version of NNMi to the next, data backups cannot be shared across versions of NNMi.

- ▶ Note the following:

- Updates to Status and Incidents could be delayed as NNMi resynchronizes following a restore from backup.
- If you see the following message during this resynchronization, it does not indicate a problem:

The Causal Engine's large queue size is causing delayed updates to Status and Incidents. This could be due to resynchronization following an upgrade, application failover, restore from backup, or a manual resynchronization.
- Do not stop NNMi during this resynchronization. To ensure resynchronization has completed, keep NNMi running for several hours following the restore from backup.

Same System Restore

You can use the backup and restore commands on a single system for data recovery. The following items must not have changed between the time of the backup and time of the restore:

- NNMi version (including any patches)
- Operating system type
- Character set (language)
- Hostname
- Domain

Different System Restore

You can use the backup and restore commands to transfer data from one NNMi management server to another. The intended uses of different system restoration include recovering from system failure and transferring NNMi to a different system during an operating system upgrade.

Best practice

Because the NNMi UUID is copied to the target system during the database restore, both source and target systems now appear to be running the same instance of NNMi. Uninstall NNMi from the source system.



To create multiple functional NNMi management servers with similar configurations, such as while deploying global network management, use the `nnmconfigexport.ovpl` and `nnmconfigimport.ovpl` commands.

For a different system restore, the following items must be identical on both systems:

- NNMi version (including any patches)
- Operating system type and version
- Character set (language)

The following items can differ between the two systems:

- Hostname
- Domain

For a different system restore, the `nnmrestore.ovpl` command does not copy the license information to the new system. Obtain and apply a new license for the new NNMi management server. For more information, see [Licensing NNMi](#) on page 121.

Backup and Restore Strategies

Back up All Data Periodically

Your disaster recovery plan should include a regularly scheduled complete backup of all NNMi data. You do not need to shut down NNMi to create this backup. If you incorporate the backup into a script, use the `-force` option to ensure that NNMi is on the correct state before the backup begins. For example:

```
nmmbackup.ovpl -force -type online -scope all -archive
               -target nmi_backups\periodic
```

If you must recover your NNMi data after a hardware failure, follow these steps:

- 1 Rebuild or acquire new hardware.
- 2 Install NNMi to the same version and patch level as were in place for the backup.
- 3 Restore the NNMi data:

- If the recovery NNMi management server meets the requirements listed in [Same System Restore](#) on page 389, run a command similar to the following example:

```
nmrestore.ovpl -force -lic
               -source nmi_backups\periodic\newest_backup
```

- If the recovery NNMi management server does not qualify for a same-system restore but meets the requirements listed in [Different System Restore](#) on page 389, run a command similar to the following example:

```
nmrestore.ovpl -force
               -source nmi_backups\periodic\newest_backup
```

Update the licensing as needed.

Back up Data Before Changing the Configuration

Perform scoped backups (as described in [Backup Scope](#) on page 385) as needed before beginning configuration changes. In this way, if your configuration changes do not have the expected effect, you will be able to revert to a known working configuration. For example:

```
nmmbackup.ovpl -type online -scope config
               -target nmi_backups\config
```

To restore this backup to the same NNMi management server, stop all NNMi processes, and then run a command similar to the following example:

```
nmrestore.ovpl -force -source nmi_backups\config\newest_backup
```

Back up Data Before Upgrading NNMi or the Operating System

Before making major system changes (including upgrading NNMi or the operating system), perform a complete backup of all NNMi data. To ensure that no changes are made to the NNMi database after the backup is made, stop all NNMi processes and create an offline backup. For example:

```
nmmbackup.ovpl -type offline -scope all
               -target nmi_backups\offline
```

If NNMi does not run correctly after the system change, roll back the change or set up a different NNMi management server and ensure that the requirements listed in [Different System Restore](#) on page 389 are met. Then run a command similar to the following example:

```
nmrestore.ovpl -lic -source nmi_backups\offline\newest_backup
```

Restore File System Files Only

To overwrite NNMi files without affecting the database tables, run a command similar to the following example:

```
nmmrestore.ovpl -partial
-source nmi_backups\offline\newest_backup
```

The command is useful when the NNMi implementation uses Oracle for the main NNMi database.

Backing up and Restoring the Embedded Database Only

NNMi provides the `nmbbackupembdb.ovpl` and `nmmrestoreembdb.ovpl` commands to back up and restore the NNMi embedded database only. This functionality is useful for creating a snapshot of the data as you experiment with NNMi configuration settings. The `nmbbackupembdb.ovpl` and `nmmrestoreembdb.ovpl` commands perform online backups only. At a minimum, the `nmsdbmgr` service must be running.

Best practice

Run the `nmmresetembdb.ovpl` command before restoring data to the embedded database. This command ensures that the database does not contain any errors, thereby eliminating the possibility of encountering database constraint violations. For information about running the embedded database reset command, see the `nmmresetembdb.ovpl` reference page, or the UNIX manpage.

Using Backup and Restore Tools in an HA Environment

Some helpful tips to consider when using backup and restore tools in an HA environment include:

Backup

- Perform a backup using the active (primary) system. (A backup of the backup (secondary) node is not recommended because configuration files could be out-of-date and no shared disk information would be included because the backup node cannot access to the shared disk.)
- The shared disk must be connected to the active node. If using a cron job, verify that the shared disk is mounted.
- Put the system into maintenance mode (so as not to trigger a failover).
- Perform an online backup using the `nmbbackup.ovpl` script on the active node only.
- Periodically, run an offline backup.
- See the `nmbbackup.ovpl` reference page, or the UNIX manpage, for more information.

Restore

- Verify that the shared disk is mounted.
- Verify that the system is in maintenance mode.
- Perform the restore using the `nnmrestore.ovpl` script.
- See the *nnmrestore.ovpl* reference page, or the UNIX manpage, for more information.

For more information on using NNMi in an HA environment, see [Configuring NNMi in a High Availability Cluster](#) on page 313.

Maintaining NNMi

After you have your NNMi management server functioning, there are maintenance tasks you can perform to optimize several of the NNMi features.

This chapter contains the following topics:

- [Administering Access Control Lists for NNMi Folders](#)
- [Administering a Custom Poller Collection Export](#)
- [Administering Incident Actions](#)
- [Blocking Trap Storms using the hosted-on-trapstorm.conf File](#)
- [Blocking Incidents using the trapFilter.conf File](#)
- [Configuring Character Set Encoding Settings for NNMi](#)
- [Configuring NNMi to Require Encryption for Remote Access](#)
- [Configuring the Auto-Trim Oldest SNMP Trap Incidents Feature](#)
- [Modifying NNMi Gauge Titles to Show SNMP MIB Variable Names](#)
- [Modifying NNMi Normalization Properties](#)
- [Modifying Simultaneous SNMP Requests](#)
- [Modifying the Embedded Database Port](#)
- [Modifying MIB Browser Parameters](#)
- [NNMi Self Monitoring](#)
- [Suppressing the Use of Discovery Protocols for Specific Nodes](#)
- [Suppressing the Use of VLAN-indexing for Large Switches](#)
- [Configuring Node Component Status](#)

Administering Access Control Lists for NNMi Folders

You might run across situations that would cause you to modify the user name that runs the HP NNM Action Server as shown in [Setting the Action Server Name Parameter](#) on page 398. If you change the user name that runs the action server without modifying the user name permissions, the HP NNM Action Server might not start, and NNMi might not log messages when running incident actions. This section includes actions to take to prevent this from happening.

NNMi (Everest) contains permission changes to the following directories:

- `/var/opt/OV/log/nnm/public`
- `/var/opt/OV/shared/perfSpi`

Although the NNMi Everest out-of-the-box permissions for the `/var/opt/OV/log/nnm/public` folder is 755, NNMi uses ACLs to adjust access permissions for the database user (`nmsdbmgr`), and the `nnmaction` user (`bin`). During the NNMi Everest post-installation (part of the installation or upgrade script), the installation script changes the `/var/opt/OV/log/nnm/public` folder permissions and adds the ACLs.

If the installation script is unable to set the ACLs on the `/var/opt/OV/log/nnm/public` folder due to some unexpected error, the script will leave the `/var/opt/OV/log/nnm/public` folder world-writable and the NNMi installation should complete successfully. Following a successful NNMi installation, if you want to restrict world-write permissions on the `/var/opt/OV/log/nnm/public` folder, see the system administration documentation for setting up ACLs for the NNMi management server's operating system.

For the `/var/opt/OV/log/nnm/public` folder, use UNIX ACLs (access control lists) to adjust user access. Configuring ACLs is a useful method to extend the `owner/group/other` permissions. ACLs are supported on all 4 UNIX platforms: RedHat, SuSE, HP-UX, and Solaris.

For example, after running the following commands, the user depicted by the `USER` variable obtains write access to the `/var/opt/OV/log/nnm/public` folder. Without running these commands, the permissions for the `/var/opt/OV/log/nnm/public` folder are 755, and files within the directory are not writable by anyone other than root.

RedHat Linux and SuSE Linux and Solaris:

```
setfacl -m user:<USER>:rwx /var/opt/OV/log/nnm/public
```

HP-UX:

```
setacl -m user:<USER>:rwx /var/opt/OV/log/nnm/public
```

Solaris ZFS:

```
chmod A+user:<USER>:read_data/add_file/write_data/  
list_directory:allow /var/opt/OV/log/nnm/public
```

For information about how to use the `setfacl`, `setacl`, or `chmod` commands, see the appropriate reference pages, or the UNIX manpages.

Administering a Custom Poller Collection Export

The **Custom Poller feature** enables you to take a proactive approach to network management by using SNMP MIB expressions to specify additional information that NNMI should poll. A **Custom Poller collection** defines the information you want to gather (poll) as well as how NNMI reacts to the gathered data. See *Create a Custom Poller Collection* and *Configure Custom Polling* in the NNMI help for more detailed information. See also the *HP Network Node Manager i Software Step-by-Step Guide to Custom Poller White Paper*.

The Custom Poller feature relies on you to remove files from the export directory as you process them. Do not use the exported files for long term storage; if they consume more than the configured maximum disk space, NNMI removes the older files and creates new ones. Unless you process these files and store them in a different location, you will lose them.

Changing the Custom Poller Collections Export Directory

NNMI writes the data from the collections you export into the following directory:

- *Windows*: %NNM_DATA%\shared\nnm\databases\custompoller\export
- *UNIX*: \$NNM_DATA/shared/nnm/databases/custompoller/export

To change the directory that NNMI writes its custom poller files into, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-custompoller.properties
 - *UNIX*: \$NNM_PROPS/nms-custompoller.properties
- 2 Look for the `exportdir` entry, which is similar to the following line:


```
#!com.hp.nnm.custompoller.exportdir=<base directory to export custom poller metrics>
```

To configure NNMI to write Custom Poller collection information into the C:\CustomPoller directory, change the line as follows:

```
com.hp.nnm.custompoller.exportdir=C:\CustomPoller
```
- 3 Restart the NNMI management server.
 - a Run the `ovstop` command on the NNMI management server.
 - b Run the `ovstart` command on the NNMI management server.

Changing the Maximum Amount of Disk Space for Custom Poller Collections Export

To change the maximum amount of disk space that NNMI uses when exporting data to `collection_name.csv` files, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-custompoller.properties
 - *UNIX*: \$NNM_PROPS/nms-custompoller.properties

- 2 Look for the `maxdiskspace` entry, which is similar to the following line:

```
#!com.hp.nnm.custompoller.maxdiskspace=1000
```

To configure NNMi to reserve up to 2000 MB (2 GB) of storage space for each `collection_name.csv` file, change the line as follows:

```
com.hp.nnm.custompoller.maxdiskspace=2000
```

- 3 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server.

Changing the Custom Poller Metric Accumulation Interval

NNMi sets the length of time, in minutes, that it accumulates Custom Poller Collection metrics before it writes data into a file.

To change the custom poller metric accumulation interval, follow these steps:

- 1 Edit the following file:

- *Windows:* %NNM_PROPS%\nms-custompoller.properties
- *UNIX:* \$NNM_PROPS/nms-custompoller.properties

- 2 Look for a line the resembles the following:

```
#!com.hp.nnm.custompoller.accumulationinterval=5
```

To configure NNMi to collect metrics for ten minutes instead of the default value of five minutes, change the line as follows:

```
com.hp.nnm.custompoller.accumulationinterval=10
```

- 3 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server.

Administering Incident Actions

You can configure actions to automatically run at any point in the incident lifecycle. For example, you might want to configure an action to occur when an incident of the type you are configuring is generated. See *Configure an Action for an Incident* in the NNMi Help for more information.

To adjust action parameters, follow the steps shown in the following sections.

- ▶ To avoid undesirable results (such as unintended memory growth, slower event action processing time), HP recommends that you do not change the default property values for event action processing.

Setting the Number of Simultaneous Actions

- ▶ Increasing the number of simultaneous actions on a Solaris NNMi management server causes NNMi performance degradation.

To modify the number of simultaneous actions that NNMi can run, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\shared\nnmaction.properties
 - *UNIX*: \$NNM_PROPS/shared/nnmaction.properties
- 2 Look for a line that resembles the following:

```
#!com.hp.ov.nms.events.action.numProcess=10
```

To configure NNMi to enable 20 simultaneous actions instead of the default value, change the line as follows:

```
com.hp.ov.nms.events.action.numProcess=20
```

- ▶ Make sure to remove the **#!** characters located at the beginning of the line.
- 3 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.

Setting the Number of Threads for Jython Actions

To modify the number of threads the action server uses to run jython scripts, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\shared\nnmaction.properties
 - *UNIX*: \$NNM_PROPS/shared/nnmaction.properties

- 2 Look for a line that resembles the following:

```
#!com.hp.ov.nms.events.action.numJythonThreads=10
```

To configure NNMi to enable 20 threads for running jython scripts instead of the default value, change the line as follows:

```
com.hp.ov.nms.events.action.numJythonThreads=20
```



Make sure to remove the **#!** characters located at the beginning of the line.

- 3 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.

Setting the Action Server Name Parameter

If you have an NNMi management server running on a Windows operating system, the HP NNM Action Server runs as a windows service with a Local System account. That means you must use the Local System account to run action server actions.

To modify the user name that runs the HP NNM Action Server windows service on a Windows NNMi management server, change the LogOn property of the HP NNM Action Server service.

