HP Network Node Manager i Software

Step-by-Step Guide to Monitoring Devices Located Behind a Static NAT Gateway

NNMi 9.1x Patch 1

You can configure NNMi to monitor devices using static Network Address Translation (Static NAT). This paper describes how to configure NNMi to monitor devices located behind the NAT gateway using SNMP and ICMP polling and SNMP traps.

CONTENTS

Problem Statement

When NNMi discovers a node located behind the Network Address Translation (NAT) gateway at a remote site, NNMi uses the public routable address that the NAT gateway assigned to it. However, the node itself cannot identify the address that the NAT gateway assigned to it. Typically, such nodes have non-globally-routable addresses assigned to them for routing within the remote site. A benefit of using NAT is that remote sites can have overlapping IP addresses because their addresses are unique *within their local domain*. However, this causes challenges for NNMi.

By default, NNMi expects to find a public routable address for a node within a node's IP address table. However, this is not the case when using NAT, as a node assigned a NAT address cannot identify its NAT address. Under these conditions, NNMi may disqualify the node from discovery and discard the node.

Additionally, problems can arise when NNMi receives traps from nodes behind the NAT gateway as these nodes may have a source address of the non-routable address rather than the NAT assigned global address. NNMi is unable to distinguish which node sent the traps.

Figure 1: Sample NAT Environment

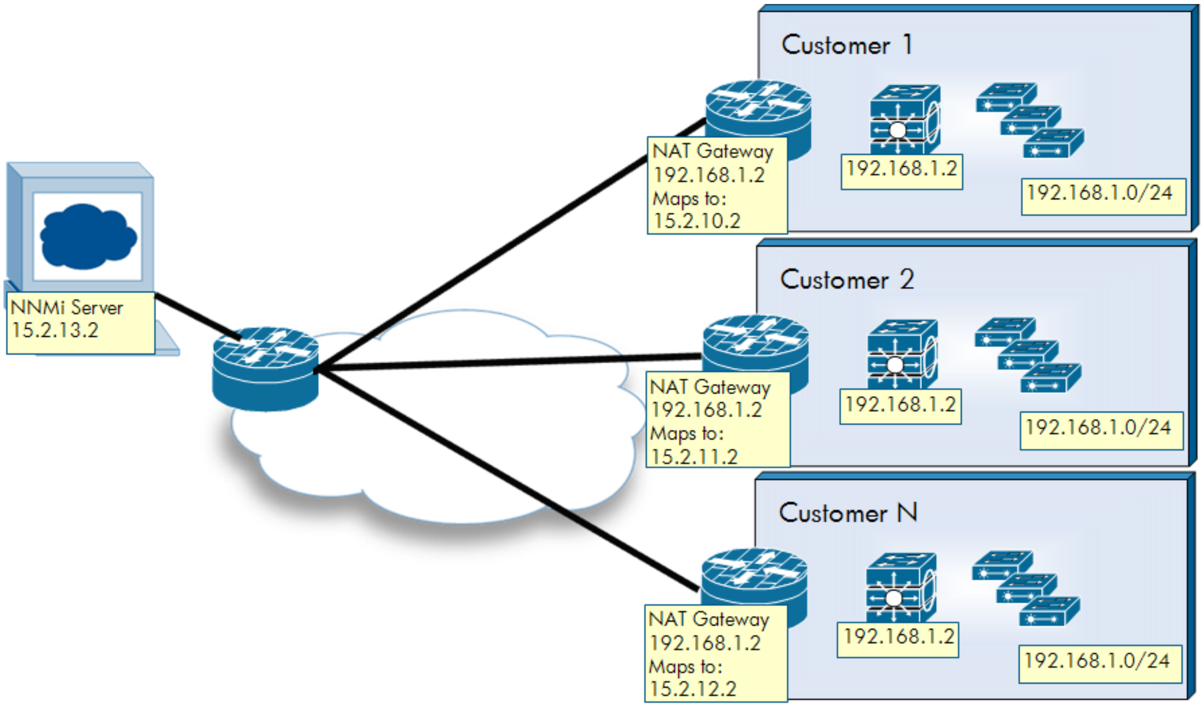

Solution

In NNMi 9.10, NNMi discovers and monitors nodes even if the management IP address is not in the IP address table of the node.

NNMi discovers layer 2 topology for nodes behind the NAT without any additional configuration changes. This is because protocols such as Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP) usually are not IP based but name based. Forwarding database analysis also works without change because it is Media Access Control (MAC) based rather than IP based.

For trap handling, you only need to make a simple change in the NNMi configuration.

You can also take advantage of the new Multi-Tenancy and Security Groups features to separate the overlapping address domains for better presentation of the nodes. See the *HP Network Node Manager i Software Deployment Reference* for more information.

With regard to this solution, note the following:

- This solution supports static NAT only.
- In addition to SNMP monitoring, now using ICMP to monitor management addresses is supported. (ICMP polling to non-routable addresses is not supported.)
- If you have overlapping IP addresses, you may need to filter layer 3 maps for proper viewing.

Summary of Steps

This document shows a simple configuration example. The basic steps include:

1. Obtain Routable Addresses
2. Set up SNMP Communication
3. Disable Small Subnets Connection Rule
4. Optionally Configure a Tenant and Security Group for each Site
5. Build a Node Group
6. Load Seeds for Discovery
7. Configure SNMP Traps

Obtain Routable Addresses

You need to know the routable address for each managed node that uses a NAT address. Obtain this information from your NAT gateway administrator.

Set up SNMP Communication

Set up SNMP communication for the routable addresses of each site as you normally would.

Disable Small Subnets Connection Rule

Because your network likely contains nodes with duplicate IP addresses in NAT environments (typically on different sites), disable the `Small Subnets` discovery rule. This rule enables NNMi to build connections based on IP addresses with /30 subnet masks. Disabling this feature may not be necessary in your environment (see the NNMi help for details). However, if you anticipate that nodes behind the NAT gateway will have some duplicate /30 subnet masks, then disable this feature. You should consider disabling other discovery rules as required by your environment.

To disable the `Small Subnets` connection rule:

1. From the workspace navigation panel, select the **Configuration** workspace, open **Discovery**, and then click **Discovery Configuration**.

Figure 2: Configuration: Discovery Configuration

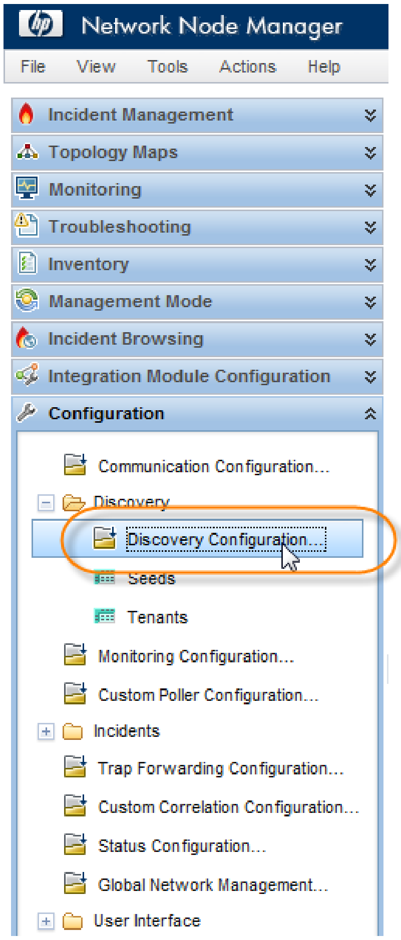
Network Node Manager
File
View
Tools
Actions
Help
Incident Management
Topology Maps
Monitoring
Troubleshooting
Inventory
Management Mode
Incident Browsing
Integration Module Configuration
Configuration
Communication Configuration...
Discovery
Discovery Configuration...
Seeds
Tenants
Monitoring Configuration...
Custom Poller Configuration...
Incidents
Trap Forwarding Configuration...
Custom Correlation Configuration...
Status Configuration...
Global Network Management...
User Interface

