HP Network Node Manager i Software

Network Bandwidth Utilization for Standalone and Global Network Management Environments

NNMi 9.20

This document provides examples of the expected network utilization for the NNMi release 9.20 standalone and Global Network Management environments. When using this document, note the following:

- Use this document as a guideline only.
- The majority of NNMi traffic is SNMP and ICMP.
- The Global Network Management traffic is TCP.
- The scenarios included in this document only measure traffic generated directly by NNMi.
- Other network traffic is likely to increase (for example, ARP and RARP) by using NNMi or any other network management software.

Contents

Introduction

This document describes the amount of network traffic generated by NNMi during different periods of common use. Use this document to determine in general how much network traffic might be used by NNMi in your network. NNMi is used in many ways, and each network is different (performance in your network environment might vary).

The first section (Standalone NNMi Network Utilization) documents network utilization of one NNMi management server. The second section (Global Network Management Network Utilization) documents the amount of network traffic generated in the Global Network Management environment. The third section documents Application Failover scenarios.

Standalone NNMi Network Utilization

Standalone System – Initial Discovery

This scenario measured the volume of SNMP traffic generated on one NNMi management server during the initial discovery cycle. The time required for this initial discovery cycle depends on your network speed, the number and type of devices in your network, and the hardware on which NNMi is installed.

This test measured utilization during initial discovery of 1,500 nodes.

- No traps were received or generated by NNMi during this scenario.
- All NNMi Monitoring (SNMP and ICMP polling) was disabled to ensure that this scenario messaged only discovery traffic.

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	Specific Settings	Objects
Communication Configuration	Communication Region (one defined): ICMP Settings: (1) ☐ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	
Discovery Configuration	Ping Sweep: none **Schedule Settings:** **Redisovery Interval: 10 days** Auto-Discovery Rules: none Subnet Connection Rules: none Excluded IP Addresses: none Excluded Interfaces: none **Discovery Seeds**: 1500	1500 Nodes
Monitoring Configuration	Disable all current polling configurations: **Enable State Polling:** ☐	
Incident Configuration	SNMP Traps: (none)	

Results: Averages of 1,030.08 packets and 2.71 Megabits per second were measured for this scenario.

Standalone System – Rediscovery

This scenario measured the volume of SNMP traffic generated on one NNMi management server during the rediscovery cycle. The time required for this rediscovery cycle depends on your network speed, the number and types of devices on your network, and the hardware on which NNMi is installed.

All polling was disabled during the rediscovery of the 1,500 nodes. This rediscovery scenario measured the utilization of a rediscovery cycle. This measurement was taken during the first rediscovery cycle, which typically uses the most bandwidth. Over time, rediscovery extends over your configured rediscovery period and average bandwidth for rediscovery decreases.

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	Specific Settings	Objects
Communication Configuration	Communication Region (one defined): ICMP Settings: (1) ☐ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	
Discovery Configuration	Ping Sweep: none Auto-Discovery Rules: none Subnet Connection Rules: none Excluded IP Addresses: none Excluded Interfaces: none **Discovery Seeds**: 1500	1500 Nodes
Monitoring Configuration	Disable all current polling configurations: **Enable State Polling:** ☐	
Incident Configuration	SNMP Traps: (none)	

Results: Averages of 2,113.329 packets and 5.6045 Megabits per second were measured for this scenario.

Standalone System – SNMP Status Polling

This scenario measured the volume of SNMP traffic generated on one NNMi management server during the device status polling cycle.

This scenario measured utilization during status polling on ~11k polled interface objects.

- No traps were being received by NNMi during this time.
- All polling except SNMP status polling was turned off.
- Discovery was turned off during this period.

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	Specific Settings	Objects
Communication Configuration	Region: ICMP Settings: (1) ☐ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	
Discovery Configuration	**Rediscovery Interval** set to 10 days (to prevent rediscovery cycles during the scenario)	
Monitoring Configuration	Default Fault Monitoring: (1) ☐ **Enable ICMP Management Address Polling** (2) ☐ **Enable ICMP Fault Polling** (3) ☑ **Enable SNMP Interface Fault Polling** (4) ☐ **Enable Card Fault Polling** (5) ☐ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	1808 Interfaces
	Default Performance Monitoring: (1) ☐ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	
Incident Configuration	SNMP Traps: none	

Results: Averages of 38.733 packets and .052 Megabits per second were measured for this scenario.

Standalone System – Performance Polling

This scenario measured the volume of SNMP traffic generated on one NNMi management server during the device performance polling cycle.

Utilization was measured during performance polling for the same ~11k interface objects polled during status polling scenario.