If you have an NNMi management server running on a HP-UX, Solaris, or Linux operating system, the action server runs with a bin user name. To modify the user name that runs the action server on these operating systems, complete the following steps:

- 1 Edit the following file:

```
$NNM_PROPS/nmaction.properties
```

- 2 Look for a line the resembles the following:

```
#!com.hp.ov.nms.events.action.userName=bin
```

To configure NNMi to have *root* run the action server instead of the default value, change the line as follows:

```
com.hp.ov.nms.events.action.userName=root
```



Make sure to remove the **#!** characters located at the beginning of the line.

- 3 Save your changes.
- 4 Restart the action server:
 - a Run the **ovstop nmaction** command on the NNMi management server.
 - b Run the **ovstart nmaction** command on the NNMi management server.

Changing the Action Server Queue Size

For actions that use a long action command string at a high execution rate, such as responding to a trap storm, the action server can use up a lot of memory. To provide better action server performance, HP places limits on the memory size that the action server can grow to.



For Solaris NNMi management servers, if the NNMi health information shows that action queue sizes are growing, reduce the maximum memory size to improve performance.

To modify these limits, follow these steps:

- 1 Edit the following file:
 - %NNM_PROPS%\shared\nnmaction.properties
 - \$NNM_PROPS/shared/nnmaction.properties
- 2 Look for two lines that resemble the following:


```
com.hp.ov.nms.events.action.jvmargs.minMemsize=-Xms6m
com.hp.ov.nms.events.action.jvmargs.maxMemsize=-Xmx30m
```
- 3 The above parameters show the minimum memory size set to 6MB and the maximum set to 30MB. Adjust these parameters to meet your needs.
- 4 Save your changes.
- 5 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.

Incident Actions Log

When an action runs, output is logged to the associated Incident Actions Log file. To view the contents of the log for a selected incident, use the **Tools > Incident Actions Log** menu option. Below are descriptions of the items contained in the log:

Table 40 Incident Actions Log Items

Item	Description
Command	Script to run when incident occurs
Incident Name	Name of incident as defined in incident configuration
Incident UUID	The UUID of the incident (from Registration tab)
Command Type	Type of command (Jython or ScriptOrExecutable)
Lifecycle State	Lifecycle state of the incident (Registered , In Process , Completed , or Closed)
Exit Code	Return code of the command (similar to an error code)
Standard Output	Standard output of the action
Standard Error	Standard error output
Execution Status	The determined status per the action

Blocking Trap Storms using the `hosted-on-trapstorm.conf` File

NNMi includes a way to block trap storms from hosted-on devices (including interfaces), while eliminating the dependency on iSPI-Net licensing for trap storm detection/suppression.

- 1 Run the `nmtrapconfig.ovpl` script. Include the appropriate values for `-hostedOnTrapstorm` and `-hostedOnThreshold`, as described in the `nmtrapconfig.ovpl` reference page or the UNIX manpage, to configure the trap service. Use the `-setProp` parameter to reconfigure the trap server to reflect the property changes.
- 2 Optionally, to change any out-of-the-box configurations, edit the following file:
 - *Windows*: `%NnmDataDir%\shared\nnm\conf\hosted-on-trapstorm.conf`
 - *UNIX*: `$NnmDataDir/shared/nnm/conf/hosted-on-trapstorm.conf`

Make changes per the format described in the `hosted-on-trapstorm.conf` reference page, or the UNIX manpage.
- 3 If you made changes to the `hosted-on-trapstorm.conf` file, you must run `nmtrapconfig.ovpl -stop` followed by `nmtrapconfig.ovpl -start` to restart the trap service. See the `nmtrapconfig.ovpl` reference page or the UNIX manpage for more information.

Blocking Incidents using the `trapFilter.conf` File

Suppose the number of incidents flowing through your NNMi management server reaches a rate that causes NNMi to block newly arriving incidents.

When this happens, NNMi generates a `TrapStorm` incident, indicating that incidents are blocked. NNMi might also generate a major health message indicating that the incident rate is high and incidents are being blocked.

To remedy this, you might try to use the `nmtrapd.conf` file to block incidents from entering NNMi in an attempt to reduce the incident traffic. However, if you use the `nmtrapd.conf` file approach, NNMi still uses these incidents to calculate the trap rate and to write to the trap binary store. By using the `nmtrapd.conf` file approach, you only stop incidents from being created or stored in the database. See the `nmtrapd.conf` reference page, or the UNIX manpage for more information.

There is a better solution to this problem than using the `nmtrapd.conf` file. NNMi provides a filtering mechanism that blocks incidents earlier in the NNMi event pipeline, preventing these incidents from being analyzed for trap rate calculations or from being stored in the NNMi trap binary store. By adding device IP addresses or OIDs to the `trapFilter.conf` file, you can block these high-volume incidents and avoid incident volume problems. See the `trapFilter.conf` reference page, or the UNIX manpage for more information.

Configuring Character Set Encoding Settings for NNMi

Depending on the locale configured for your NNMi management server, you might need to configure the source encodings NNMi uses to interpret SNMP OCTETSTRING data. To do this, edit the `nms-jboss.properties` file as follows:

- 1 Edit the following file:
 - *Windows*: `%NNM_PROPS%\nms-jboss.properties`
 - *UNIX*: `$NNM_PROPS/nms-jboss.properties`
- 2 Search for the text block containing the following line:
`#!com.hp.nnm.sourceEncoding=UTF-8`
- 3 Uncomment and edit the following line to read as follows:
`com.hp.nnm.sourceEncoding=UTF-8`
- 4 Modify the UTF-8 property value shown in [step 3](#) using the instructions and examples shown in the `nms-jboss.properties` file.
- 5 Save your work.
- 6 Restart NNMi:
 - a **ovstop**
 - b **ovstart**

Configuring NNMi to Require Encryption for Remote Access

An administrator can disable HTTP and other unencrypted access from the network to NNMi.



Before configuring NNMi to permit only encrypted remote access, make sure Global Network Management, NNMi iSPIs, and other integrations support SSL. Configure them for SSL before configuring NNMi to permit only encrypted remote access.

To disable HTTP and other unencrypted access from the network to NNMi, edit the `server.properties` file as follows:

- 1 Edit the following file (you may need to create it if it does not exist):
 - *Windows*: `%NnmDataDir%\nmsas\NNM\server.properties`
 - *UNIX*: `$NnmDataDir/nmsas/NNM/server.properties`
- 2 Add the following four lines to the `server.properties` file:


```
nmsas.server.net.bind.address = 127.0.0.1
nmsas.server.net.bind.address.ssl = 0.0.0.0
nmsas.server.net.hostname = localhost
nmsas.server.net.hostname.ssl = ${com.hp.ov.nms.fqdn}
```
- 3 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.

- b Run the `ovstart` command on the NNMi management server.

With the modification just described, NNMi will not “listen” to HTTP requests from a remote system; however, HTTP requests would still be supported for localhost access.

Configuring the Auto-Trim Oldest SNMP Trap Incidents Feature

To keep NNMi performing at a high level, NNMi drops incoming SNMP traps (including syslog messages) after storing a specific number of SNMP traps in its database. You can use the auto-trim oldest SNMP trap incidents feature to control the number of SNMP traps stored in the NNMi database and to retain important incoming SNMP traps.

The auto-trim oldest SNMP trap incidents feature defaults to being disabled. After enabling the auto-trim oldest SNMP trap incidents feature, NNMi removes the oldest SNMP trap incidents from the NNMi database.



To manually trim SNMP trap incidents from the NNMi database, use the `nmtrimincidents.ovpl` script. See the `nmtrimincidents.ovpl` reference page, or the UNIX manpage, for more information.

Enabling the Auto-Trim Oldest SNMP Trap Incidents Feature (No Incident Archive)

Suppose you want to enable the auto-trim oldest SNMP trap incidents feature to trim 30,000 SNMP trap incidents (including syslog messages) after the number of SNMP trap incidents in the NNMi database exceeds 60,000. For this example, you do not want NNMi to archive the SNMP trap incidents before trimming them. Complete the following steps:

- 1 Edit the following file:
 - *Windows*: `%NNM_PROPS\nms-jboss.properties`
 - *UNIX*: `$NNM_PROPS/nms-jboss.properties`
- 2 Locate the text block containing the following line:
`#!com.hp.nnm.events.snmpTrapAutoTrimStartPercentage=50`
- 3 Uncomment and edit the line to read as follows:
`com.hp.nnm.events.snmpTrapAutoTrimStartPercentage=60`
- 4 Locate the text block containing the following line:
`#!com.hp.nnm.events.snmpTrapAutoTrimPercentageToDelete=25`
- 5 Uncomment and edit the line to read as follows:
`com.hp.nnm.events.snmpTrapAutoTrimPercentageToDelete=50`
- 6 Locate the text block containing the following line:
`#!com.hp.nnm.events.snmpTrapAutoTrimSetting=Disabled`
- 7 Uncomment and edit the line to read as follows:
`com.hp.nnm.events.snmpTrapsAutoTrimSetting=TrimOnly`
- 8 Restart NNMi:
 - a Run the `ovstop` command on the NNMi management server.

- b Run the **ovstart** command on the NNMi management server.

The default value of **com.hp.nnm.events.snmpTrapMaxStoreLimit** is 100,000. With this configuration, after NNMi stores 60,000 SNMP trap incidents (including syslog messages) from the NNMi database, it trims 30,000 SNMP trap incidents from the NNMi database using the following formula:

```
(com.hp.nnm.events.snmpTrapAutoTrimStartPercentage X
com.hp.nnm.events.snmpTrapMaxStoreLimit X
com.hp.nnm.events.snmpTrapAutoTrimPercentageToDelete
```

Enabling the Auto-Trim Oldest SNMP Trap Incidents Feature (Incident Archive Enabled)

Suppose you want to enable the auto-trim oldest SNMP trap incidents feature to trim 60,000 SNMP trap incidents (including syslog messages) after the number of SNMP trap incidents in the NNMi database exceeds 80,000. For this example, you want NNMi to archive the SNMP trap incidents before trimming them. Complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS\nms-jboss.properties
 - *UNIX*: \$NNM_PROPS/nms-jboss.properties
- 2 Locate the text block containing the following line:
#!com.hp.nnm.events.snmpTrapAutoTrimStartPercentage=50
- 3 Uncomment and edit the line to read as follows:
com.hp.nnm.events.snmpTrapAutoTrimStartPercentage=80
- 4 Locate the text block containing the following line:
#!com.hp.nnm.events.snmpTrapAutoTrimPercentageToDelete=25
- 5 Uncomment and edit the line to read as follows:
com.hp.nnm.events.snmpTrapAutoTrimPercentageToDelete=75
- 6 Locate the text block containing the following line:
#!com.hp.nnm.events.snmpTrapAutoTrimSetting=Disabled
- 7 Edit the line to read as follows:
com.hp.nnm.events.snmpTrapsAutoTrimSetting=TrimAndArchive
- 8 Restart NNMi:
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.

The default value of **com.hp.nnm.events.snmpTrapMaxStoreLimit** is 100,000. With this configuration, after NNMi stores 80,000 SNMP trap incidents (including syslog messages) from the NNMi database, it archives, then trims 60,000 SNMP trap incidents from the NNMi database using the following formula:

```
(com.hp.nnm.events.snmpTrapAutoTrimStartPercentage X
com.hp.nnm.events.snmpTrapMaxStoreLimit X
com.hp.nnm.events.snmpTrapAutoTrimPercentageToDelete
```

Reducing the Number of Stored SNMP Trap Incidents

If you do not need NNMi to keep SNMP trap incidents for a long time period, you might consider reducing the number of SNMP trap incidents stored in the NNMi database.



NNMi begins dropping SNMP traps (including syslog messages) after the number of SNMP trap incidents in its database reaches 100,000. Setting this limit to a higher number is not supported, as doing so can cause NNMi performance degradation.

Suppose you want to reduce the maximum number of stored SNMP trap incidents (including syslog messages) to 50,000 SNMP trap incidents. To do this, complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS\nms-jboss.properties
 - *UNIX*: \$NNM_PROPS/nms-jboss.properties
- 2 Locate the text block containing the following line:


```
#!com.hp.nnm.events.snmpTrapMaxStoreLimit=100000
```
- 3 Uncomment and edit the line to read as follows:


```
com.hp.nnm.events.snmpTrapMaxStoreLimit=50000
```
- 4 Restart NNMi:
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.

Monitoring the Auto-Trim Oldest SNMP Trap Incidents Feature

From the **NNMi console**, click **Help > System Information > Health** to check the health of the auto-trim oldest SNMP trap incidents feature. NNMi also generates the following alarms regarding the auto-trim oldest SNMP trap incidents feature.

- NNMi generates a critical alarm after the number of stored SNMP trap incidents (including syslog messages) reaches 100% of the **com.hp.nnm.events.snmpTrapMaxStoreLimit** value.
- NNMi generates an **snmpTrapLimitMajorAlarm** alarm after the number of stored SNMP trap incidents (including syslog messages) reaches 95% of the **com.hp.nnm.events.snmpTrapMaxStoreLimit** value.
- NNMi generates an **snmpTrapLimitWarningAlarm** alarm after the number of stored SNMP trap incidents (including syslog messages) reaches 90% of the **com.hp.nnm.events.snmpTrapMaxStoreLimit** value.

Disabling the Auto-Trim Oldest SNMP Trap Incidents Feature

To disable the auto-trim oldest incidents feature, complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS\nms-jboss.properties
 - *UNIX*: \$NNM_PROPS/nms-jboss.properties

- 2 Locate the text block containing the following:
`com.hp.nnm.events.snmpTrapAutoTrimSetting`
- 3 Uncomment and edit the line to read as follows:
`com.hp.nnm.events.snmpTrapAutoTrimSetting=Disabled`
- 4 Restart NNMi:
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server.

Modifying NNMi Gauge Titles to Show SNMP MIB Variable Names

The **Node Component** tab in the NNMi analysis pane contains gauges showing the NNMi component name as the MIB OID being polled. This helps you understand which gauge goes with which component. The node component name helps differentiate gauges if NNMi shows a lot of gauges for a node. For example, if a node contains a large number of CPUs, NNMi shows different names for the individual CPUs. With this feature disabled, NNMi shows the SNMP MIB variable name to be the same for all CPUs.