2. Click the **Subnet Connection Rules** tab; then double-click the **Small Subnets** rule.

Figure 3: Subnet Connection Rules Tab: Open the Small Subnets Rule

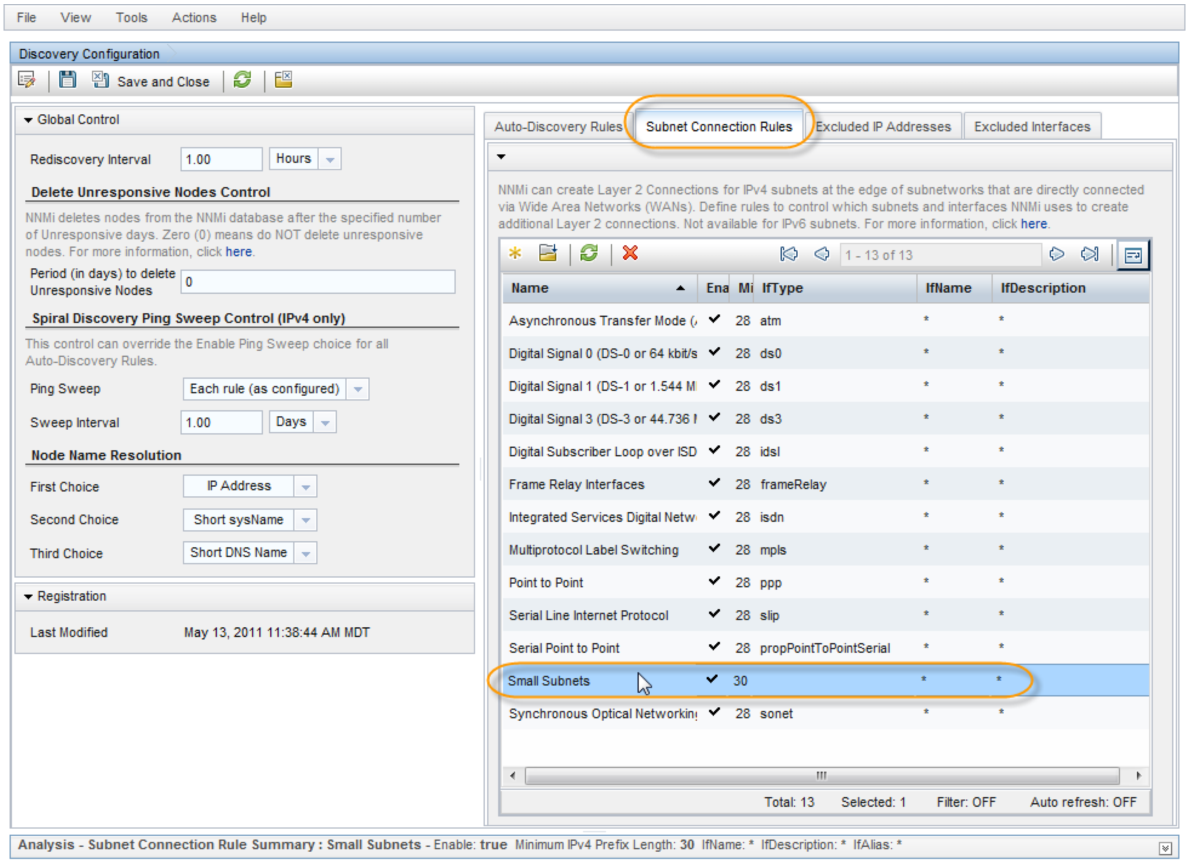

Global Control

Rediscovery Interval 1.00 Hours

Delete Unresponsive Nodes Control

NNMi deletes nodes from the NNMi database after the specified number of Unresponsive days. Zero (0) means do NOT delete unresponsive nodes. For more information, click here.

Period (in days) to delete Unresponsive Nodes 0

Spiral Discovery Ping Sweep Control (IPv4 only)

This control can override the Enable Ping Sweep choice for all Auto-Discovery Rules.

Ping Sweep: Each rule (as configured)

Sweep Interval 1.00 Days

Node Name Resolution

First Choice: IP Address

Second Choice: Short sysName

Third Choice: Short DNS Name

Registration

Last Modified May 13, 2011 11:38:44 AM MDT

NNMi can create Layer 2 Connections for IPv4 subnets at the edge of subnetworks that are directly connected via Wide Area Networks (WANs). Define rules to control which subnets and interfaces NNMi uses to create additional Layer 2 connections. Not available for IPv6 subnets. For more information, click here.

Name	Ena	Mi	IfType	IfName	IfDescription
Asynchronous Transfer Mode (/	✓	28	atm	*	*
Digital Signal 0 (DS-0 or 64 kbit/s	✓	28	ds0	*	*
Digital Signal 1 (DS-1 or 1.544 M	✓	28	ds1	*	*
Digital Signal 3 (DS-3 or 44.736 I	✓	28	ds3	*	*
Digital Subscriber Loop over ISD	✓	28	idsl	*	*
Frame Relay Interfaces	✓	28	frameRelay	*	*
Integrated Services Digital Netw	✓	28	isdn	*	*
Multiprotocol Label Switching	✓	28	mpls	*	*
Point to Point	✓	28	ppp	*	*
Serial Line Internet Protocol	✓	28	slip	*	*
Serial Point to Point	✓	28	propPointToPointSerial	*	*
Small Subnets	✓	30		*	*
Synchronous Optical Networkin	✓	28	sonet	*	*

Total: 13 Selected: 1 Filter: OFF Auto refresh: OFF

Analysis - Subnet Connection Rule Summary : Small Subnets - Enable: true Minimum IPv4 Prefix Length: 30 IfName: * IfDescription: * IfAlias: *

3. Clear the **Enable** check box. Click **Save and Close**, and then click **Save and Close** again.

Figure 4: Subnet Connection Rule: Clear the Enable Check Box

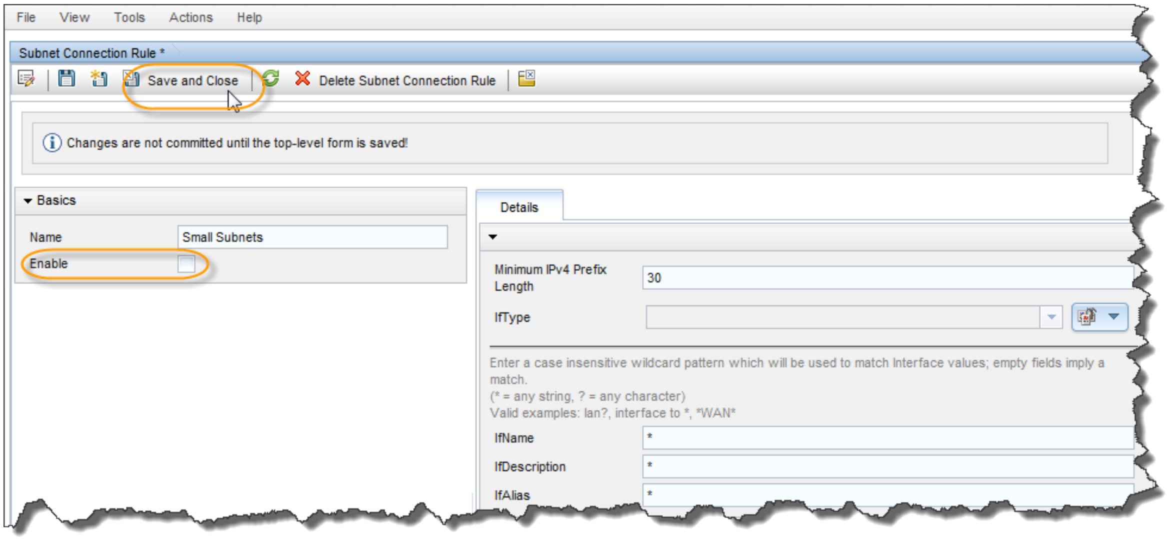

Changes are not committed until the top-level form is saved!

Basics

Name Small Subnets

Enable □

Minimum IPv4 Prefix Length 30

IfType

Enter a case insensitive wildcard pattern which will be used to match Interface values; empty fields imply a match.
(* = any string, ? = any character)
Valid examples: lan?, interface to *, *WAN*

IfName *

IfDescription *

IfAlias *

Optionally Configure a Tenant and Security Group for each Site

It can be very helpful to use the new Multi-Tenancy and Security Group feature to separate the different sites that are behind NAT firewalls. By using this feature, operators can have separate views of the nodes in each site including incidents, tables, and maps. You can create a separate Security Group for the nodes from each site. Then you can associate a Tenant with each Security Group. Set up User Groups to view the Security Groups as desired. After you configure a Tenant and Security Group for each site, each operator can have distinct views of nodes based on the site responsibilities of the operator. You can also set up Node Groups based on the Tenants or Security Groups to help separate sites for administrators or operators that can access multiple sites.

Use the Security Wizard to easily configure a Tenant and a Security Group.