- No traps were being received by NNMi during this time.
- All polling except SNMP performance polling was turned off.

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	Specific Settings	Objects
Communication Configuration	Region: ICMP Settings: (1) ☐ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	
Discovery Configuration	**Rediscovery Interval** set to 10 days (to prevent rediscovery cycles during the test)	
Monitoring Configuration	Default Fault Monitoring: (1) ☐ **Enable ICMP Management Address Polling**	

	(2) ☐ **Enable ICMP Fault Polling** (3) ☐ **Enable SNMP Interface Fault Polling** (4) ☐ **Enable Card Fault Polling** (5) ☐ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	
	Default Performance Monitoring: (1) ☑ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	1808 Interfaces
Incident Configuration	SNMP Traps: none	

Results: Averages of 159.397 packets and 0.44 Megabits per second were measured for this scenario.

Standalone System – ICMP Status Polling

This scenario measure the volume of ICMP traffic generated on one NNMi management server during the ICMP fault polling cycle.

Utilization was measured during ICMP polling on 1,808 polled addresses.

- Only ICMP traffic was measured—all other traffic was excluded.

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	Specific Settings	Objects
Communication Configuration	Region: ICMP Settings: (1) ☑ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	1808 addresses
Discovery Configuration	**Rediscovery Interval** set to 10 days (to prevent rediscovery cycles during the test)	
Monitoring Configuration	Default Fault Monitoring: (1) ☐ **Enable ICMP Management Address Polling** (2) ☑ **Enable ICMP Fault Polling** (3) ☐ **Enable SNMP Interface Fault Polling** (4) ☐ **Enable Card Fault Polling** (5) ☐ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	1808 addresses
	Default Performance Monitoring: (1) ☐ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	
Incident Configuration	SNMP Traps: none	

Results: Averages of 17.87 packets and .008 Megabits per second were measured for this scenario.

Standalone System – Traps

This scenario measured the volume of SNMP traffic generated on one NNMi management server under a steady-state trap load. A Cisco Link Down/Cisco Link Up trap was sent at a rate of 10 per second.

- Only Cisco Link Up/Cisco Link Down traps were sent from interface 1 (which was polled by NNMi), and were randomly sent from all 1,500 nodes.
- Cisco Link Up/Cisco Link Down traps caused the following NNMi actions that resulted in additional ICMP and SNMP traffic:
    - Rediscovery of each node that sent a link down trap
    - Immediate status poll of each interface that sent a trap
- Some trap de-duplication was occurring so not every trap caused the secondary NNMi actions.

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	**Specific Settings**	**Object**
Communication Configuration	Region: ICMP Settings: (1) ☐ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	
Discovery Configuration	**Rediscovery Interval** set to 10 days (to prevent rediscovery cycles during the test)	1500 nodes
Monitoring Configuration	Default Fault Monitoring: (1) ☐ **Enable ICMP Management Address Polling** (2) ☑ **Enable ICMP Fault Polling** (3) ☑ **Enable SNMP Interface Fault Polling** (4) ☐ **Enable Card Fault Polling** (5) ☐ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	1,808 IP addresses 11,000 Interfaces
	Default Performance Monitoring: (1) ☐ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	
Incident Configuration	SNMP Traps: (1) Cisco Link Up (2) Cisco Link Down	

Results: Averages of 220.63 packets and .506 Megabits per second were measured for this scenario.

Standalone System – Custom Polling

This scenario measured the volume of SNMP traffic generated on one NNMi management server during a Custom Polling cycle. Utilization was measured during custom polling of 11K interfaces (if%util was polled).

The NNMi management server was configured as described in the following table:

Configuration settings on the NNMi management server:

Configuration Workspace	Specific Settings	Object
Communication Configuration	Region: ICMP Settings: (1) ☐ **Enable ICMP Communication** (2) **ICMP Timeout** set to 5 seconds Default SNMPv1/v2 Community Strings: (1) Only one Community String configured	
Discovery Configuration	**Rediscovery Interval** set to 10 days so that no discovery would occur	1500 nodes
Monitoring Configuration	Default Fault Monitoring: (1) ☐ **Enable ICMP Management Address Polling** (2) ☐ **Enable ICMP Fault Polling** (3) ☐ **Enable SNMP Interface Fault Polling** (4) ☐ **Enable Card Fault Polling** (5) ☐ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	
	Default Performance Monitoring: (1) ☐ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	
Incident Configuration	SNMP Traps: none	
Custom Poller Configuration	☑ Enable Custom Poller One Custom Poller Collection defined for if%util	11000 interfaces

Results: Averages of 103.241 packets and .189 Megabits per second were measured for this scenario.