If you want to change this property to show gauge titles as SNMP MIB variable names rather than NNMi component names, complete the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS\nms-ui.properties
 - *UNIX*: \$NNM_PROPS/nms-ui.properties
- 2 Locate the text block containing the following line:
`com.hp.nnm.ui.analysisGaugeTitleIsNodeComponentName = true`
- 3 Edit the following line to read as follows:
`com.hp.nnm.ui.analysisGaugeTitleIsNodeComponentName = false`
- 4 Save your changes.
- 5 Restart NNMi:
 - a `ovstop`
 - b `ovstart`

Modifying NNMi Normalization Properties

NNMi stores both hostnames and node names in case-sensitive form. This means that all searches, sorts, and filters that the NNMi console provides return case-sensitive results. If the DNS servers you use return a variety of case-preserving node names and hostnames, including all uppercase, all lowercase, and a mixture of uppercase and lowercase, this can cause less-than-optimal results.

You can change several NNMi normalization properties to meet your specific needs. A good practice is to make these changes before seeding NNMi for its initial discovery. HP recommends that you adjust the settings in this section during deployment, but before running the initial discovery.

If you run an initial discovery, then decide to change the normalization properties later, you can run the `nnmnodediscover.ovpl -all` script to initiate a full discovery. See the `nnmnodediscover.ovpl` reference page, or the UNIX manpage, for more information.

You can change the following properties:

- Normalize discovered node names to UPPERCASE, LOWERCASE, or OFF.
- Normalize discovered hostnames to UPPERCASE, LOWERCASE, or OFF.

To change normalization properties follow these steps:

1 Edit the following file:

- *Windows:* %NNM_PROPS%\nms-topology.properties
- *UNIX:* \$NNM_PROPS/nms-topology.properties

2 To configure NNMi to normalize discovered names, look for a line the resembles the following:

```
#!com.hp.ov.nms.topo.NAME_NORMALIZATION=OFF
```

a Un-comment the property:

```
com.hp.ov.nms.topo.NAME_NORMALIZATION=OFF
```



To un-comment a property, remove the `#!` characters from the beginning of a line.

b Change OFF to LOWERCASE or UPPERCASE.

c Save your changes.

3 To configure NNMi to normalize discovered hostnames, look for a line the resembles the following:

```
#!com.hp.ov.nms.topo.NAME_NORMALIZATION=OFF
```

a Un-comment the property:

```
com.hp.ov.nms.topo.HOSTNAME_NORMALIZATION=OFF
```

b Change OFF to LOWERCASE or UPPERCASE.

c Save your changes.

4 Restart the NNMi management server.

a Run the `ovstop` command on the NNMi management server.

b Run the `ovstart` command on the NNMi management server.

Changing Normalization Properties Following an Initial Discovery

Changing normalization properties following an initial discovery causes NNMi to be inconsistent with the property changes until the next discovery. To remedy this, run the `nnmnodediscover.ovpl -all` script to initiate a full discovery after changing NNMi normalization properties.

After NNMi completes a full discovery, the behaviors shown below should return to normal. These examples are not exhaustive, and are intended to provide a few examples of things to consider when changing NNMi normalization properties.

Modifying Simultaneous SNMP Requests


NNMi maintains a limit of three simultaneous SNMP requests to a node. This reduces the risk of a node's SNMP agents dropping responses.

You can adjust this value higher, resulting in increased discovery speed. However, if you set the value too high, you increase the risk of dropped responses and reduced discovery accuracy.

If you want to modify this limit, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS%\nms-communication.properties
 - *UNIX*: \$NNM_PROPS/nms-communication.properties
- 2 To increase the current number of simultaneous SNMP requests to a node, do the following:
 - a Look for a line the resembles the following:

```
#!com.hp.ov.nms.comm.snmp.maxConcurrentRequests=3
```
 - b Un-comment the property:

```
com.hp.ov.nms.comm.snmp.maxConcurrentRequests=3
```
 -  To un-comment a property, remove the #! characters from the beginning of a line.
 - c Change the existing value to the number of desired simultaneous SNMP requests to a node.
 - d Save your changes.
- 3 Restart the NNMi management server.
 - a Run the **ovstop** command on the NNMi management server.
 - b Run the **ovstart** command on the NNMi management server.

Modifying the Embedded Database Port

If you want to configure NNMi to use a different port for the embedded database, follow these steps:

- 1 Edit the following file:
 - *Windows*: %NNM_CONF%\nnm\props\nms-local.properties
 - *UNIX*: \$NNM_CONF/nnm/props/nms-local.properties
- 2 Look for a line the resembles the following:

```
#!com.hp.ov.nms.postgres.port=5432
```

- 3 Un-comment the property:
`com.hp.ov.nms.postgres.port=5432`



To un-comment a property, remove the `#!` characters from the beginning of a line.

- 4 Change the existing value to the new port number.
- 5 Save your changes.
- 6 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server

Modifying MIB Browser Parameters

If you use the NNMi MIB browser (**Action** > **MIB Information** > **Browse MIB** menu) to obtain information about a node, and provide an optional SNMP community string for that node, the NNMi MIB browser uses MIB browser parameters located in the `nms-ui.properties` file for MIB Browser SNMP communication.



If you do not provide a community string when using the MIB Browser, NNMi uses the **Communication Configuration** settings established for the node (if any). These settings are configured in the NNMi console using the **Communications Settings** view in the **Configuration** workspace. See *Configuring Communication Protocol* in the NNMi help for more information.

To modify the MIB Browser parameters in the `nms-ui.properties` file, follow these steps:

- 1 Edit the following file:
 - *Windows*: `%NNM_PROPS\nms-ui.properties`
 - *UNIX*: `$NNM_PROPS/nms-ui.properties`
- 2 Locate the text block containing the following line:


```
# MIB Browser Parameters
```
- 3 Locate the MIB browser parameters located below `# MIB Browser Parameters` by searching for lines containing the following text:


```
mibbrowser
```
- 4 Modify the MIB browser parameters by following instructions within the `nms-ui.properties` file.
- 5 Save your changes.
- 6 Restart NNMi:
 - a `ovstop`
 - b `ovstart`

NNMi Self Monitoring

NNMi performs self-monitoring checks, including memory, CPU, and disk resources. NNMi generates an incident after the NNMi management server becomes low on resources or detects a serious condition.

To view NNMi health information, use one of the following methods:

- From the NNMi console, click **View > System Information**; then click the **Health** tab.
- For a detailed self-monitoring report, select **Tools > NNMi System Health Report**
- Run the **nnmhealth.ovpl** script.

NNMi displays a status message at the bottom of the NNMi console and on the top of forms after an NNMi detects a self-monitoring health exception. You can disable this warning message by completing the following steps:

- 1 Edit the following file:
 - *Windows*: %NNM_PROPS\nms-ui.properties
 - *UNIX*: \$NNM_PROPS/nms-ui.properties
- 2 Locate the text block containing the following line:
`#!com.hp.nms.ui.health.disablewarning=false`
- 3 Uncomment and edit the following line to read as follows:
`com.hp.nms.ui.health.disablewarning==true`
- 4 Save your changes.
- 5 Restart NNMi:
 - a **ovstop**
 - b **ovstart**

Suppressing the Use of Discovery Protocols for Specific Nodes

NNMi uses several protocols to discover layer 2 connectivity between and among network devices. There are many defined discovery protocols. For example, *Link Layer Discovery Protocol* (LLDP) is an industry standard protocol, while there are many vendor-specific protocols like *Cisco Discovery Protocol* (CDP) for Cisco devices.

You can configure NNMi to suppress discovery protocol collections for devices you specify. There are special circumstances that might be remedied by suppressing discovery protocol collections.

Here are some examples:

- *Enterasys devices*: Using SNMP to collect information from the *Enterasys Discovery Protocol* (EnDP) and LLDP tables on some Enterasys devices might cause issues with NNMi running out of memory. You could prevent this by

configuring NNMi to skip EnDP and LLDP processing on these devices. To do this, add the management address of the devices to the `disco.SkipXdpProcessing` file as shown in [Suppressing the Use of Discovery Protocol Collections](#).



New operating system versions on some Enterasys devices support the `set snmp timefilter break` command. On those Enterasys devices, run the `set snmp timefilter break` command. If you configure the device using this command, you do not need to list the device in the `disco.SkipXdpProcessing` file.

- *Nortel devices:* Many Nortel devices use *SynOptics Network Management Protocol* (SONMP) to discover layer 2 layout and connectivity. Some of these devices use the same MAC address on multiple interfaces, and do not work well with this protocol. You might experience this problem if two interconnected Nortel devices show a layer 2 connection between the wrong set of interfaces and the connection shows a connection source of SONMP. For this example, it is best to configure NNMi to not use the SONMP protocol to derive layer 2 connections for the devices shown as participating in the wrong connection. To do this, add the management address of the two devices to the `disco.SkipXdpProcessing` file as shown in [Suppressing the Use of Discovery Protocol Collections](#).

Suppressing the Use of Discovery Protocol Collections

If you want to suppress this collection, follow these steps:

- 1 Create the following file:
 - *Windows:*
`%NnmDataDir%\shared\nnm\conf\disco\disco.SkipXdpProcessing`
 - *UNIX:* `$NnmDataDir/shared/nnm/conf/disco/disco.SkipXdpProcessing`

The `disco.SkipXdpProcessing` file is case-sensitive.
- 2 Add the device IP addresses to the `disco.SkipXdpProcessing` file for all of the devices you want to suppress protocol collection for. Follow the instructions show in the *disco.SkipXdpProcessing* reference page, or the UNIX manpage.
- 3 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server.



Suppressing the discovery protocol processing of a node or nodes might cause some inaccuracies in the layer 2 layout of the managed network. HP is not responsible for these inaccuracies.



The `ovjboss` service reads the `disco.SkipXdpProcessing` file on startup. If you make any changes after starting NNMi, restart NNMi as shown in this step.



If you ran the `set snmp timefilter break` command on any Enterasys devices, remove the device addresses from the `disco.SkipXdpProcessing` file, then restart NNMi as shown in this step. NNMi displays more accurate layer 2 maps when it uses discovery protocols.

See the *disco.SkipXdpProcessing* reference page, or the UNIX manpage, for more information.

Suppressing the Use of VLAN-indexing for Large Switches

One of the methods NNMi uses to learn layer 2 connectivity between and among switch devices in a managed network is to retrieve the `dot1dTpFdbTable` (FDB) from the switches. However, for Cisco switches, NNMi must use a VLAN-indexing method to retrieve the entire FDB. If there is a large number of VLANs configured on each device, retrieving the FDB with VLAN-indexing might take hours to complete.

Cisco switches are often configured to use the Cisco Discovery Protocol (CDP). CDP is considered to be a superior method for learning layer 2 connectivity. Large switches located in the in the core of the network might contain many VLANs. These switches typically do not have end nodes connected directly to them. If the switches you want to manage do not have end nodes connected directly to them, you might want to suppress the collection of the FDB on these large switches. NNMi still completes the layer 2 discovery using data collected from CDP. These large switches are prime candidates for suppression of VLAN-indexing. Do not suppress VLAN-indexing on smaller switches located at the network's edge (often known as access switches) that have many end nodes attached to them.

You can configure NNMi to suppress VLAN-indexing. To do this, the NNMi administrator needs to create and add management addresses or address ranges of the large switches to the `disco.NoVLANIndexing` file as shown in [Suppressing the Use of VLAN-indexing](#) on page 411. The `ovjboss` service reads the `disco.NoVLANIndexing` file when it starts. If the NNMi administrator makes changes to the `disco.NoVLANIndexing` file after the `ovjboss` service starts, those changes will not take effect until the next time the `ovjboss` service starts. By default, the `disco.NoVLANIndexing` file does not exist. If the `disco.NoVLANIndexing` does not exist, this feature is disabled and NNMi attempts to use VLAN-indexing to collect the entire FDB table on all devices.

Suppressing the Use of VLAN-indexing

If you want to disable this `vlan-indexing`, follow these steps:

- 1 Create the following file:
 - *Windows*: `%NnmDataDir%\shared\nnm\conf\disco\disco.NoVLANIndexing`
 - *UNIX*: `$NnmDataDir/shared/nnm/conf/disco/disco.NoVLANIndexing`

The `disco.NoVLANIndexing` file is case-sensitive.
- 2 Add the device IP addresses or address ranges to the `disco.NoVLANIndexing` file for all of the devices you want to disable `vlan-indexing` for. Follow the instructions show in the *disco.NoVLANIndexing* reference page, or the UNIX manpage.
- 3 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server.



Suppressing `vlan-indexing` of a node or nodes might cause some inaccuracies in the layer 2 layout of the managed network. HP is not responsible for these inaccuracies.



The `ovjboss` service reads the `disco.NoVLANIndexing` file on startup. If you make any changes after starting NNMi, restart NNMi as shown in this step.

See the `disco.Disco.NoVLANIndexing` reference page, or the UNIX manpage, for more information.

Configuring Node Component Status

NNMi includes the following node components, which might be monitored to help determine a node's status:

- 1 CPU
- 2 BUFFERS
- 3 VOLTAGE
- 4 TEMPERATURE
- 5 DISK_SPACE
- 6 FAN
- 7 POWER_SUPPLY
- 8 BACK_PLANE
- 9 MEMORY

By default, the last four node components in the previous list, FAN, POWER_SUPPLY, BACK_PLANE, and MEMORY, propagate their status to the node level. For example, if a fan has a red status indicator, then its corresponding node receives a status indicator of yellow. A user, in this case, viewing the status of a node would be alerted to the fact that a component of that node has some kind of failure.

By default, the first five node components in the previous list do not propagate status to the node level.

Propagating Node Component Status to a Node

You can configure a node component to have its status propagate to the node level by performing the following steps:

- 1 If not already present, create a new properties file with the name `nmm-apa.properties` in the following directory:
 - *Windows*: `%NnmDataDir%\shared\nnm\conf\props`
 - *UNIX*: `$NnmDataDir/shared/nnm/conf/props`
- 2 Within the properties file, use a text editor to include the following text:


```
com.hp.ov.nms.apa.NodeComponentPropagateToNodeStatus_<Type>:
true
```

where `<Type>` = a node component as shown in the list at the start of [Configuring Node Component Status](#) on page 412.