1. From the workspace navigation panel, select the **Configuration** workspace, open **Security**, and then click **Security Wizard**.

Figure 5: Configuration: Click Security Wizard

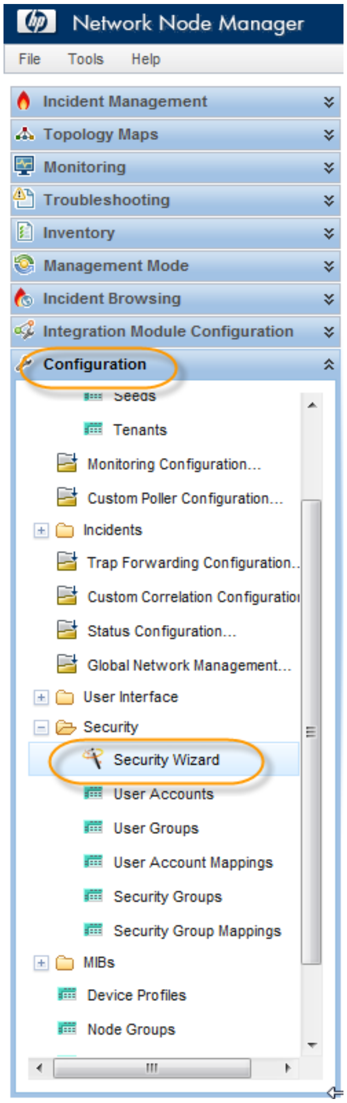

In this example, we have two general sites, Customer1 and Customer2, to represent each remote customer network behind a NAT.

2. Create an Operator User Account and an Operator User Group for each customer site. Then map the User Accounts to the User Groups as shown below. Do this for each customer site.

Figure 6: Security Wizard: Create Operator User and User Group for each Customer Site

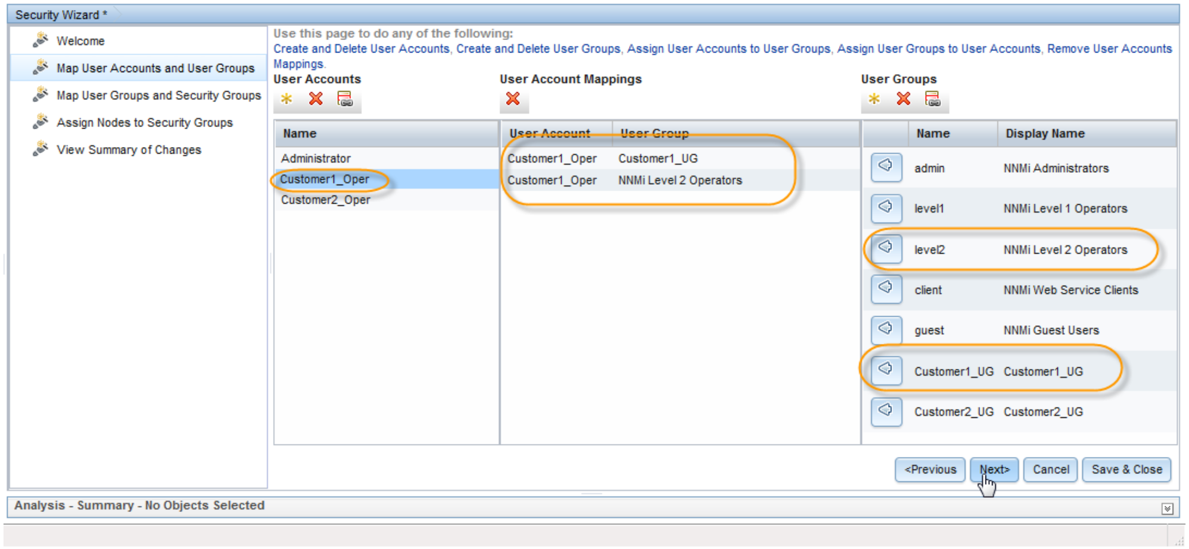
Security Wizard *
Welcome
Map User Accounts and User Groups
Map User Groups and Security Groups
Assign Nodes to Security Groups
View Summary of Changes
<Previous
Next>
Cancel
Save & Close
Analysis - Summary - No Objects Selected

Use this page to do any of the following:
Create and Delete User Accounts, Create and Delete User Groups, Assign User Accounts to User Groups, Assign User Groups to User Accounts, Remove User Accounts Mappings.

User Accounts

Name
Administrator
Customer1_Oper
Customer2_Oper

User Account Mappings

User Account	User Group
Customer1_Oper	Customer1_UG
Customer1_Oper	NNMi Level 2 Operators

User Groups

Name	Display Name
admin	NNMi Administrators
level1	NNMi Level 1 Operators
level2	NNMi Level 2 Operators
client	NNMi Web Service Clients
guest	NNMi Guest Users
Customer1_UG	Customer1_UG
Customer2_UG	Customer2_UG

3. Create a Security Group for each customer site and map the appropriate customer site User Group to the Security Group as shown in the following figure.

Figure 7: Security Wizard: Create Security Group for each Customer Site

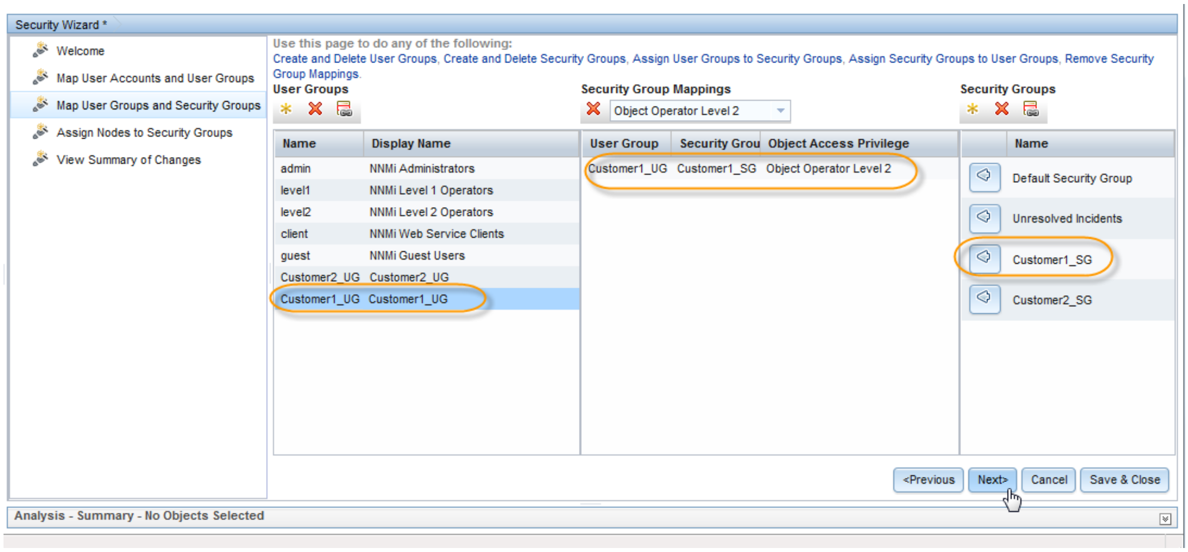
Security Wizard *
Welcome
Map User Accounts and User Groups
Map User Groups and Security Groups
Assign Nodes to Security Groups
View Summary of Changes
<Previous
Next>
Cancel
Save & Close
Analysis - Summary - No Objects Selected

Use this page to do any of the following:
Create and Delete User Groups, Create and Delete Security Groups, Assign User Groups to Security Groups, Assign Security Groups to User Groups, Remove Security Group Mappings.

User Groups

Name	Display Name
admin	NNMi Administrators
level1	NNMi Level 1 Operators
level2	NNMi Level 2 Operators
client	NNMi Web Service Clients
guest	NNMi Guest Users
Customer2_UG	Customer2_UG
Customer1_UG	Customer1_UG

Security Group Mappings
Object Operator Level 2

User Group	Security Grou	Object Access Privilege
Customer1_UG	Customer1_SG	Object Operator Level 2

Security Groups

Name
Default Security Group
Unresolved Incidents
Customer1_SG
Customer2_SG

4. Click **Save and Close** because, at this point, you cannot assign nodes to the Security Groups because nodes have not yet been loaded into NNMi.