Global Network Management (GNM) Network Utilization

Configuration

The Global Network Management (GNM) feature of NNMi allows for central collection of several remote network management stations. The GNM scenario described in this section included one Global Manager and three Regional Managers.

Traffic Measured for the Global Network Management Scenarios:

- Traffic was measured on the Global Manager. The Global Manager was not responsible for discovering or monitoring any nodes. It was only responsible for displaying data received from the three Regional Managers.
- The Global Manager's database contained a total of 30,000 nodes.
- Three Regional Managers (each managing from 6,000 to 25,000 nodes) forwarded data to the Global Manager.
- Across the three Regional Managers was a total of 420,000 fault and performance polled interfaces, 400,000 node components (sometimes referred to as node health components) and 30,000 IP addresses.

GNM - Discovery Scenario

This scenario measured the volume of traffic generated during the period directly after configuring one Regional Manager to forward data to the Global Manager. This is the traffic bandwidth required to send all of the topology information for 25,000 nodes from the Regional Manager to the Global Manager. Initial discovery on the Regional Manager was completed. No other Regional Managers were connected to the Global Manager during this period.

Configuration Settings on the Global Manager (NNMi management server)

Configuration Workspace	Specific Settings	Objects
Global Network Management	Regional Manager Connections (one configured)	1 Regional Manager
Incident Configuration	SNMP Traps: none	

Results: Averages of 252.13 TCP packets and 2.48 Megabits per second were measured during this scenario.

GNM – Steady State SNMP Polling

This scenario measured the TCP traffic volume received by the Global Manager from the three Regional Managers.

During this scenario all of the polled objects mentioned in the Configuration section above were polled at the default interval of 5 minutes. The Global Manager had completed discovering (transferring) all the topology data from the three Regional Managers.

The three Regional Managers were configured as described in the following table:

Configuration Settings on the Three Regional Managers (NNMi management servers)

Configuration Workspace	**Specific Settings on Regional Managers (each with these settings)**	**Cumulative Object Count forwarded Global Manager**
Discovery Configuration	**Rediscovery Interval** on all Regional Managers was set to 24 hours	30,000 nodes
Monitoring Configuration	Default Fault Monitoring: (1) ☑ **Enable ICMP Management Address Polling** (2) ☑ **Enable ICMP Fault Polling** (3) ☑ **Enable SNMP Interface Fault Polling** (4) ☑ **Enable Card Fault Polling** (5) ☑ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	30,000 IP addresses -and- 420,000 Interfaces -and- 400,000 Node Components
	Default Performance Monitoring: (1) ☑ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	420,000 Interfaces
Incident Configuration	SNMP Traps: none	

Results: Averages of 412.26 TCP packets and 4.01 Megabits per second were measured in the lighter loading scenario.

GNM - Heavy Load Fault and Performance Polling

This scenario measured the TCP traffic volume received by the Global Manager from the three Regional Managers. During this scenario 80,000 of the 420,000 interfaces were fault and performance polled every minute. The remaining interfaces were fault and performance polled every 5 minutes.
Each of the Regional Manager configuration settings are described in the following tables. The only difference in the Regional Manager settings is the interval settings noted in red.

HEAVY TRAFFIC CONFIGURATION SETTINGS:
Configuration Settings on *One* of the Regional Managers (NNMi management server)

Configuration Workspace	**Specific Settings on Regional Managers (each with these settings)**	**Cumulative Object Count forwarded Global Manager**
Discovery Configuration	**Rediscovery Interval** was set to 24 hours	10,000 Nodes
Monitoring Configuration	Default Fault Monitoring: (1) **Enable ICMP Management Address Polling** (2) ☑ **Enable ICMP Fault Polling** (3) ☑ **Enable SNMP Interface Fault Polling** (4) ☑ **Enable Card Fault Polling** (5) ☑ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 1 minute	5,000 IP addresses -and- 80,000 Interfaces -and- 50,000 Node Components

	Default Performance Monitoring: (1) ☑ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 1 minute	80,000 Interfaces
Incident Configuration	SNMP Traps: none	

NORMAL TRAFFIC CONFIGURATION SETTINGS:

Configuration Settings on *Five* of the Regional Managers (NNMi management servers)