- 3 Save the properties file.
- 4 Run the following command sequence on the NNMi management server:

- a ovstop
- b ovstart

Configuring a Node Component's Status so that it does not Propagate to the Node

You can configure a node component such that it does not propagate its status to the node level by performing the following steps:

- 1 If not already present, create a new properties file with the name `nmm-apa.properties` in the following directory:
 - *Windows*: `%NnmDataDir%\shared\nnm\conf\props`
 - *UNIX*: `$NnmDataDir/shared/nnm/conf/props`
- 2 Within the properties file, use a text editor to include the following text:

```
com.hp.ov.nms.apa.NodeComponentNoPropagateToNodeStatus_<Type>:
true
```

where `<Type>` = a node component as shown in the list at the start of [Configuring Node Component Status](#) on page 412.
- 3 Save the properties file.
- 4 Run the following command sequence on the NNMi management server:
 - a ovstop
 - b ovstart

Overriding Node Component Status Values

By default, three node component state values (None, Warning, and Unavailable) map up to a Normal status by the Causal Engine. You can override these default state mappings such that None, Warning, and Unavailable map to Critical with the following method:

- `com.hp.ov.nms.apa.NodeComponentValueRemappedToDown_None`
- `com.hp.ov.nms.apa.NodeComponentValueRemappedToDown_Warning`
- `com.hp.ov.nms.apa.NodeComponentValueRemappedToDown_Unavailable`

To override node component status values:

- 1 If not already present, create a new properties file with the name `nmm-apa.properties` in the following directory:
 - *Windows*: `%NnmDataDir%\shared\nnm\conf\props`
 - *UNIX*: `$NnmDataDir/shared/nnm/conf/props`
- 2 Within the properties file, use a text editor to include one, two, or all three of the following lines, as applicable:


```
com.hp.ov.nms.apa.NodeComponentValueRemappedToDown_None: true
com.hp.ov.nms.apa.NodeComponentValueRemappedToDown_Warning: true
com.hp.ov.nms.apa.NodeComponentValueRemappedToDown_Unavailable:
true
```
- 3 Save the properties file.
- 4 Run the following command sequence on the NNMi management server:
 - a `ovstop`
 - b `ovstart`



You can map an Unavailable state to an Unpolled status (since Unavailable means that the measurement facility is not available). This situation can often occur because the sensor is non-functional as opposed to the component being non-functional. To map Unavailable to Unpolled, use the same procedure as just described, except in [step 2](#), use the following text:

```
com.hp.ov.nms.apa.NodeComponentValueRemappedToUnpolled_Unavailable:
true
```

NNMi Logging

NNMi Log Files

To investigate HP Network Node Manager i Software (NNMi) performance, or to observe how NNMi processes and services are behaving, you can view log files that show a history of process and service activity. These files are available at the following location:

- *Windows*: %NnmDataDir%\log\nnm\
- *UNIX*: \$NnmDataDir/log/nnm

NNMi stores these log files in a *name.log* file name format. Any archived log file has a number appended to it in the form *name.log.%g*.

- *name* is the log file base name.
- *%g* relates to the archive number of the archived log file. The highest appended archive number represents the oldest file.

A log file can become an archived log file after the size of the log file exceeds the configured limit. After a log file exceeds the configured limit, the last active log file is archived. For example, after NNMi archives the *nmn.log* file as the *nmn.log.1* file, NNMi begins logging to a new *nmn.log* file.

NNMi logs messages at the following logging levels:

- **SEVERE**: Events that relate to abnormal NNMi behavior.
- **WARNING**: Events that indicate potential problems and all messages included in the SEVERE logging level.
- **INFO**: Messages written to the NNMi console (or its equivalent) and all messages included in the WARNING logging level.

Changing Logging File Properties

NNMi includes some features that can change NNMi logging. The instructions included in this section explain how to adjust these features.

Sign-in and Sign-out Logging

NNMi 9.20 is not configured to generate a log entry for each user that signs in to or out of the NNMi console. If you want to configure NNMi to log sign-in and sign-out activity, do the following:

- 1 Edit the following file:
 - Windows: %NNM_DATA%\shared\nnm\conf\props\nnm-logging.properties
 - UNIX: \$NNM_DATA/shared/nnm/conf/props/nnm-logging.properties
- 2 Search for the text block containing the following line:

```
com.hp.ov.nnm.log.signin.level = OFF
```
- 3 Modify the line to read as follows:

```
com.hp.ov.nnm.log.signin.level = INFO
```
- 4 Save your changes.
- 5 Restart NNMi:
 - a Run **ovstop**
 - b Run **ovstart**

NNMi Security

This chapter contains the following topics:

- [Configuring SSL Communications for Web Access and RMI Communications](#)
- [Allowing Non-Root UNIX Users to Start and Stop NNMi](#)
- [Providing a Password for Embedded Database Tools](#)

Configuring SSL Communications for Web Access and RMI Communications

NNMi includes a suite of default ciphers that are used in configuring Secure Sockets Layer (SSL) in Web access and Java Remote Method Invocation (RMI) communications. The ciphers are listed in the `nms-jboss.properties` file.



Adding or removing ciphers from the cipher list without the approval of HP is not supported; doing so may cause damage to the product or cause the product to become inoperable.

Allowing Non-Root UNIX Users to Start and Stop NNMi



If the `/opt/OV` directory is on a partition with the `nosuid` option set, the non-root user feature is not available. See `/etc/fstab` to determine if the partition is configured with the `nosuid` option set.

NNMi provides a way to allow non-root UNIX users to start and stop NNMi. Do the following:

- 1 As root, edit the following file:
`$NnmDataDir/shared/nmm/conf/ovstart.allow`
- 2 Include the non-root users (one per line) that you want to be able to start and stop NNMi.
- 3 Save your changes.

Providing a Password for Embedded Database Tools

To run embedded database tools (such as `psql`), NNMi requires a password. NNMi provides a default password, which the user should change using the `nnmchangeembdbpw.ovpl` script.

You must be logged in as administrator on Windows systems or root on UNIX systems to run the `nnmchangeembdbpw.ovpl` script. For more information, see the `nnmchangeembdbpw.ovpl` reference page, or the UNIX manpage

In an HA environment you run the `nnmchangeembdbpw.ovp1` script on the Primary Cluster Node only. The application automatically copies the password to the Secondary Cluster Node; no further action is required.

Additional Information

This section contains the following appendices:

- [Manually Configuring NNMi for Application Failover](#)
- [NNMi Environment Variables](#)
- [NNMi 9.20 and Well-Known Ports](#)
- [NNMi 9.20 iSPI Well-Known Ports](#)
- [Suggested Configuration Changes](#)

Manually Configuring NNMi for Application Failover

The steps contained in this appendix provide an alternative to using the NNMi Cluster Setup Wizard to configure application failover.

- ▶ If you are using application failover with Oracle as your database, you must follow the configuration steps in this appendix.

To manually configure application failover, perform the following steps:

- 1 Run **ovstop** on both nodes.
- 2 Configure server X (active) and server Y (standby) for the application failover feature using guidance from the detailed instructions contained in the `nms-cluster.properties` file. Use the following procedure:

- ▶ **Edit** in the following steps means to uncomment the lines in the text block within the file and to modify the text.

- a Edit the following file:

- *Windows*: %NnmDataDir%\shared\nnm\conf\props\nms-cluster.properties
- *UNIX*: \$NnmDataDir/shared/nnm/conf/props/nms-cluster.properties

- b Declare a unique name for the NNMi cluster. Use the same name when configuring both the active and standby servers.

```
com.hp.ov.nms.cluster.name=MyCluster
```

- c Add the hostnames of all nodes in the cluster to the `com.hp.ov.nms.cluster.member.hostnames` parameter in the `nms-cluster.properties` file:

```
com.hp.ov.nms.cluster.member.hostnames = fqdn_for_active,  
fqdn_for_standby
```

- ▶ In NNMi 9.0x, the application failover feature supported a UDP solution where cluster hosts were automatically discovered on the network. Beginning with NNMi 9.1x, HP eliminated the UDP solution and only supports the TCP solution. If you are migrating from NNMi 9.0x you must define the cluster hostnames by completing [step c](#) for application failover to work.

- d *Optional*. Define other `com.hp.ov.nms.cluster*` parameters within the `nms-cluster.properties` file. Follow the instructions contained within the `nms-cluster.properties` file for modifying each parameter

- ▶ If you are using application failover with Oracle as your database, NNMi ignores the database parameters contained in the `nms-cluster.properties` file.

- 3 Depending on the approach you take, complete the instructions shown in [Configuring Application Failover to use Self-Signed Certificates](#) on page 130 or the instructions shown in [Configuring Application Failover to use a Certificate Authority](#) on page 132.



When configuring the application failover feature, you must merge the `nmn.keystore` and `nmn.truststore` file content for both nodes into a single `nmn.keystore` and `nmn.truststore` file. *You must choose your approach and complete one set of instructions from [step 3](#)*

- 4 Copy the following file from server X to server Y:
 - **Windows:**
`%NnmDataDir%\shared\nnm\conf\nnmcluster\cluster.keystore`
 - **UNIX:**
`$NnmDataDir/shared/nnm/conf/nnmcluster/cluster.keystore`
- 5 Run the following command on both server X and server Y: **nmncluster**
 Each server should display something similar to the following:

```
===== Current cluster state =====
State ID: 000000001000000005
Date/Time: 15 Mar 2011 - 09:37:58 (GMT-0600)
Cluster name: ThisCluster (key CRC:626,187,650)
Automatic failover: Enabled
NNM database type: Embedded
NNM configured ACTIVE node is: NO_ACTIVE
NNM current ACTIVE node is: NO_ACTIVE
Cluster members are:

  Local?   NodeType  State           OvStatus      Hostname/Address
  -----  -
* REMOTE   ADMIN     n/a             n/a           serverX.xxx.yyy.yourcompany.com/
16.78.61.68:7800
  (SELF)   ADMIN     n/a             n/a           serverY.xxx.yyy.yourcompany.com/
16.78.61.71:7800
```

The display should list both server X and server Y. If information about both nodes are not displayed, the nodes are not communicating with each other. Here are some things to check for and correct before continuing:

- The Cluster names might be different on server X and server Y.
- The key CRCs might be different on server X and server Y. Check the contents of the following files on both server X and server Y:
 - Windows:**
`%NnmDataDir%\shared\nnm\conf\nnmcluster\cluster.keystore`
 - UNIX:** `$NnmDataDir/shared/nnm/conf/nnmcluster/cluster.keystore`
- A firewall on server X or server Y might be preventing the nodes from communicating.
- Make sure you merged the `nmn.keystore` and `nmn.truststore` files. You should see this error displayed after running the **nmncluster** command.
- Server X and server Y are running different operating systems. For example, suppose server X is running a Linux operating system and server Y is running a Windows operating system. You should see this error displayed after running the **nmncluster** command.

- Server X and server Y are running different NNMi versions. For example, suppose server X is running NNMi 9.20 and server Y is running NNMi 9.20 patch 1 (after it is available). You should see this error displayed after running the **nnmcluster** command.

6 On server X, start the NNMi cluster manager:

```
nnmcluster -daemon
```



After you run the **nnmcluster -daemon** command on NNMi management server X, the NNMi cluster manager goes through the following startup routine:

- Connects NNMi management server X to the cluster.
- Detects that there are no other NNMi management servers present.
- NNMi management server X assumes the active state.
- Starts the NNMi services on NNMi management server X (the active server).
- Creates a database backup.

For more information, see the *nnmcluster* reference page, or the UNIX manpage.

7 Wait a few minutes for server X to become the first active node in the cluster. Run the **nnmcluster -display** command on server X and search the displayed results for the term ACTIVE as in ACTIVE_NNM_STARTING or ACTIVE_SomeOtherState. Do not continue with [step 8](#) until you know that server X is the active node.

8 On server Y, start the NNMi cluster manager:

```
nnmcluster -daemon
```



After you run the **nnmcluster -daemon** command on NNMi management server Y, the NNMi cluster manager goes through the following startup routine:

- Connects NNMi management server Y to the cluster.
- Detects that NNMi management server X is present and is in the active state. The display shows STANDBY_INITIALIZING.
- Compares the database backup on NNMi management server Y to the backup on NNMi management server X. If these do not match, a new database backup is sent from NNMi management server X (active) to NNMi management server Y (standby). The display shows STANDBY_RECV_DBZIP.
- NNMi management server Y receives a minimal set of transaction logs which is the minimum necessary for the backup to be applicable for its standby state. The display shows STANDBY_RECV_TXLOGS.
- NNMi management server Y goes into a waiting state, continuously receiving new transaction logs and heartbeat signals from NNMi management server X. The display shows STANDBY_READY.

For more information, see the *nnmcluster* reference page, or the UNIX manpage.

- 9 If a failover occurs, the NNMi console for server X no longer functions. Close the NNMi console session for server X and log on to server Y (the new active server). Instruct NNMi users to store two bookmarks in their browsers, one to server X (the active NNMi management server) and one to server Y (the standby NNMi management server). If a failover occurs, users can connect to server Y (the standby NNMi management server).
- 10 Instruct network operations center (NOC) personnel to configure their devices to send traps to both server X and server Y. While server X (active) is running, it processes the forwarded traps and server Y (standby) ignores the forwarded traps.

NNMi Environment Variables

HP Network Node Manager i Software (NNMi) provides many environment variables that are available for your use in navigating the file system and writing scripts.

This appendix contains the following topics:

- [Environment Variables Used in This Document](#)
- [Other Available Environment Variables](#)

Environment Variables Used in This Document

This document primarily uses the following two NNMi environment variables to reference file and directory locations. This list shows the default values. Actual values depend on the selections that you made during NNMi installation.

- *Windows Server 2008:*

- %NnmInstallDir%: <drive>\Program Files (x86)\HP\HP BTO Software
- %NnmDataDir%: <drive>\ProgramData\HP\HP BTO Software



On Windows systems, the NNMi installation process creates these system environment variables, so they are always available to all users.

- *UNIX:*

- \$NnmInstallDir: /opt/OV
- \$NnmDataDir: /var/opt/OV



On UNIX systems, you must manually create these environment variables if you want to use them.