Figure 8: Security Wizard: Nodes not Loaded

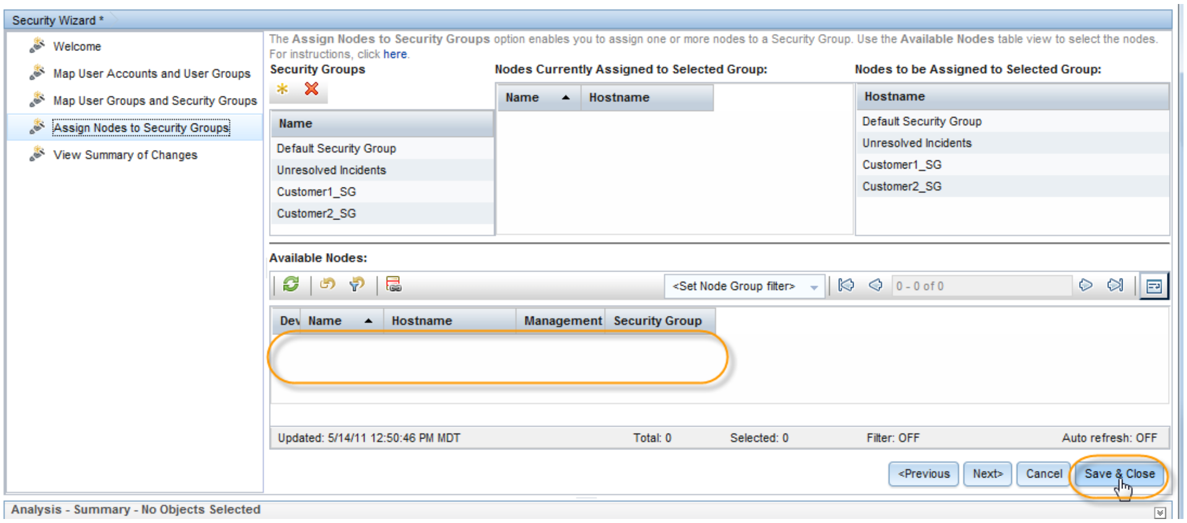
Security Wizard *
Welcome
Map User Accounts and User Groups
Map User Groups and Security Groups
Assign Nodes to Security Groups
View Summary of Changes
Security Groups
Nodes Currently Assigned to Selected Group:
Nodes to be Assigned to Selected Group:
Available Nodes:
<Set Node Group filter>
0 - 0 of 0
<Previous
Next>
Cancel
Save & Close
Analysis - Summary - No Objects Selected

The **Assign Nodes to Security Groups** option enables you to assign one or more nodes to a Security Group. Use the **Available Nodes** table view to select the nodes. For instructions, click here.

Name
Default Security Group
Unresolved Incidents
Customer1_SG
Customer2_SG

Name	Hostname

Hostname
Default Security Group
Unresolved Incidents
Customer1_SG
Customer2_SG

Dev	Name	Hostname	Management	Security Group

Updated: 5/14/11 12:50:46 PM MDT Total: 0 Selected: 0 Filter: OFF Auto refresh: OFF

Now you need to create a Tenant to which you will assign nodes from each customer site. You must create a Tenant so that at the moment a node is loaded with a Tenant assignment, it will automatically go into the correct Security Group.

5. From the workspace navigation panel, select the **Configuration** workspace, open **Discovery**, and then click **Tenants**. Click the ✱ icon.

Figure 9: Discovery Configuration: Create New Tenant

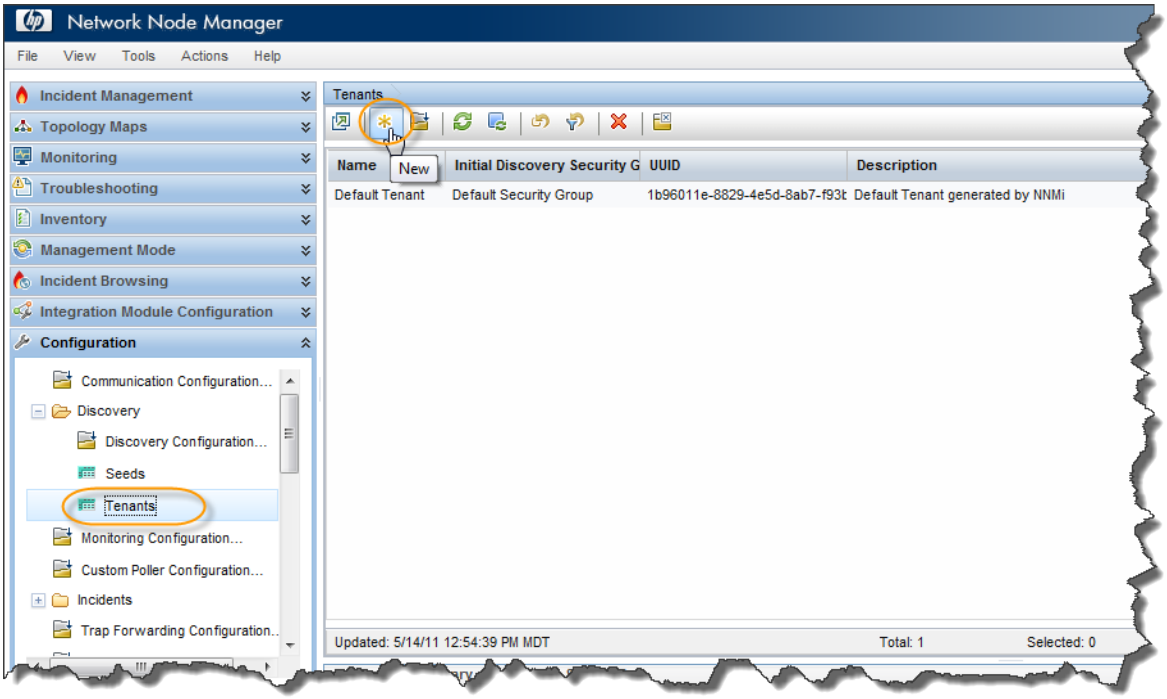

Name	Initial Discovery Security G	UUID	Description
Default Tenant	Default Security Group	1b96011e-8829-4e5d-8ab7-f93t	Default Tenant generated by NNMi

Updated: 5/14/11 12:54:39 PM MDT Total: 1 Selected: 0

6. Enter in the information for this Tenant. Remember to assign an `Initial Discovery Security Group`. Do this for both customer sites in this example.

Figure 10: Tenant Form: Assign Initial Discovery Security Group

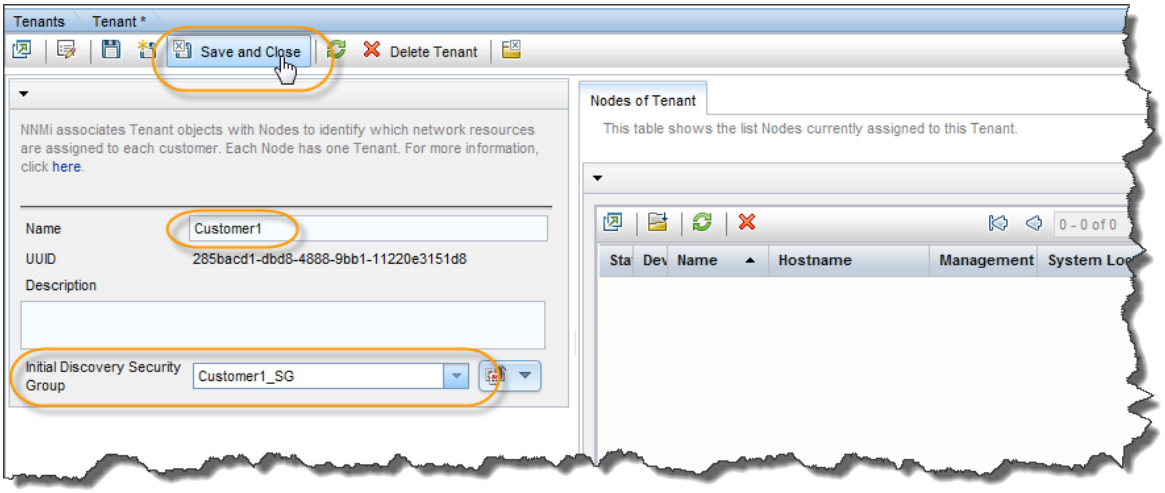

NNMi associates Tenant objects with Nodes to identify which network resources are assigned to each customer. Each Node has one Tenant. For more information, click here.

Name	Customer1
UUID	285bacd1-dbd8-4888-9bb1-11220e3151d8
Description	
Initial Discovery Security Group	Customer1_SG

This table shows the list Nodes currently assigned to this Tenant.