Configuration Workspace	**Specific Settings on Regional Managers (each with these settings)**	**Cumulative Object Count forwarded Global Manager**
Discovery Configuration	**Rediscovery Interval** on all Regional Managers was set to 24 hours	55,000 Nodes
Monitoring Configuration	Default Fault Monitoring: (1) **Enable ICMP Management Address Polling** (2) ☑ **Enable ICMP Fault Polling** (3) ☑ **Enable SNMP Interface Fault Polling** (4) ☑ **Enable Card Fault Polling** (5) ☑ **Enable Node Component Fault Polling** (6) **Fault Polling Interval** set to 5 minutes	25,000 IP addresses -and- 340,000 Interfaces -and- 450,000 Node Components
	Default Performance Monitoring: (1) ☑ Enable SNMP Interface Performance Polling (2) **Performance Polling Interval** set to 5 minutes	340,000 Interfaces
Incident Configuration	SNMP Traps: none	

Results: Averages of 602.46 TCP packets and 6.32 Megabits per second were measured during this scenario.

GNM - Resynchronization Scenario

This scenario measures the TCP traffic volume received by the Global Manager from the three Regional Managers during a resynchronization of the NNMi database. GNM resynchronization happens when all state and status is recalculated and sent from the Regional Managers to the Global Manager. A resynchronization happens after major events such as upgrading NNMi versions or after the Global Manager fails over. Network traffic was measured on the Global Manager during a resynchronization.

Configuration Settings on the Regional Managers (NNMi management servers)

Configuration Workspace	**Specific Settings on Regional Managers (each with these settings)**	**Cumulative Object Count forwarded Global Manager**
Discovery Configuration	**Rediscovery Interval** on all Regional Managers was set to 24 hours	55,000 Nodes
Monitoring Configuration	Default Fault Monitoring: (7) **Enable ICMP Management Address Polling** (8) ☑ **Enable ICMP Fault Polling** (9) ☑ **Enable SNMP Interface Fault Polling** (10) ☑ **Enable Card Fault Polling** (11) ☑ **Enable Node Component Fault Polling** (12) **Fault Polling Interval** set to 5 minutes	25,000 IP addresses -and- 340,000 Interfaces -and- 450,000 Node Components
	Default Performance Monitoring: (3) ☑ Enable SNMP Interface Performance Polling (4) **Performance Polling Interval** set to 5 minutes	340,000 Interfaces
Incident Configuration	SNMP Traps: none	

Averages of 1154.76 TCP packets and 11.474 Megabits per second were measured during this scenario.

NNMi Application Failover

Configuration

NNMi's Application Failover functionality replicates the Postgres database from the Active server to the Standby server on port 5432 and sends transaction logs and database .zip files on port 7810. The following scenarios measured the TCP traffic on these two ports.

For these scenarios, NNMi was scaled to the upper end of the Extra High single-system tier: 25,000 nodes, 2,000 polled objects (performance and fault polling), and 100,000 Custom Polled objects. All objects were polled at the default interval.

The NNMi Active Server configuration that applies to each scenario is described in the following table:

Configuration Settings on the Extra High Single-System Tier (NNMi management servers)

Configuration Workspace	**Specific Settings on Regional Managers (each with these settings)**	**Cumulative Object Count forwarded Global Manager**
Discovery Configuration	**Rediscovery Interval** on all Regional Managers was set to 24 hours	55,000 Nodes
Monitoring Configuration	Default Fault Monitoring: (13) **Enable ICMP Management Address Polling** (14) ☑ **Enable ICMP Fault Polling** (15) ☑ **Enable SNMP Interface Fault Polling** (16) ☑ **Enable Card Fault Polling** (17) ☑ **Enable Node Component Fault Polling** (18) **Fault Polling Interval** set to 5 minutes	25,000 IP addresses -and- 340,000 Interfaces -and- 450,000 Node Components
	Default Performance Monitoring: (5) ☑ Enable SNMP Interface Performance Polling (6) **Performance Polling Interval** set to 5 minutes	340,000 Interfaces
Incident Configuration	SNMP Traps: none	

Application Failover - Discovery

This scenario measured the TCP traffic volume received on the Application Failover ports during initial discovery on the Active server.

Results: Averages of 1330.81 packets and 41.2 Megabits per second were measured during this scenario

Application Failover - Steady State SNMP Polling

This scenario measured traffic after discovery had completed and all polling was set to the 5 minutes default interval. No resynchronization, heavy rediscovery, outages or other events that would cause heavy load on the system were occurring.

Results: Averages of 20.35 packets and 0.142 Megabits per second were measured during this scenario

Application Failover - Database Synchronization

Periodically, Application Failover does a full database synchronization from the Active server to the Standby server. Network traffic was measured during this period of increased traffic between the Active and Standby servers.

Results: Averages of 3595.60 packets and 127.156 Megabits per second were measured during this scenario

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