Additionally, this document references some of the NNMi environment variables that you can source as part of your user log-on configuration on the NNMi management server. These variables are of the form `NNM_*`. For information about this extended list of NNMi environment variables, see [Other Available Environment Variables](#) on page 425.

Other Available Environment Variables

NNMi administrators access some NNMi file locations regularly. NNMi provides a script that sets up many environment variables for navigating to commonly accessed locations.

To set up the extended list of NNMi environment variables, use a command similar to the following examples:

- **Windows:** "C:\Program Files (x86)\HP\HP BTO Software\bin\nnm.envvars.bat"
- **UNIX:** . /opt/OV/bin/nnm.envvars.sh

After you run the command for your operating system, you can use the NNMi environment variables shown in [Table 41](#) (Windows) or [Table 42](#) (UNIX) to get to commonly used NNMi file locations.

Table 41 Environment Variable Default Locations for the Windows Operating System

Variable	Windows (example)
%NNM_BIN%	C:\Program Files (x86)\HP\HP BTO Software\bin
%NNM_CONF%	C:\ProgramData\HP\HP BTO Software\conf
%NNM_DATA%	C:\ProgramData\HP\HP BTO Software\
%NNM_DB%	C:\ProgramData\HP\HP BTO Software\shared\nnm\databases
%NNM_JAVA%	C:\Program Files (x86)\HP\HP BTO Software\nonOV\jdk\nnm\bin\java.exe
%NNM_JAVA_DIR%	C:\Program Files (x86)\HP\HP BTO Software\java
%NNM_JAVA_PATH_SEP%	;
%NNM_JBOSS%	C:\Program Files (x86)\HP\HP BTO Software\nmsas
%NNM_JBOSS_DEPLOY%	C:\Program Files (x86)\HP\HP BTO Software\nmsas\server\nms\deploy
%NNM_JBOSS_LOG%	C:\Program Files (x86)\HP\HP BTO Software\nmsas\server\nms\log
%NNM_JBOSS_SERVERCONF%	C:\Program Files (x86)\HP\HP BTO Software\nmsas\server\nms
%NNM_JRE%	C:\Program Files (x86)\HP\HP BTO Software\nonOV\jdk\nnm
%NNM_LOG%	C:\ProgramData\HP\HP BTO Software\log
%NNM_LRF%	C:\ProgramData\HP\HP BTO Software\shared\nnm\lrf
%NNM_PRIV_LOG%	C:\ProgramData\HP\HP BTO Software\log
%NNM_PROPS%	C:\ProgramData\HP\HP BTO Software\shared\nnm\conf\props
%NNM_SHARED_CONF%	C:\ProgramData\HP\HP BTO Software\shared\nnm\conf
%NNM_SHARE_LOG%	C:\ProgramData\HP\HP BTO Software\log
%NNM_SNMP_MIBS%	C:\Program Files (x86)\HP\HP BTO Software\misc\nnm\snmp_mibs
%NNM_SUPPORT%	C:\Program Files (x86)\HP\HP BTO Software\support
%NNM_TMP%	C:\ProgramData\HP\HP BTO Software\tmp
%NNM_USER_SNMP_MIBS%	C:\ProgramData\HP\HP BTO Software\shared\nnm\user-snmp-mibs
%NNM_WWW%	C:\ProgramData\HP\HP BTO Software\shared\nnm\www

Table 42 Environment Variable Default Locations for UNIX Operating Systems

Variable	HP-UX
\$NNM_BIN	/opt/OV/bin
\$NNM_CONF	/var/opt/OV/conf
\$NNM_DATA	/var/opt/OV
\$NNM_DB	/var/opt/OV/shared/nnm/databases
\$NNM_JAVA	/opt/OV/nonOV/jdk/nnm/bin/java
\$NNM_JAVA_DIR	/opt/OV/java
\$NNM_JAVA_PATH_SEP	:
\$NNM_JBOSS	/opt/OV/nmsas
\$NNM_JBOSS_DEPLOY	/opt/OV/nmsas/server/nms/deploy
\$NNM_JBOSS_LOG	/opt/OV/nmsas/server/nms/log
\$NNM_JBOSS_SERVERCONF	/opt/OV/nmsas/server/nms
\$NNM_JRE	/opt/OV/nonOV/jdk/nnm
\$NNM_LOG	/var/opt/OV/log
\$NNM_LRF	/var/opt/OV/shared/nnm/lrf
\$NNM_PRIV_LOG	/var/opt/OV/log
\$NNM_PROPS	/var/opt/OV/shared/nnm/conf/props
\$NNM_SHARED_CONF	/var/opt/OV/shared/nnm/conf
\$NNM_SHARE_LOG	/var/opt/OV/log
\$NNM_SNMP_MIBS	/opt/OV/misc/nnm/snmp_mibs
\$NNM_SUPPORT	/opt/OV/support
\$NNM_TMP	/var/opt/OV/tmp
\$NNM_USER_SNMP_MIBS	/var/opt/OV/shared/nnm/user-snmp-mibs
\$NNM_WWW	/var/opt/OV/shared/nnm/www

NNMi 9.20 and Well-Known Ports

Table 43 shows the ports NNMi uses on the management server. NNMi listens on these ports. If port conflicts occur, you can change most of these port numbers as shown in the *Change Configuration* column. See the *nmm.ports* reference page, or the UNIX manpage, for more information.

- ▶ For application failover to work successfully, open TCP ports 7800-7810. For the application failover feature to function correctly, the active and standby NNMi management servers must have unrestricted network access to each other.

Table 43 Ports Used on the NNMi Management Server

Port	Type	Name	Purpose	Change Configuration
80	TCP	nmsas.server.port.web.http	Default HTTP port - used for Web UI & Web Services - In GNM configurations NNMi uses this port to establish communication from the global manager to the regional manager - Once this port is open, it becomes bi-directional	Modify the <code>nms-local.properties</code> file You can also change this during installation
162	UDP	trapPort	SNMP trap port	Modify using the <code>nmtrapconfig.ovpl</code> Perl script. See the <i>nmtrapconfig.ovpl</i> reference page, or the UNIX manpage, for more information.
443	TCP	nmsas.server.port.web.https	Default secure HTTPS port (SSL) - used for Web UI & Web Services	Modify the <code>nms-local.properties</code> file

Table 43 Ports Used on the NNMi Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
1098	TCP	nmsas.server.port.naming.rmi	<ul style="list-style-type: none"> - Used by NNMi command line tools to communicate with a variety of services used by NNMi - HP recommends configuring the system firewall to restrict access to these ports to localhost only 	Modify the nms-local.properties file
1099	TCP	nmsas.server.port.naming.port	<ul style="list-style-type: none"> - Used by NNMi command line tools to communicate with a variety of services used by NNMi - HP recommends configuring the system firewall to restrict access to these ports to localhost only 	Modify the nms-local.properties file
3873	TCP	nmsas.server.port.remoting.ejb3	<ul style="list-style-type: none"> - Used by NNMi command line tools to communicate with a variety of services used by NNMi - HP recommends configuring the system firewall to restrict access to these ports to localhost only 	Modify the nms-local.properties file
4444	TCP	nmsas.server.port.jmx.jrmp	<ul style="list-style-type: none"> - Used by NNMi command line tools to communicate with a variety of services used by NNMi - HP recommends configuring the system firewall to restrict access to these ports to localhost only 	Modify the nms-local.properties file

Table 43 Ports Used on the NNMi Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
4445	TCP	nmsas.server.port.jmx.rmi	<ul style="list-style-type: none"> - Used by NNMi command line tools to communicate with a variety of services used by NNMi - HP recommends configuring the system firewall to restrict access to these ports to localhost only 	Modify the <code>nms-local.properties</code> file
4446	TCP	nmsas.server.port.invoker.unified	<ul style="list-style-type: none"> - Used by NNMi command line tools to communicate with a variety of services used by NNMi - HP recommends configuring the system firewall to restrict access to these ports to localhost only 	Modify the <code>nms-local.properties</code> file
4457	TCP	nmsas.server.port.hq	<ul style="list-style-type: none"> - Used for un-encrypted Global Network Management traffic. - Messaging travels from the global manager to the regional manager - Once this port is open, it becomes bi-directional 	Modify the <code>nms-local.properties</code> file

Table 43 Ports Used on the NNMi Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
4459	TCP	nmsas.server.port.hq.ssl	<ul style="list-style-type: none"> - Used for encrypted Global Network Management traffic. - Messaging travels from the global manager to the regional manager - Once this port is open, it becomes bi-directional 	Modify the <code>nms-local.properties</code> file
4712	TCP	nmsas.server.port.ts.recovery	Internal transaction service port	Modify the <code>nms-local.properties</code> file
4713	TCP	nmsas.server.port.ts.status	Internal transaction service port	Modify the <code>nms-local.properties</code> file
4714	TCP	nmsas.server.port.ts.id	Internal transaction service port	Modify the <code>nms-local.properties</code> file
5432	TCP	com.hp.ov.nms.postgres.port	This PostgreSQL port is the port the embedded database listens on for this NNMi management server.	Modify the <code>nms-local.properties</code> file
7800-7810	TCP		<ul style="list-style-type: none"> - JGroups ports for application failover - If application failover is not used, HP recommends configuring the system firewall to restrict access to these ports 	Modify the <code>nms-cluster.properties</code> file
8886	TCP	OVSPMD_MGMT	NNMi <code>ovspmd</code> (process manager) management port	Modify the <code>/etc/services</code> file
8887	TCP	OVSPMD_REQ	NNMi <code>ovsmpd</code> (process manager) request port	Modify the <code>/etc/services</code> file

Table 44 shows some of the ports NNMi uses to communicate with other systems. If a firewall separates NNMi from these systems, you must open many of these ports in the firewall. The actual set of ports depends on the set of integrations you configured to use with NNMi and how you configured those integrations. If column 4 indicates *Client*, NNMi connects or sends to this port; if column 4 indicates *Server*, NNMi listens on this port.

Table 44 Ports Used for Communication Between the NNMi Management Server and Other Systems

Port	Type	Purpose	Client, Server
80	TCP	Default HTTP port for NNMi; used for Web UI and Web Services	Server
80	TCP	Default HTTP port for NNMi connecting to other applications. The actual port depends on NNMi configuration.	Client
161	UDP	SNMP request port	Client
162	UDP	SNMP trap port - traps received by NNMi	Server
162	UDP	SNMP trap port; Trap Forwarding, Northbound Interface, or NetCool integrations	Client
389	TCP	Default LDAP port	Client
395	UDP	nGenius Probe SNMP trap port	Client
443	TCP	Default secure HTTPS port for NNMi connecting to other applications; the actual port depends on NNMi configuration. Default HTTPS port for HP OM on Windows	Client
443	TCP	Default secure HTTPS port; used for Web UI and Web Services	Server
636	TCP	Default secure LDAP port (SSL)	Client
1741	TCP	Default CiscoWorks LMS web services port	Client
4457	TCP	Used for un-encrypted Global Network Management traffic. The connection is from the global manager to the regional manager.	Client, Server
4459	TCP	Used for encrypted Global Network Management traffic. The connection is from the global manager to the regional manager.	Client, Server
7800-7810	TCP	JGroups ports for application failover	Client and Server
8004	TCP	Default HTTP port for NNMi if another web server already has port 80. Used for Web UI and Web Services. Verify the actual HTTP port for your NNMi management server.	Server
8080	TCP	Default HTTP port for connecting to NA if installed on the same system as NNMi. Default HTTPS port for HP UCMDB web services	Client
8443 or 8444	TCP	Default HTTP port for connecting to HP OM for UNIX	Client
9300	TCP	Default HTTP port for connecting to NNM iSPI Performance for Metrics	Client
50000	TCP	Default HTTPS port for connecting to SIM	Client

- ▶ If you configure NNMi to use ICMP fault polling or ping sweep for discovery, configure the firewall to pass ICMP packets through the firewall.
- ▶ The Web Services approach for the NNMi-HP OM integration does not work through a firewall, however the NNMi-HP OM integration using the Northbound Interface does work through a firewall.

If you plan to use the global network management feature, [Table 45](#) shows the well-know ports that need to be accessible from a global NNMi management server to a regional NNMi management server. The global network management feature requires these ports to be open for TCP access from the global NNMi management server to the regional NNMi management server. The regional NNMi management server will not open sockets back to the global NNMi management server.

Table 45 Required Accessible Sockets for Global Network Management

Security	Parameter	TCP Port
non-SSL	jboss.http.port	80
	jboss.bisocket.port	4457
SSL	jboss.https.port	443
	jboss.sslbisocket.port	4459

NNMi 9.20 iSPI Well-Known Ports

Table 46 shows the ports the HP Network Node Manager iSPI for MPLS Software uses on the management server. In case of port conflicts, almost all of these port numbers can be changed using the `server.properties` file located at: `%NnmDataDir%/nmsas/mpls/server.properties`.

Table 46 Ports Used on the HP Network Node Manager iSPI for MPLS Software Management Server

Port	Type	Name	Purpose	Change Configuration
5432	TCP	<code>com.hp.ov.nms.postgres.port</code>	This PostgreSQL port is the port the embedded database listens on for this NNMi management server. The port is expected to be the same as that configured for NNMi in the <code>nms-local.properties</code> file.	N/A
24040	TCP	<code>nmsas.server.port.web.http</code>	Default HTTP port - used for Web UI.	Modify the <code>server.properties</code> file. You can also change this during installation.
24041	TCP	<code>nmsas.server.port.remoting.ejb3</code>	Default EJB3 remoting connector port	Modify the <code>server.properties</code> file.
24043	TCP	<code>nmsas.server.port.web.https</code>	Default secure HTTPS port (SSL) - used for Web UI.	Modify the <code>server.properties</code> file. You can also change this during installation.
24044	TCP	<code>nmsas.server.port.jmx.jrmp</code>	Default RMI object port (JRMP invoker)	Modify the <code>server.properties</code> file.
24045	TCP	<code>nmsas.server.port.invoker.unified</code>	Default RMI remoting server connector port	Modify the <code>server.properties</code> file.
24046	TCP	<code>nmsas.server.port.naming.port</code>	Default bootstrap JNP service port (JNDI provider)	Modify the <code>server.properties</code> file. You can also change this during installation.
24047	TCP	<code>nmsas.server.port.hq</code>	Used for un-encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
24048	TCP	<code>nmsas.server.port.jmx.rmi</code>	Default RMI pooled invoker port	Modify the <code>server.properties</code> file.