Sta	Dev	Name	Hostname	Management	System Lo

Now you can begin loading nodes into NNMi.

Figure 11: Tenants Form: Initial Discovery Security Groups Assigned

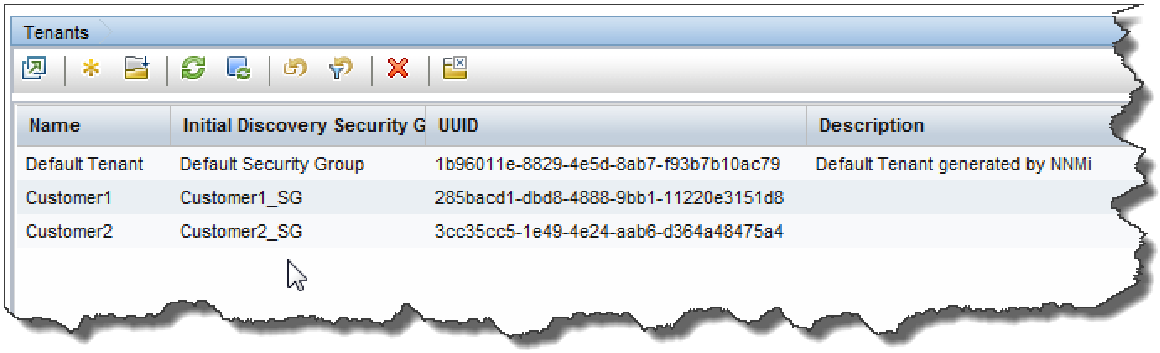

Name	Initial Discovery Security G	UUID	Description
Default Tenant	Default Security Group	1b96011e-8829-4e5d-8ab7-f93b7b10ac79	Default Tenant generated by NNMi
Customer1	Customer1_SG	285bacd1-dbd8-4888-9bb1-11220e3151d8	
Customer2	Customer2_SG	3cc35cc5-1e49-4e24-aab6-d364a48475a4	

Tip: You can use the command line to accomplish steps in this example if you prefer. The primary tool is `nnmsecurity.ovpl`. See the command's reference page for further details.

Build a Node Group

Create node groups based on Tenant for better separation of sites for all users.

Figure 12: Node Group Form: Node Group based on Tenant

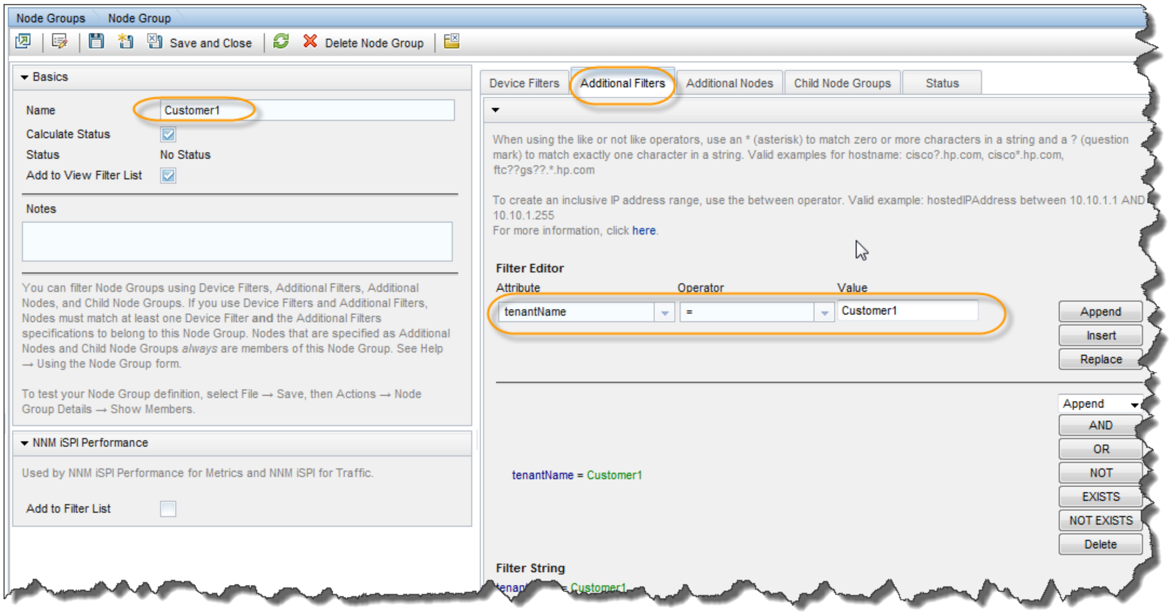

Name: Customer1
Calculate Status: ☑
Status: No Status
Add to View Filter List: ☑

Notes

You can filter Node Groups using Device Filters, Additional Filters, Additional Nodes, and Child Node Groups. If you use Device Filters and Additional Filters, Nodes must match at least one Device Filter **and** the Additional Filters specifications to belong to this Node Group. Nodes that are specified as Additional Nodes and Child Node Groups *always* are members of this Node Group. See Help → Using the Node Group form.

To test your Node Group definition, select File → Save, then Actions → Node Group Details → Show Members.

NNM iSPI Performance

Used by NNM iSPI Performance for Metrics and NNM iSPI for Traffic.

Add to Filter List: □

When using the like or not like operators, use an * (asterisk) to match zero or more characters in a string and a ? (question mark) to match exactly one character in a string. Valid examples for hostname: cisco?.hp.com, cisco*.hp.com, ftc??gs??.*.hp.com

To create an inclusive IP address range, use the between operator. Valid example: hostedIPAddress between 10.10.1.1 AND 10.10.1.255
For more information, click here.

Attribute	Operator	Value
tenantName	=	Customer1

```
tenantName = Customer1
```

Filter String

```
tenan = Customer1
```

Load Seeds for Discovery

Now you can begin discovering seeds using "loading seeds" in Discovery, which is the only way to assign a Tenant and Node Group at initial discovery. This is usually the preferred choice for monitoring customer sites behind NAT firewalls.

Tip: You can load discovery seeds using the graphical user interface (one at a time) or by using the command line.

This example shows the command line method. Repeat the following procedure for each customer site.

1. Create a file with a line for each node containing the routable address of each node. For example:

   seeds_cust1.txt:
   ```
   172.20.4.6
   172.20.4.4
   172.20.4.13
   172.20.4.11
   172.20.4.8
   ```

2. Use the nnmloadseeds.ovpl command line tool to load these seeds into NNMi. Include the "-t" option to assign the Tenant.
   ```
   # nnmloadseeds.ovpl -f seeds_cust1.txt -t Customer1

   5 seeds added
   0 seeds invalid
   0 seeds duplicated
   ```

   After some time, you see the nodes discovered in NNMi. Notice how they have the appropriate Tenant and Security Group assigned.

Figure 13: Nodes Form: Discovered Nodes

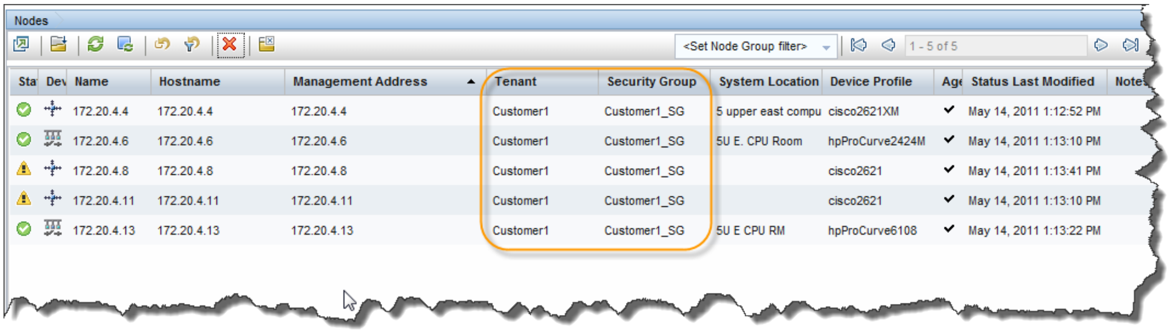

Nodes

<Set Node Group filter> 1 - 5 of 5

Sta	Dev	Name	Hostname	Management Address	Tenant	Security Group	System Location	Device Profile	Age	Status Last Modified	Notes
✓		172.20.4.4	172.20.4.4	172.20.4.4	Customer1	Customer1_SG	5 upper east compu	cisco2621XM	✓	May 14, 2011 1:12:52 PM	
✓		172.20.4.6	172.20.4.6	172.20.4.6	Customer1	Customer1_SG	5U E. CPU Room	hpProCurve2424M	✓	May 14, 2011 1:13:10 PM	
⚠		172.20.4.8	172.20.4.8	172.20.4.8	Customer1	Customer1_SG		cisco2621	✓	May 14, 2011 1:13:41 PM	
⚠		172.20.4.11	172.20.4.11	172.20.4.11	Customer1	Customer1_SG		cisco2621	✓	May 14, 2011 1:13:10 PM	
✓		172.20.4.13	172.20.4.13	172.20.4.13	Customer1	Customer1_SG	5U E CPU RM	hpProCurve6108	✓	May 14, 2011 1:13:22 PM	

3. Open one of these nodes to see that it has a routable management address that is not in the `IP Addresses` table.

Figure 14: Node Form: Open a Node

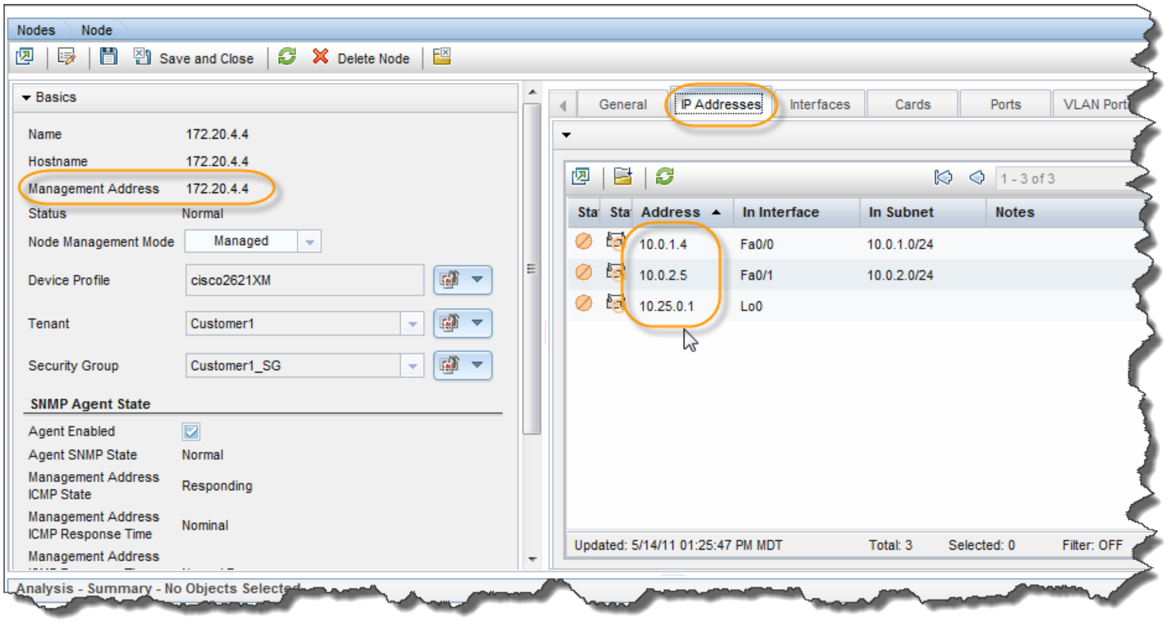

Basics

Name	172.20.4.4
Hostname	172.20.4.4
Management Address	172.20.4.4
Status	Normal
Node Management Mode	Managed
Device Profile	cisco2621XM
Tenant	Customer1
Security Group	Customer1_SG

SNMP Agent State

Agent Enabled	☑
Agent SNMP State	Normal
Management Address ICMP State	Responding
Management Address ICMP Response Time	Nominal
Management Address	

1 - 3 of 3

Sta	Sta	Address	In Interface	In Subnet	Notes
		10.0.1.4	Fa0/0	10.0.1.0/24	
		10.0.2.5	Fa0/1	10.0.2.0/24	
		10.25.0.1	Lo0		

Updated: 5/14/11 01:25:47 PM MDT Total: 3 Selected: 0 Filter: OFF

Analysis - Summary - No Objects Selected

Note: Connectivity may take a few hours to discover.

In this example, you can see in Figure 15 and Figure 16 that NNMi accurately discovered the connectivity. NNMi discovered some connections using CDP and others using Forwarding Database (FDB).

Figure 15: Layer 2 Neighbor View: Properly Connected Nodes (CDP)

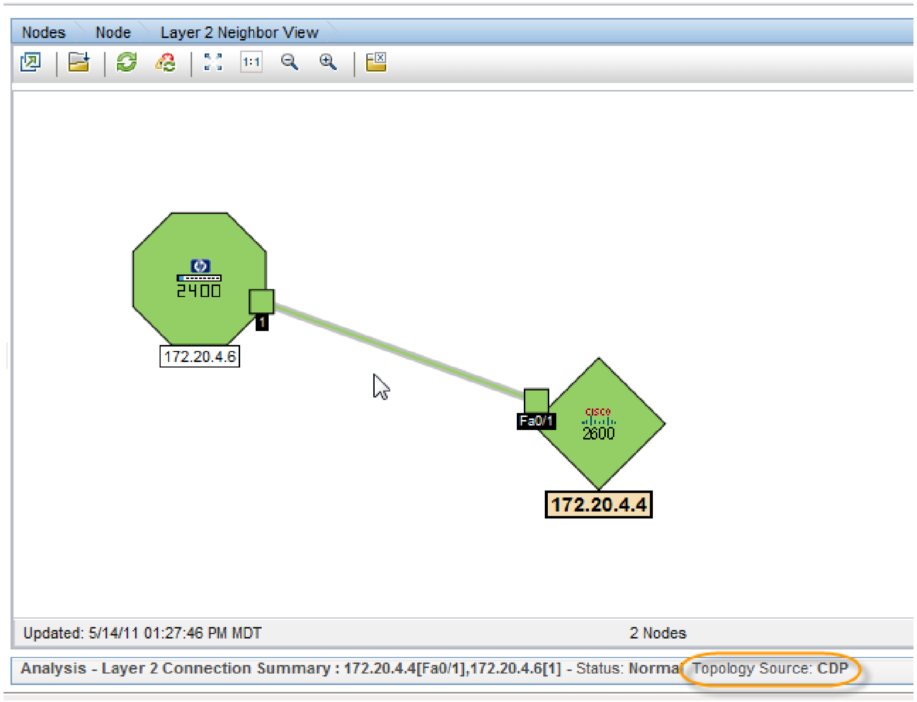

Updated: 5/14/11 01:27:46 PM MDT 2 Nodes

Analysis - Layer 2 Connection Summary : 172.20.4.4[Fa0/1],172.20.4.6[1] - Status: Normal Topology Source: CDP

Figure 16: Layer 2 Neighbor View: Properly Connected Nodes (FDB)

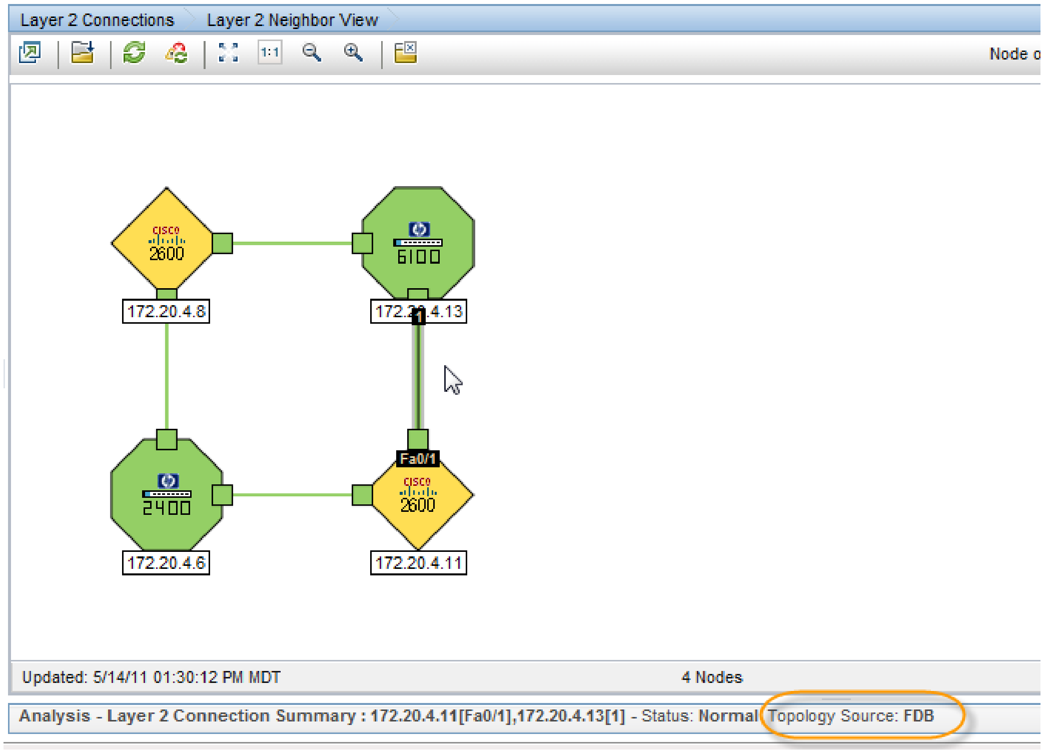

Updated: 5/14/11 01:30:12 PM MDT 4 Nodes

Analysis - Layer 2 Connection Summary : 172.20.4.11[Fa0/1],172.20.4.13[1] - Status: Normal Topology Source: FDB