Table 46 Ports Used on the HP Network Node Manager iSPI for MPLS Software Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
24049	TCP	nmsas.server.port.naming.rmi	Default port for RMI naming service	Modify the server.properties file.
24092	TCP	nmsas.server.port.hq.ssl	Used for encrypted Global Network Management traffic.	Modify the server.properties file.
24712	TCP	nmsas.server.port.ts.recovery	Default recovery port used by the Transaction service.	Modify the server.properties file.
24713	TCP	nmsas.server.port.ts.status	Default status port used by the Transaction service.	Modify the server.properties file.
24714	TCP	nmsas.server.port.ts.id	Default port used by the Transaction service.	Modify the server.properties file.

Table 47 shows the ports the NNM iSPI for IP Telephony uses on the management server. In case of port conflicts, almost all of these port numbers can be changed using the server.properties file located at: %NnmDataDir%/nmsas/ipt/server.properties.

Table 47 Ports Used on the NNM iSPI for IP Telephony Management Server

Port	Type	Name	Purpose	Change Configuration
5432	TCP	com.hp.ov.nms.postgres.port	This PostgreSQL port is the port the embedded database listens on for this NNMi management server. The port is expected to be the same as that configured for NNMi in the nms-local.properties file.	N/A
10080	TCP	nmsas.server.port.web.http	Default HTTP port - used for Web UI.	Modify the server.properties file. You can also change this during installation.
10083	TCP	nmsas.server.port.naming.rmi	Default port for RMI naming service	Modify the server.properties file
10084	TCP	nmsas.server.port.jmx.jrmp	Default RMI object port (JRMP invoker)	Modify the server.properties file.
10085	TCP	nmsas.server.port.jmx.rmi	Default RMI pooled invoker port	Modify the server.properties file.
10086	TCP	nmsas.server.port.invoker.unified	Default RMI remoting server connector port	Modify the server.properties file.

Table 47 Ports Used on the NNM iSPI for IP Telephony Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
10087	TCP	nmsas.server.port.hq	Used for un-encrypted Global Network Management traffic.	Modify the server.properties file.
10089	TCP	nmsas.server.port.remoting.ejb3	Default EJB3 remoting connector port	Modify the server.properties file.
10092	TCP	nmsas.server.port.hq.ssl	Used for encrypted Global Network Management traffic.	Modify the server.properties file.
10099	TCP	nmsas.server.port.naming.port	Default bootstrap JNP service port (JNDI provider)	Modify the server.properties file. You can also change this during installation.
10443	TCP	nmsas.server.port.web.https	Default secure HTTPS port (SSL) - used for Web UI.	Modify the server.properties file. You can also change this during installation.
14712	TCP	nmsas.server.port.ts.recovery	Default recovery port used by the Transaction service.	Modify the server.properties file.
14713	TCP	nmsas.server.port.ts.status	Default status port used by the Transaction service.	Modify the server.properties file.
14714	TCP	nmsas.server.port.ts.id	Default port used by the Transaction service.	Modify the server.properties file.

Table 48 shows the ports the NNM iSPI for IP Multicast uses on the management server. In case of port conflicts, almost all of these port numbers can be changed using the `server.properties` file located at: `%NmDataDir%/nmsas/multicast/server.properties`.

Table 48 Ports Used on the NNM iSPI for IP Multicast Management Server

Port	Type	Name	Purpose	Change Configuration
5432	TCP	<code>com.hp.ov.nms.postgres.port</code>	This PostgreSQL port is the port the embedded database listens on for this NNMi management server. The port is expected to be the same as that configured for NNMi in the <code>nms-local.properties</code> file.	N/A
8084	TCP	<code>nmsas.server.port.web.http</code>	Default HTTP port - used for Web UI.	Modify the <code>server.properties</code> file. You can also change this during installation.
14083	TCP	<code>nmsas.server.port.naming.rmi</code>	Default port for RMI naming service	Modify the <code>server.properties</code> file.
14084	TCP	<code>nmsas.server.port.jmx.jrmp</code>	Default RMI object port (JRMP invoker)	Modify the <code>server.properties</code> file.
14085	TCP	<code>nmsas.server.port.jmx.rmi</code>	Default RMI pooled invoker port	Modify the <code>server.properties</code> file.
14086	TCP	<code>nmsas.server.port.invoker.unified</code>	Default RMI remoting server connector port	Modify the <code>server.properties</code> file.
14087	TCP	<code>nmsas.server.port.hq</code>	Used for un-encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
14089	TCP	<code>nmsas.server.port.remoting.ejb3</code>	Default EJB3 remoting connector port	Modify the <code>server.properties</code> file.
14092	TCP	<code>nmsas.server.port.hq.ssl</code>	Used for encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
14099	TCP	<code>nmsas.server.port.naming.port</code>	Default bootstrap JNP service port (JNDI provider)	Modify the <code>server.properties</code> file. You can also change this during installation.
14102	TCP	<code>nmsas.server.port.ts.id</code>	Default port used by the Transaction service.	Modify the <code>server.properties</code> file.

Table 48 Ports Used on the NNM iSPI for IP Multicast Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
14103	TCP	nmsas.server.port.ts.recovery	Default recovery port used by the Transaction service.	Modify the server.properties file.
14104	TCP	nmsas.server.port.ts.status	Default status port used by the Transaction service.	Modify the server.properties file.
14443	TCP	nmsas.server.port.web.https	Default secure HTTPS port (SSL) - used for Web UI.	Modify the server.properties file. You can also change this during installation.

Table 49 shows the ports the NNM iSPI Performance for Traffic (Traffic Master component) uses on the management server. In case of port conflicts, almost all of these port numbers can be changed using the server.properties file located at: %NnmDataDir%/nmsas/traffic-master/server.properties.

Table 49 Ports Used on the NNM iSPI Performance for Traffic Management Server (Traffic Master)

Port	Type	Name	Purpose	Change Configuration
5432	TCP	com.hp.ov.nms.postgres.port	This PostgreSQL port is the port the embedded database listens on for this NNMi management server. The port is expected to be the same as that configured for NNMi in the nms-local.properties file	N/A
12080	TCP	nmsas.server.port.web.http	Default HTTP port - used for Web UI.	Modify the server.properties file. You can also change this during installation.
12081	TCP	nmsas.server.port.web.https	Default secure HTTPS port (SSL) - used for Web UI.	Modify the server.properties file. You can also change this during installation.
12083	TCP	nmsas.server.port.naming.rmi	Default port for RMI naming service	Modify the server.properties file.
12084	TCP	nmsas.server.port.jmx.jrmp	Default RMI object port (JRMP invoker)	Modify the server.properties file.
12085	TCP	nmsas.server.port.jmx.rmi	Default RMI pooled invoker port	Modify the server.properties file.

Table 49 Ports Used on the NNM iSPI Performance for Traffic Management Server (Traffic Master) (cont'd)

Port	Type	Name	Purpose	Change Configuration
12086	TCP	<code>nmsas.server.port.invoker.unified</code>	Default RMI remoting server connector port	Modify the <code>server.properties</code> file.
12087	TCP	<code>nmsas.server.port.hq</code>	Used for un-encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
12089	TCP	<code>nmsas.server.port.remoting.ejb3</code>	Default EJB3	Modify the <code>server.properties</code> file.
12092	TCP	<code>nmsas.server.port.hq.ssl</code>	Used for encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
12099	TCP	<code>nmsas.server.port.naming.port</code>	Default bootstrap JNP service port (JNDI provider)	Modify the <code>server.properties</code> file. You can also change this during installation.
12712	TCP	<code>nmsas.server.port.ts.recovery</code>	Default recovery port used by the Transaction service.	Modify the <code>server.properties</code> file.
12713	TCP	<code>nmsas.server.port.ts.status</code>	Default status port used by the Transaction service.	Modify the <code>server.properties</code> file.
12714	TCP	<code>nmsas.server.port.ts.id</code>	Default port used by the Transaction service.	Modify the <code>server.properties</code> file.

Table 50 shows the ports the NNM iSPI Performance for Traffic (Traffic Leaf component) uses on the management server. In case of port conflicts, almost all of these port numbers can be changed using the `server.properties` file located at: `%NnmDataDir%/nmsas/traffic-leaf/server.properties`.

Table 50 Ports Used on the NNM iSPI Performance for Traffic Management Server (Traffic Leaf)

Port	Type	Name	Purpose	Change Configuration
5432	TCP	<code>com.hp.ov.nms.postgres.port</code>	This PostgreSQL port is the port the embedded database listens on for this NNMi management server. The port is expected to be the same as that configured for NNMi in the <code>nms-local.properties</code> file.	N/A
11080	TCP	<code>nmsas.server.port.web.http</code>	Default HTTP port - used for Web UI.	Modify the <code>server.properties</code> file. You can also change this during installation.
11081	TCP	<code>nmsas.server.port.web.https</code>	Default secure HTTPS port (SSL) - used for Web UI.	Modify the <code>server.properties</code> file. You can also change this during installation.
11083	TCP	<code>nmsas.server.port.naming.rmi</code>	Default port for RMI naming service	Modify the <code>server.properties</code> file.
11084	TCP	<code>nmsas.server.port.jmx.jrmp</code>	Default RMI object port (JRMP invoker)	Modify the <code>server.properties</code> file.
11085	TCP	<code>nmsas.server.port.jmx.rmi</code>	Default RMI pooled invoker port	Modify the <code>server.properties</code> file.
11086	TCP	<code>nmsas.server.port.invoker.unified</code>	Default RMI remoting server connector port	Modify the <code>server.properties</code> file.
11087	TCP	<code>nmsas.server.port.hq</code>	Used for un-encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
11089	TCP	<code>nmsas.server.port.remoting.ejb3</code>	Default EJB3 remoting connector port	Modify the <code>server.properties</code> file.
11092	TCP	<code>nmsas.server.port.hq.ssl</code>	Used for encrypted Global Network Management traffic.	Modify the <code>server.properties</code> file.
11099	TCP	<code>nmsas.server.port.naming.port</code>	Default bootstrap JNP service port (JNDI provider)	Modify the <code>server.properties</code> file. You can also change this during installation.

Table 50 Ports Used on the NNM iSPI Performance for Traffic Management Server (Traffic Leaf) (cont'd)

Port	Type	Name	Purpose	Change Configuration
11712	TCP	nmsas.server.port.ts.recovery	Default recovery port used by the Transaction service.	Modify the server.properties file.
11713	TCP	nmsas.server.port.ts.status	Default status port used by the Transaction service.	Modify the server.properties file.
11714	TCP	nmsas.server.port.ts.id	Default port used by the Transaction service.	Modify the server.properties file.

Table 51 shows the ports the NNM iSPI Performance for QA uses on the management server. In case of port conflicts, almost all of these port numbers can be changed using the server.properties file located at: %NnmDataDir%/nmsas/qa/server.properties.

Table 51 Ports Used on the NNM iSPI Performance for QA Management Server

Port	Type	Name	Purpose	Change Configuration
5432	TCP	com.hp.ov.nms.postgres.port	This PostgreSQL port is the port the embedded database listens on for this NNMi management server. The port is expected to be the same as that configured for NNMi in the nms-local.properties file	N/A
54040	TCP	nmsas.server.port.web.http	Default HTTP port - used for Web UI.	Modify the server.properties file. You can also change this during installation.
54043	TCP	nmsas.server.port.web.https	Default secure HTTPS port (SSL) - used for Web UI.	Modify the server.properties file. You can also change this during installation.
54046	TCP	nmsas.server.port.naming.port	Default bootstrap JNP service port (JNDI provider)	Modify the server.properties file. You can also change this during installation.
54047	TCP	nmsas.server.port.naming.rmi	Default port for RMI naming service	Modify the server.properties file.

Table 51 Ports Used on the NNM iSPI Performance for QA Management Server (cont'd)

Port	Type	Name	Purpose	Change Configuration
54084	TCP	nmsas.server.port.jmx.jrmp	Default RMI object port (JRMP invoker)	Modify the server.properties file.
54085	TCP	nmsas.server.port.jmx.rmi	Default RMI pooled invoker port	Modify the server.properties file.
54086	TCP	nmsas.server.port.invoker.unified	Default RMI remoting server connector port	Modify the server.properties file.
54087	TCP	nmsas.server.port.hq	Used for un-encrypted Global Network Management traffic.	Modify the server.properties file.
54088	TCP	nmsas.server.port.hq.ssl	Used for encrypted Global Network Management traffic.	Modify the server.properties file.
54089	TCP	nmsas.server.port.remoting.ejb3	Default EJB3 remoting connector port	Modify the server.properties file.
54712	TCP	nmsas.server.port.ts.recovery	Default recovery port used by the Transaction service.	Modify the server.properties file.
54713	TCP	nmsas.server.port.ts.status	Default status port used by the Transaction service.	Modify the server.properties file.
54714	TCP	nmsas.server.port.ts.id	Default port used by the Transaction service.	Modify the server.properties file.

Table 52 shows the ports required for NNM iSPI Performance for Metrics and Network Performance Server (NPS). In case of port conflicts, almost all of these port numbers can be changed.

Table 52 Required Ports for NNM iSPI Performance for Metrics and NPS

Port	Type	Name	Purpose	Change Configuration
9300	TCP	NPS UI	Default HTTP port - used for Web UI & BI Web Services.	Change using 'configureWebAccess.ovpl'.
9305	TCP	NPS UI - SSL	Default Secure HTTPS port (SSL) - used for Web UI & BI Web Services.	Change using 'configureWebAccess.ovpl'.

NOTE: If NNM and NPS are not coexisting, then the network ports used for the OS network file sharing are also required (NFS services on Linux, Windows File Sharing on Windows).

Ports used by processes running on the same server (in other words, not used for communication between servers over the network)

9301	TCP	Sybase ASE	Sybase ASE (BI Content Manager Database)	Change not supported.
9302	TCP	Sybase IQ Agent	Sybase IQ Agent service	Change not supported.