Sign in to NNMi as **`Customer1_Oper`** and you should see only nodes and incidents related to the Customer1 site, whereas an administrator will see nodes from all sites since administrators are not restricted by Security Groups.

Figure 17: Nodes Form: Customer1_Oper View of Nodes and Incidents

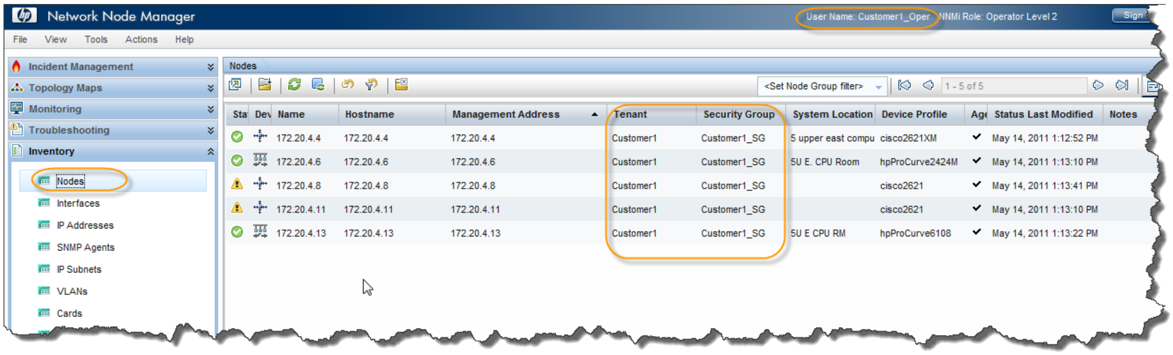

Sta	Dev	Name	Hostname	Management Address	Tenant	Security Group	System Location	Device Profile	Age	Status Last Modified	Notes
✓		172.20.4.4	172.20.4.4	172.20.4.4	Customer1	Customer1_SG	5 upper east compu	cisco2621XM	✓	May 14, 2011 1:12:52 PM	
✓		172.20.4.6	172.20.4.6	172.20.4.6	Customer1	Customer1_SG	5U E. CPU Room	hpProCurve2424M	✓	May 14, 2011 1:13:10 PM	
⚠		172.20.4.8	172.20.4.8	172.20.4.8	Customer1	Customer1_SG		cisco2621	✓	May 14, 2011 1:13:41 PM	
⚠		172.20.4.11	172.20.4.11	172.20.4.11	Customer1	Customer1_SG		cisco2621	✓	May 14, 2011 1:13:10 PM	
✓		172.20.4.13	172.20.4.13	172.20.4.13	Customer1	Customer1_SG	5U E CPU RM	hpProCurve6108	✓	May 14, 2011 1:13:22 PM	

Tip: When working with nodes behind NAT gateways, when selecting a node, a layer 2 neighbor view works well. A layer 3 neighbor view may not give accurate results because multiple nodes share the overlapping IP addresses. NNMi shows these overlapping IP addresses *connected together* in the layer 3 neighbor view; however they are not connected when they are located behind different NAT gateways.

The issue previously discussed will not affect any monitoring or fault analysis because NNMi does not base that analysis on a layer 3 neighbor relationship. But if you have separate NNMi operators with Tenants and Security Group configured, even the layer 3 views will work well because the overlapping nodes will not be visible to the specific operator.

Note: In NNMi version 9.10, the fault monitoring selection, **Enable Management Address ICMP Polling**, is supported for nodes behind NAT firewalls (as shown in the following figure). This means that nodes will be monitored using SNMP and IMCP to determine status. This helps eliminate false node down notifications. Always select the **Enable Management Address ICMP Polling** check box on your polling policies, assuming that you are able to "ping" the management addresses.

Figure 18: Node Settings: Enable Management Address ICMP Polling Check Box

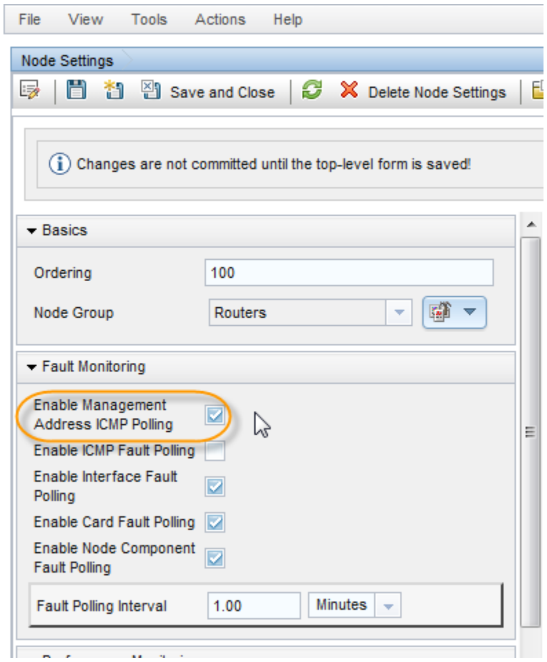

Changes are not committed until the top-level form is saved!

Basics

Ordering: 100

Node Group: Routers

Fault Monitoring

Enable Management Address ICMP Polling ☑

Enable ICMP Fault Polling □

Enable Interface Fault Polling ☑

Enable Card Fault Polling ☑

Enable Node Component Fault Polling ☑

Fault Polling Interval: 1.00 Minutes

Configure SNMP Traps

You must make changes to the managed nodes in order for the NNMi management server to receive SNMP traps from nodes behind the NAT gateway. This example covers two types of SNMP traps: SNMPv2c traps and SNMPv1 traps. This example also shows changes specific to Cisco devices. Other vendors may require similar changes.

Note: Source address resolution is a common challenge with traps. NNMi must unambiguously resolve the source address of each trap that it receives. This problem manifests itself differently depending on the SNMP version (v1 or v2c).

SNMPv2c Traps

Table 1 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.

Version etc.
Source Address
Destination Address
PDU-Type: 4
Request Identifier
Error Status
Error Index
PDU Variable bindings

Table 1 SNMPv2c Trap Format

Since SNMPv2c traps do not have an `Agent Address` field in the PDU, the only source field of the trap is within the IP Packet Header. NAT routers properly translate the source field.
Only one step is required on the source node: make sure that the interface associated with the private inside IP address sources all traps from devices behind the NAT router. The NAT gateway can then translate the trap to the correct public address.

Figure 19 shows an example of this correct translation from the NAT gateway. You can see that 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 19: SNMPv2c Example

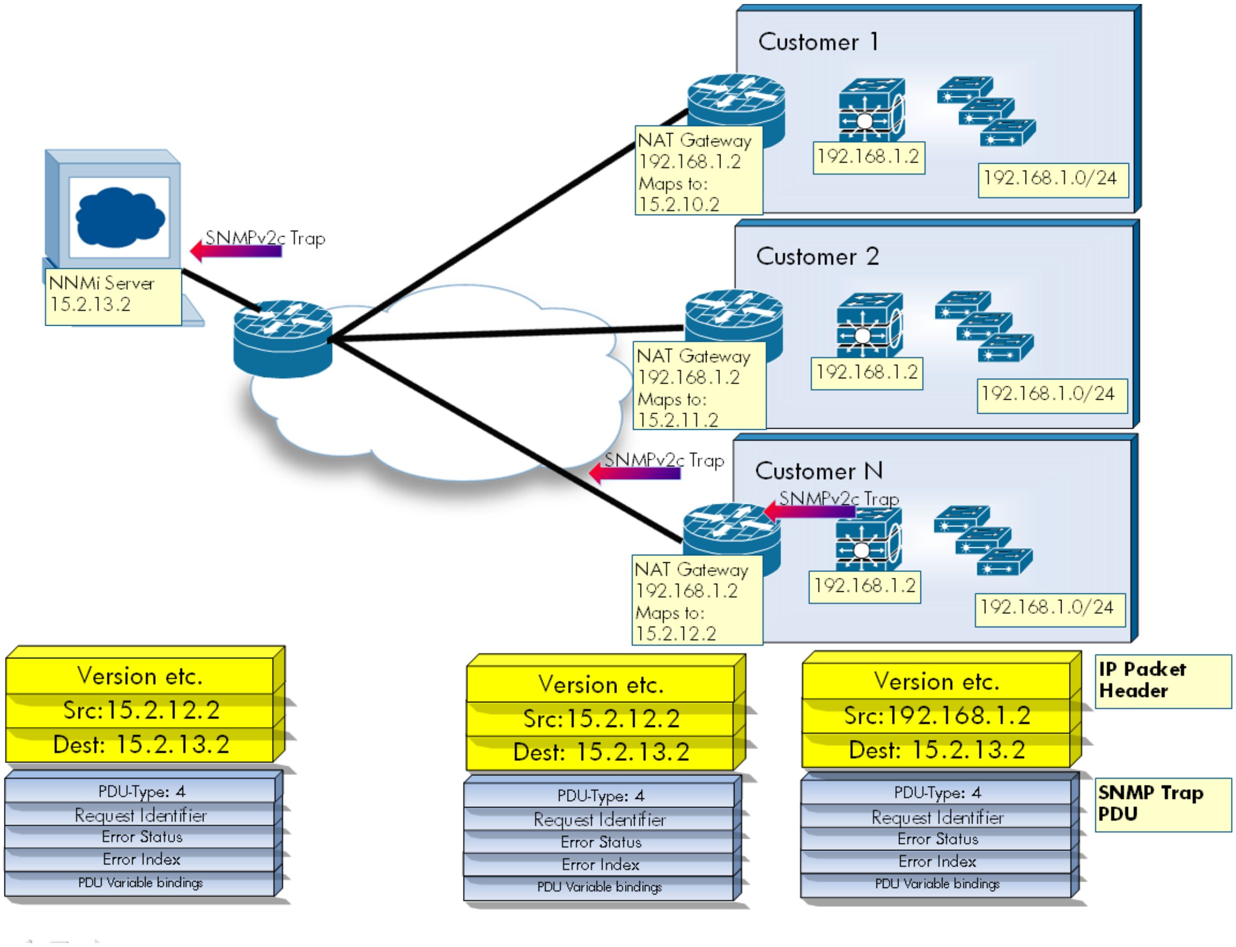

SNMPv1 Traps

SNMPv1 traps are more complex because they embed the `Agent Address` inside the SNMP Trap PDU. Table 2 shows the format of an SNMPv1 trap, with the IP Header forming the top section of the table and the SNMP Trap PDU forming the lower section of the table.

Version etc.
Source Address
Destination Address
PDU-Type: 4
Enterprise
Agent Address
Generic Trap Code
Specific Trap Code
Timestamp
PDU Variable Bindings

Table 2: SNMPv1 Trap Format

Because the `Agent Address` is embedded in the PDU rather than the Header, usually the NAT router will *not* translate this value. NNMi can note the address in the Header and ignore the Agent Address in the payload. To enable this change, do the following:

1. Edit the file `$DataDir/shared/nnm/conf/props/nms-jboss.properties`. Find the line:

   ```
   #!com.hp.nnm.trapd.useUdpHeaderIpAddress=false
   ```

2. Change the value to **`true`** and remove the #! characters as shown below:

   ```
   com.hp.nnm.trapd.useUdpHeaderIpAddress=true
   ```

3. Save the file and restart NNMi.

Figure 20 shows an example of an SNMPv1 trap where NNMi ignores the conflicting IP address fields.

Figure 20: SNMPV1 Example

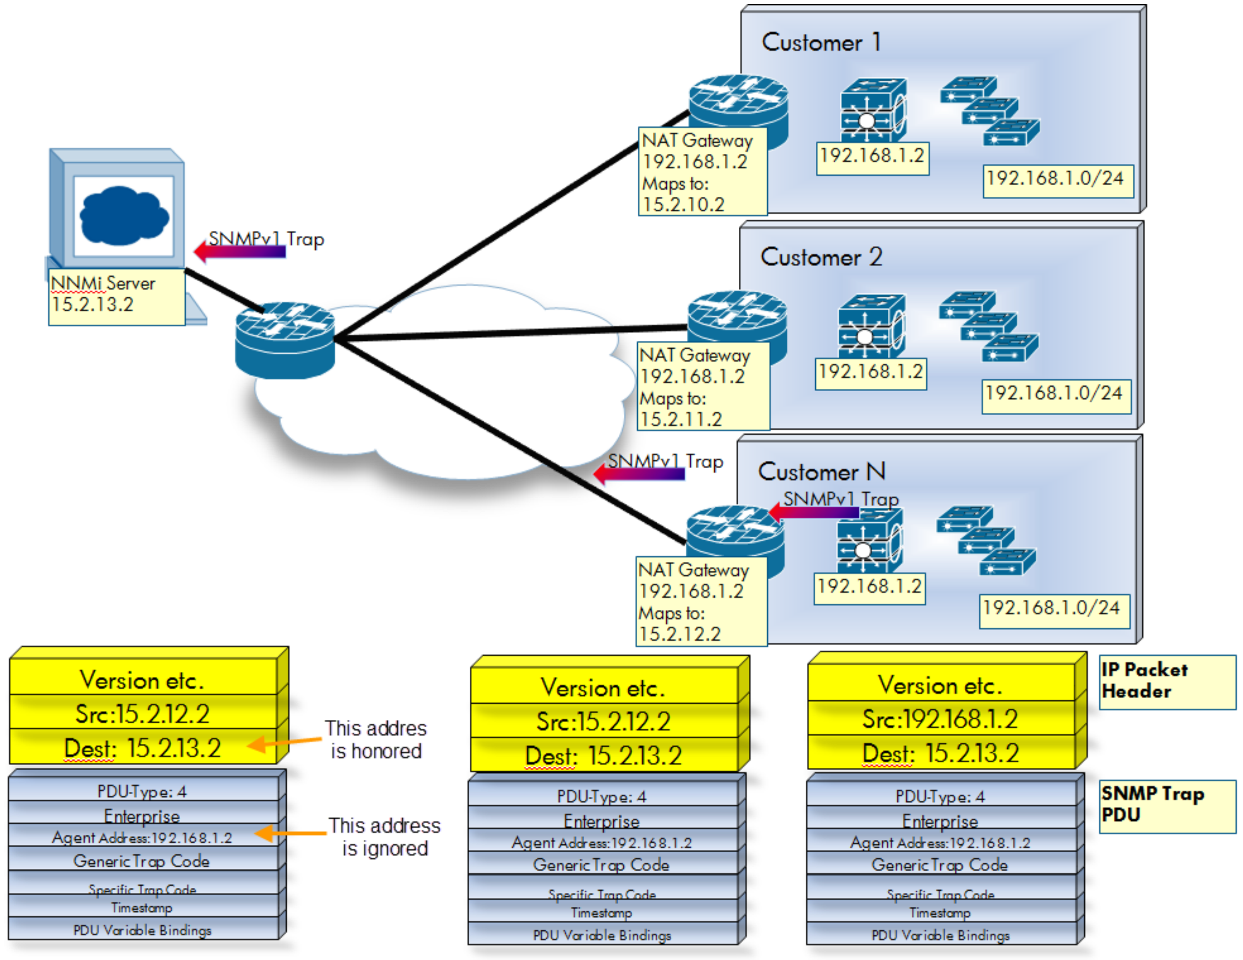

Conclusion

This document has presented the steps necessary to configure NNMi to monitor devices located behind the NAT gateway. By following the steps in this document, you can more effectively monitor networks that contain devices using static NAT.

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 Notices

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)

Support

Visit the HP Software Support 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 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