Table 52 Required Ports for NNM iSPI Performance for Metrics and NPS (cont'd)

Port	Type	Name	Purpose	Change Configuration
9303	TCP	Sybase IQ - PerfSPI DB	Sybase IQ database used to store all NPS extensionPack data.	Change not supported.
9304	TCP	Sybase IQ - PerfSPI DEMO DB	Sybase IQ database used to store extensionPack DEMO data.	Change not supported.
9306	TCP	Database SQL Rewrite Proxy - PerfSPI DB	SQL Rewrite proxy for the Perfspi database - used by BI Server.	Change not supported.
9307	TCP	Database SQL Rewrite Proxy - PerfSPI DEMO DB	SQL Rewrite proxy for the Perfspi DEMO database - used by BI Server.	Change not supported.
9308	TCP	Sybase ASE Backup Server	Sybase ASE backup server for the BI content manager database.	Change not supported.

Table 53 shows the ports used by the NNM iSPI NET diagnostics server. The NNM iSPI NET diagnostic server installs HP Operations Orchestration (HP OO). For more information, see the *HP Operations Orchestration Administrator's Guide*.

Table 53 Ports Used by the NNM iSPI NET Diagnostics Server

Port	Type	Name	Purpose	Change Configuration
3306	TCP	MySQL database port	Provides access to MySQL database.	Change not supported.
8080	TCP	jetty http port	Default HTTP port - used for Web UI & Web Services.	Post-install modifications not supported.
8443	TCP	jetty SSL/https port	Default HTTPS port - used for Web UI & Web Services.	Post-install modifications not supported.
9004	TCP	HP OO RAS port	Provides access to HP OO Remote Action Service.	Change not supported.

Suggested Configuration Changes

Some common issues and how to address them.

Problems and Solutions

Problem: NNMi does not always interpret and display SNMP data and MIB strings correctly.

Solution: This is caused by NNMi not always knowing which character set to use to interpret this data. The result is that NNMi displays garbled strings from some SNMP traps and other octetstring data, such as `sysDescription`, `sysContact` and other data. The solution is to use the correct character set to interpret this data.

For SNMP traps and other octetstring data that result in garbled text displays due to using improper character sets, do the following:

- a Edit the following file:
 - *Windows:* %NNM_PROPS%\nms-jboss.properties
 - *UNIX:* \$NNM_PROPS/nms-jboss.properties
- b Remove the comment (`#!` characters) from the line that begins as follows:


```
#!com.hp.nnm.sourceEncoding=
```
- c Set the `com.hp.nnm.sourceEncoding` JVM property to a comma-separated list of source encodings that your environment currently supports using the examples shown in the `nms-jboss.properties` file. These examples show combinations of the Shift_JIS, EUC_JP, UTF-8, and ISO-8859-1 character sets.
- d Save your changes.
- e From a command prompt, run **ovstop**.
- f From a command prompt, run **ovstart**.
- g To test your changes, resend the suspect trap to NNMi and make sure the garbled display problem no longer occurs.

If the garbled text involves binary data or data that cannot be interpreted for any reason, do the following to configure NNMi to display the strings in hexadecimal format:

- a Edit the following file:
 - *Windows:* %NNMDATADIR%\shared\nnm\conf\nnmvbnosrcenc.conf
 - *UNIX:* \$NNMDATADIR/shared/nnm/conf/nnmvbnosrcenc.conf
- b Add the trap OID, varbind OID value combinations that NNMi displays in a garbled format. Also add the combinations from any varbind values you do not want NNMi to decode, such as binary data. Use the examples shown in the `nnmvbnosrcenc.conf` file as templates to configure your combinations. This tells NNMi to display the Custom Incident Attribute values in the Incident form using a hexadecimal value.
- c Save your changes.

- d From a command prompt, run **ovstop**.
- e From a command prompt, run **ovstart**.
- f Test your changes to make sure these changes result in a hexadecimal display of the formerly garbled strings.

Problem: NNMi displays messages about license keys not matching the host (the NNMi Management Server)

Solution: This happens after someone installs an NNMi license key created with an IP address that does not match the IP address of the NNMi management server. The solution is to remove the invalid license keys:

- 1 At a command prompt, enter the following command to open the Autopass user interface:

```
nnmlicense.ovpl NNM -gui
```

- 2 On the left side of the Autopass window, click **Remove License Key**.
- 3 Select the invalid license keys.
- 4 Click **Remove**.

Repeat [step 1](#) through [step 4](#) for any other affected NNMi product integrations by replacing **NNM** with the affected product. For example, to work with licenses related to the NNM iSPI Network Engineering Toolset Software, use the following command to open the Autopass user interface:

```
nnmlicense.ovpl iSPI-NET -gui
```

For additional information about licensing, see [Licensing NNMi](#) on page 121.

Problem: NNMi maps show an ESXi server and the virtual machines and servers running on the ESXi server. NNMi shows all of these systems connected by a cloud symbol. This is only a problem if you do not want to see the ESXi server, including the virtual machines and servers, on the NNMi map.

Solution: If you do not want NNMi showing ESXi servers, including the virtual machines and servers, do the following:

- 1 Open the NNMi console.
- 2 Go to the topology map showing the nodes you want to delete; delete the nodes representing the ESXi server and the virtual machines and servers.
- 3 Click **Discovery Configuration** in the **Configuration** workspace.
- 4 Click the **Auto-Discovery Rules** tab.
- 5 Create a new auto-discovery rule.
- 6 Enter a relatively low number in the **Ordering** field to give this rule a high precedence. Make sure the **Discover Included Nodes** check box is not checked.
- 7 Add a new IP address range for this rule.
- 8 For the nodes representing the ESXi server and the virtual machines and servers, add either the individual IP addresses or the IP address ranges for these nodes; then change the **Range Type** to be **Include by Rule** rather than **Ignore by Rule**.

- 9 Click **Save and Close** three times to save your work.



These steps will not delete any existing nodes; however, it prevents future discovery of nodes within the excluded IP address range.

Problem: NNMi maps show a Linux server instead of ESXi servers and nodes.

Solution: You have deployed VMWARE on a Linux server with the Net-SNMP agent enabled. If you want NNMi to discover and show ESXi servers, you must complete a bare metal installation for the ESXi servers and nodes. For more information see <http://www.vmware.com>.

Problem: NNMi maps show ESXi devices as having No SNMP instead of showing them as ESXi devices.

Solution: The ESXi SNMP agent must be installed and enabled for NNMi to discover and map ESXi servers and nodes. Perhaps you uninstalled or disabled the ESXi SNMP agent. To remedy this, install or enable the ESXi SNMP agent. For more information see <http://www.vmware.com>.

Problem: I am using NNMi with an Oracle database. I configured a large node group that results in an error when generating a node group map.

Solution: This could occur if you configure NNMi as follow:

- You use NNMi with an Oracle database.
- You create a top level node group containing child node groups.
- Any of the child node groups contain 1000 or more members.
- You select either or both of the following selections in the **Node Group Map Settings->Connectivity->Node Group Connectivity** section for these node groups:
 - **Nodes to Node Groups**
 - **Node Groups to Node Groups**

To remedy this, limit the child node groups to less than 1000 members or do not select either or both **Nodes to Node Groups** or **Node Groups to Node Groups** in the **Node Group Map Settings->Connectivity->Node Group Connectivity** section for these node groups.

Problem: For some Cisco devices using PAGP (Port Aggregation Protocol), if a link goes down that is part of a port aggregation, NNMi might consider the port on that device to no longer be part of the port aggregation. This can result in NNMi not reporting the degraded state of the port aggregation.

Solution: Beginning with NNMi 9.0x Patch 4, there is a feature that helps NNMi better manage Cisco devices that use PAGP. You can configure this NNMi feature to attempt to determine if a down interface is still configured to be a part of a port aggregation. To enable this feature, do the following:

- 1 Edit the following file:
 - *Windows:* %NNM_PROPS%\nms-disco.properties
 - *UNIX:* \$NNM_PROPS/nms-disco.properties

- 2 Look for the `enablePagpOperDownHeuristic` entry, which is similar to the following line:

```
#!com.hp.ov.nms.disco.enablePagpOperDownHeuristic=false
```

To enable the `enablePagpOperDownHeuristic`, change the line as follows:

```
com.hp.ov.nms.disco.enablePagpOperDownHeuristic=true
```



Make sure to remove the `#!` characters located at the beginning of the line.

- 3 Restart the NNMi management server.
 - a Run the `ovstop` command on the NNMi management server.
 - b Run the `ovstart` command on the NNMi management server.

Problem: You encounter pop-up dialog issues when using Internet Explorer 8 and the Internet Explorer ESC (Enhanced Security Configuration)

Solution: Windows 2003 and Windows 2008 server offer a feature called Internet ESC (Explorer Enhanced Security Configuration) in Internet Explorer 8. After this feature is enabled (this feature is currently enabled by default) all pop-up dialogs and windows are tested against a list of trusted sites. If the URL associated with the pop-up is not in the list of trusted sites, all of the controls in the dialog or window are disabled. For example, when this happens, clicking on the **OK**, **Apply**, and **Cancel** buttons has no effect.

Generally, with ESC enabled, whenever you open a dialog, the browser prompts you as to whether you want to enable the URL associated with the pop-up as a trusted site. To proceed, you must allow the URL. If you do not allow the URL, the dialog controls will not work and you will see the prompt whenever the dialog or window is opened. Eventually the nagging ceases because all important URLs will have been added to the list of Trusted Sites. One special URL that must be placed in the list is `about:blank`.

You can get into a situation where the NNMi console does not work. If, at some point, you click the **Don't show me this message again** check box in the Trusted Sites prompt, subsequent prompts will not be issued. If you had done this before installing NNMi, the NNMi console would hardly function: Dialogs would not pop up, but the controls in the dialog would not work. For example, if you opened the **Help->About** dialog, the **OK** button would not close the dialog. Also, all of the table view filter dialogs would not work. In the latter case, this is because the `about:blank` URL is not in the list of Trusted Sites.

There are several ways to resolve this problem:

- "Disable the Internet Explorer Enhanced Security Configuration feature using Server Manager.
- "Add the required URLs to the Trusted Security Sites using the **IE->Tools->Options->Security** tab; specifically, add `about:blank`.
- "Make sure the IE pop up window to permit additions to the Trusted Security Sites is enabled.

Problem: I accidentally removed the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files library from the NNMi management server.

You can specify a privacy protocol to use for communication with SNMPv3 devices on the **SNMPv3 Settings** form in the NNMi console. The AES-192, AES-256, and TripleDES protocols are available for selection only when the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files library is installed on the NNMi management server.

If you accidentally removed the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files library, and need to enable NNMi to use the AES-192, AES-256, and TripleDES privacy protocols for SNMPv3 communication, follow these steps:

- 1 Download the Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files library from the Oracle Technology Network web site for Java developers (<http://www.oracle.com/technetwork/java/index.html>). A direct link is:
https://cds.sun.com/is-bin/INTERSHOP.enfinity/WFS/CDS-CDS_Developer-Site/en_US/-/USD/ViewProductDetail-Start?ProductRef=jce_policy-6-oth-JPR@CDS-CDS_Developer
- 2 Uncompress the download package, and then copy both JAR files (`local_policy.jar` and `US_export_policy.jar`) to the following location:
 - *Windows:* `%NnmInstallDir%\nonOV\jdk\nnm\jre\lib\security`
 - *UNIX:* `$NnmInstallDir/nonOV/jdk/nnm/jre/lib/security`
- 3 Restart NNMi by running the following commands:
 - a **ovstop**
 - b **ovstart**

Glossary

A

account

See [user account](#).

active cluster node

See [active server](#).

active server

The server currently running the NNMi processes in an application failover or high availability configuration.

address hint

See [discovery hint](#).

application failover

In NNMi, the optional capability (configured by the user and utilizing jboss clustering support) that transfers control of NNMi processes to a standby server if the currently active server fails.

ARP cache

The ARP (Address Resolution Protocol) cache is an operating system table that maps Data Link Layer (OSI Layer 2) addresses to Network Layer (OSI Layer 3) addresses. Data Link Layer addresses are typically MAC addresses, while Network Layer addresses are typically IP addresses. In [rule-based discovery](#), NNMi uses ARP cache entries on discovered nodes (as well as other techniques) to find additional nodes that can be checked against the current discovery rules.

auto-discovery

See [rule-based discovery](#).

C

Causal Engine

NNMi technology that applies [root cause analysis](#) (RCA) to network symptoms, using a [causality](#)-based approach. Causal Engine RCA is triggered by certain occurrences, including changes detected as a

result of [state polling](#), [SNMP traps](#), and specific [incidents](#). The Causal Engine uses RCA to determine the [status](#) of managed objects, to formulate [conclusions](#) about them, and to generate [root cause incidents](#).

causality

Denotes the relationship between one event (the cause) and another event (the effect) which is the direct consequence (result) of the first. NNMi uses causality analysis algorithms to analyze event cycles and identify solutions for resolving network issues.

cluster

In an NNMi context, a grouping of hardware and software, linked by high availability technology or by using jboss clustering capabilities, that works together to ensure functional and data continuity if components overload or fail. The computers in a cluster are commonly connected to each other through high speed LANs. Clusters are usually deployed to improve availability, performance, or both.

cluster member or node

In an NNMi context, a system within a high availability or jboss cluster that has been or will be configured to support NNMi high availability or application failover.

community string

A password-like mechanism used in [SNMPv1](#) and [SNMPv2c](#) implementations to authenticate SNMP queries to SNMP agents. The community string is passed in cleartext in SNMP packets, making it vulnerable to packet sniffing. [SNMPv3](#) provides stronger security mechanisms for authentication.

conclusion

In NNMi, supporting detail generated and used by the [Causal Engine](#) that sheds further light on how the Causal Engine determined [status](#) and [root cause incidents](#) for a managed object.

console

See [NNMi console](#).

controller

In NNMi application failover, a JGroups term for the cluster member that has the master cluster state. JGroups determines which member of the cluster is the controller based on the lowest IP address.

D**discovery hint**

An IP address found by NNMi using an SNMP ARP cache query; a CDP, EDP, or other discovery protocol query; or a ping sweep. NNMi further queries IP addresses found as discovery hints, then checks the results against the current discovery rules in [rule-based discovery](#).

discovery process

The process by which NNMi gathers information about network [nodes](#) so that they can be placed under management. Initial discovery runs as a two-phase process, returning device inventory information and then network connectivity information.

After initial discovery, the discovery process is ongoing. In [list-based discovery](#), this means devices in the list of seeds will be updated if their configuration changes. In [rule-based discovery](#), new devices will also be added if they match current [discovery rules](#). Discovery can also be initiated on demand for a device or set of devices from the [NNMi console](#) or from the command line.

See also [spiral discovery](#), [rule-based discovery](#), and [list-based discovery](#).

discovery rule

A range of user-defined IP addresses, system object IDs ([Object Identifiers](#)), or both used to limit the [rule-based discovery](#) process. Configure discovery rules in the **Discovery Configuration** portion of the NNMi console under **Auto-Discovery Rules**. See also [rule-based discovery](#).

discovery seed

See [seed](#).

E**embedded database**

The database included with NNMi. NNMi can also be configured to use an external Oracle database

instead of the embedded database for most of its tables. See also [PostgreSQL](#).

episode

A term used in NNMi [root cause analysis](#) that refers to a specific duration, triggered by a primary failure, during which secondary failures are suppressed or are correlated under the primary failure.

F**fault polling**

A key NNMi monitoring activity, in which NNMi issues ICMP pings, SNMP read-only queries of status MIBs, or both for its managed interfaces, IP addresses, and SNMP agents to determine the [state](#) of each managed object. Users can customize the types of fault polling performed for different interface groups, node groups, and nodes under **Monitoring Configuration** in the **Configuration** workspace of the [NNMi console](#). Fault polling is a subset of [state polling](#).

G**global manager**

The NNMi management server in a global network management deployment that consolidates data from distributed NNMi regional manager servers. The global manager provides a unified view of topology and incidents across the whole environment. A global manager must have an NNMi Advanced license.

global network management

A distributed deployment of NNMi with one or more global managers consolidating data from one or more geographically distributed regional managers.

H**HA**

See [high availability](#).

HA resource group

In modern [high availability](#) environments such as HP ServiceGuard, Veritas Cluster Server, or Microsoft Cluster Services, applications are represented as compounds of resources, such as the application itself, its shared file systems and a virtual IP address. The resources consist of an *HA resource group*, which represents an application running in a cluster environment.

high availability

Used in this guide to mean a hardware and software configuration that provides for uninterrupted service if part of the configuration fails. High availability (HA) means that the configuration has redundant components to keep applications running at all times even if a component fails. NNMi can be configured to support one of several commercially available HA solutions. Contrast with [application failover](#).

HP Network Node Manager i Software

An HP software product (abbreviated NNMi) designed to aid network administration and to consolidate network management activities, including the ongoing discovery of network nodes, monitoring events, and network fault management. Primarily accessed from the [NNMi console](#).

I

ICMP

See [Internet Control Message Protocol](#).

incident

In NNMi, a notification of an occurrence related to your network, displayed in [NNMi console](#) incident views and forms. NNMi includes a number of **Incident Management** and **Incident Browsing** views that enable users to filter incidents based on incident attributes. Most incident views display incidents generated directly by NNMi (sometimes called *management events*). NNMi also includes views for browsing incidents generated from [SNMP traps](#) and from [NNM 6.x/7.x events](#).

interface

A physical port used to connect a node to the network.

interface group

One of NNMi's primary filtering techniques, where interfaces are grouped together to apply settings to a group or filter visualizations by group. Interface groups can be used for any or all of the following: configuring monitoring, filtering table views, and customizing map views. See also [node group](#).

Internet Control Message Protocol

One of the core protocols of the Internet protocol suite (TCP/IP). ICMP ping is used by NNMi along with SNMP queries for [state polling](#).

iSPI

See [NNM iSPI](#).

L

L2

See [Layer 2](#).

L3

See [Layer 3](#).

Layer 2

Refers to the Data Link Layer of the multi-layered communication model, Open Systems Interconnection (OSI). The data link layer moves data across the physical links in the network. NNMi Layer 2 views provide information about the physical connectivity of devices.

Layer 3

Refers to the Network Layer of the multi-layered communication model, Open Systems Interconnection (OSI). The network layer is concerned with knowing the address of the neighboring nodes in the network, selecting routes, and quality of service. NNMi Layer 3 views provide information about connectivity from a routing perspective.

list-based discovery

A process, based on a list of seeds, that discovers and returns detailed network information *only about the nodes that you specify as seeds*. List-based discovery maintains a limited network inventory for specific queries and tasks. Contrast with [rule-based discovery](#). See also [discovery process](#) and [spiral discovery](#).

logical volume

A computer storage virtualization term referring to an arbitrarily sized space in a [volume group](#) that can be used as a separate file system or as a device swap space. Several of the [high availability](#) products supported by NNMi use logical volumes in their shared file systems.

M

Management Information Base

In SNMP, the collection of data about the managed network, organized hierarchically. The data objects within the management information base refer to characteristics of managed devices. NNMi collects

network management information by making SNMP queries to and receiving SNMP traps from managed [nodes](#) using MIB data objects (sometimes called “MIB objects,” “objects,” or “MIBs”).

management server

The NNMi management server is the computer system on which the NNMi software is installed. The NNMi processes and services run on the NNMi management server. (Prior NNM revisions used the term “NNM management station” for this system.)

MIB

See [Management Information Base](#).

N

NNM 6.x/7.x events

An NNMi term for events forwarded from older NNM management stations to NNMi. NNMi provides incident views for browsing the incidents that NNMi generates from these forwarded events.

NNM iSPI

A Smart Plug-in within the [I](#) family. An NNM iSPI adds functionality to NNMi for a specific technology such as MPLS or for a specific domain such as network engineering.

NNMi

See [HP Network Node Manager i Software](#).

NNMi console

The NNMi user interface. Operators and administrators use the NNMi console for network management tasks in NNMi.

node

In the network context, a computer system or device (for example, printer, router, or bridge) in a network. While nodes that are able to respond to SNMP queries provide NNMi with the most comprehensive management information, NNMi can also perform restricted management of non-SNMP nodes.

node group

One of NNMi’s primary filtering techniques, where nodes are grouped together to apply settings to a group or filter visualizations by group. Node groups can be used for any or all of the following: configuring monitoring, filtering table views, and customizing map views. See also [interface group](#).

O

Object Identifier

In SNMP, a numerical sequence that identifies a [Management Information Base](#) data object. An OID consists of numbers separated by dots in which each number represents a particular data object at that level of the MIB hierarchy. The OID is the numerical equivalent of the MIB object name, for example, the MIB object name

`iso.org.dod.internet.mgmt.mib-2.`

`bgp.bgpTraps.bgpEstablished` is equivalent to its OID `1.3.6.1.2.1.15.0.1`.

OID

See [Object Identifier](#).

ovstart command

A command that starts the NNMi managed processes. Invoked at a command prompt. See the [ovstart](#) reference page, or the UNIX manpage.

ovstatus command

A command that reports the current status of the NNMi managed processes. Can be invoked from the NNMi console (**Tools > NNMi Status**) or at a command prompt. See the [ovstatus](#) reference page, or the UNIX manpage.

ovstop command

A command that stops the NNMi managed processes. Invoked at a command prompt. See the [ovstop](#) reference page, or the UNIX manpage.

P

ping sweep

A network probe technique that sends ICMP ECHO requests to multiple IP addresses to determine which addresses are assigned to responsive nodes. When enabled in [rule-based discovery](#), NNMi can use ping sweep on configured IP address ranges to find additional nodes. Some network administrators block ICMP ECHO requests because ping sweeps can be used in denial-of-service attacks.

port

In a network hardware context, a connector for passing information into and out of a network device.

PostgreSQL

An open source relational database which NNMi uses by default to store information such as topology,

incidents, and configuration information. NNMi can also be configured to use Oracle instead of PostgreSQL for most of its tables.

public key certificate

Used in network security and encryption, a file that incorporates a digital signature to bind together a public key with identity information. A certificate is used to verify that a public key belongs to an individual or organization. NNMi uses SSL certificates, which contain a public key and a private key, for authentication and encryption of client-server communication.

R

RCA

See [root cause analysis](#).

region

In NNMi, a grouping of devices for the purpose of configuring communication settings such as timeout values and access credentials.

regional manager

The NNMi management server in a global network management deployment that provides discovery, polling and trap reception for devices and forwards information to the global manager.

role

See [user role](#).

root cause analysis

In NNMi, root cause analysis (RCA) refers to a class of problem solving methods used by NNMi to determine root causes for network issues. In NNMi, the root cause is the actionable issue that will resolve associated problem symptoms if it is addressed. NNMi uses the identification of the root cause in two key ways: to notify the user of the actionable problem and to suppress reporting of secondary problem symptoms until the root cause issue has been resolved. Determination of root cause might result in status changes for managed objects, generation of [root cause incidents](#), or both.

A example of how NNMi uses RCA is the scenario in which a managed router fails, and managed nodes on the other side of the router from the NNMi [management server](#) can no longer respond to [state polling](#) queries. NNMi uses RCA to determine that the state polling failures are secondary problem symptoms. It reports the router failure as the root cause incident and refrains from reporting the

problem symptoms for the downstream nodes until the root cause router failure is resolved.

root cause incident

An NNMi [incident](#) in which the *Correlation Nature* attribute is set to *Root Cause*. NNMi uses [root cause analysis](#) (RCA) to establish the root cause incident as the actionable issue that will resolve associated problem symptoms if it is addressed. See [root cause analysis](#).

rule

See [discovery rule](#).

rule-based discovery

Often called *auto-discovery*, NNMi can use rule-based discovery to seek out [nodes](#) that NNMi should add to its database, following user-specified [discovery rules](#). NNMi looks for [discovery hints](#) in data from discovered nodes, then checks these candidates against the specified discovery rules. Configure discovery rules in the **Discovery Configuration** portion of the NNMi console under **Auto-Discovery Rules**. Contrast with [list-based discovery](#).

S

seed

A network node that helps NNMi discover your network by acting as a starting point for the network discovery process. For example, a seed might be a core router in your management environment. Each seed is identified by an IP address or host name. Unless [rule-based discovery](#) has been configured, NNMi's discovery process is limited to [list-based discovery](#) of specified seeds.

seeded discovery

See [list-based discovery](#).

Simple Network Management Protocol

A simple protocol operating at the application layer (Layer 7) of the OSI model, by which management information for a network element can be inspected or altered by remote users. SNMP is the predominant protocol used by NNMi to exchange network management information with agent processes on managed nodes. NNMi supports the three most common versions of SNMP: SNMPv1, SNMPv2c, and SNMPv3.

SNMP

See [Simple Network Management Protocol](#).

SNMP trap

Network management using polling (solicited responses from SNMP agents) is an SNMP design principle that promotes simplicity. However, the protocol does provide for communication of unsolicited messages from SNMP agents to the SNMP manager process (in this case, NNMi). Unsolicited agent messages are known as “traps” and are generated by SNMP agents in response to internal state changes or fault conditions. NNMi generates [incidents](#) from received SNMP traps, displayed in the **SNMP Traps** incident browsing view.

SNMP trap storm

A high number of unsolicited SNMP agent messages that can overwhelm an SNMP manager process (in this case, NNMi). You can configure SNMP trap storm thresholds in NNMi, using the `nmtrapconfig.ovpl` command. NNMi blocks traps when incoming trap rates exceed the specified threshold rate, until the trap rates fall below the re-arm rate.

spiral discovery

NNMi’s ongoing refinement of network topology information, which includes information about inventory, containment, relationships, and connectivity in networks managed by NNMi. See also [discovery process](#), [rule-based discovery](#), and [list-based discovery](#).

state

NNMi generally uses the term **state** for self-reported managed object responses related to MIB II `ifAdminStatus`, MIB II `ifOperStatus`, performance, or availability. Contrast with [status](#).

state polling

The directed monitoring performed by NNMi’s State Poller, which uses ICMP ping and SNMP queries to retrieve fault, performance, component health, and availability data from managed objects. See also [fault polling](#).

status

In NNMi, an attribute of a managed object that indicates its overall health. The status is calculated by the [Causal Engine](#) from the managed object’s outstanding [conclusions](#). Contrast with [state](#).

sysObjectID

See [system object ID](#).

system account

In NNMi, a special account provided for use during NNMi installation. After installation, the NNMi system account should only be used for command-line security and for recovery purposes. Contrast with [user account](#).

system object ID

In NNMi, a specialized term for an SNMP [Object Identifier](#) that identifies a model or type of network element. The system object ID is part of a network element’s [Management Information Base](#) object, which is queried by NNMi from individual nodes during discovery. Examples of network element types that can be classified by their system object IDs include any member of the HP ProCurve switch family, an HP J8715A ProCurve Switch, and an HP SNMP agent for HP IPF systems. Other vendors’ network elements can be likewise classified according to their system object IDs. A key use for the system object ID is in defining NNMi Device Profiles, which specify characteristics of network elements that can be deduced once a network element’s type is known.

T**topology (network)**

In communication networks, a schematic description of the arrangement of a network, including its nodes and connections.

trap

See [SNMP trap](#).

U**unconnected interface**

From NNMi’s perspective, an unconnected interface is an interface that is not connected to another device discovered by NNMi. By default, the only unconnected interfaces that NNMi monitors are those that have IP addresses *and* are contained in nodes from the **Routers** node group.

user account

In NNMi, a way to provide access to NNMi for users or groups of users. NNMi user accounts are set up in the NNMi console and implement predetermined user roles. See [system account](#) and [user role](#).

user role

As part of setting up user access, the NNMi administrator assigns a pre-configured user role to

each NNMi user account. User roles determine which user accounts can access the NNMi console, as well as which workspaces and actions are available to each user account. NNMi provides the following hierarchical user roles, which are predefined by the program and cannot be modified: *Administrator*, *Web Service Client*, *Operator Level 2*, *Operator Level 1*, *Guest*. See also [user account](#).

V

virtual host name

The host name associated with a [virtual IP address](#).

virtual IP address

An IP address that is not tied to any particular network hardware, used in high availability configurations to send uninterrupted network traffic to the most appropriate server based on current failover or load-balancing needs.

volume group

A computer storage virtualization term referring to one or more disk drives that are configured to form a single large storage area. Several of the [high availability](#) products supported by NNMi use volume groups in their shared file systems.

We appreciate your feedback!

If an email client is configured on this system, by default an email window opens when you click *here*.

If no email client is available, copy the information below to a new message in a web mail client, and then send this message to **ovdoc-nsm@hp.com**.

Product name and version: NNMi 9.20

Document title: *NNMi Deployment Reference*

Feedback:

