Hewlett Packard
Enterprise

HPE Network Node Manager iSPI Performance for Metrics Software

Software Version: 10.20
Windows® and Linux® operating systems

Online Help

Document Release Date: July 2016
Software Release Date: July 2016

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise 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. HPE 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 HPE 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 2009-2016 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe® is a trademark of Adobe Systems Incorporated.

Apple is a trademark of Apple Computer, Inc., registered in the U.S. and other countries.

AMD is a trademark of Advanced Micro Devices, Inc.

Google™ is a registered trademark of Google Inc.

Intel®, Intel® Itanium®, Intel® Xeon®, and Itanium® are trademarks of Intel Corporation in the U.S. and other countries.

Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries.

Internet Explorer, Lync, Microsoft, Windows, and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

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

Red Hat® Enterprise Linux Certified is a registered trademark of Red Hat, Inc. in the United States and other countries.

sFlow is a registered trademark of InMon Corp.

UNIX® is a registered trademark of The Open Group.

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

This product includes software developed by the Visigoth Software Society (http://www.visigoths.org/).

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/group/softwaresupport/search-result?keyword=.

This site requires an HP Passport account. If you do not have one, click the **Create an account** button on the HP Passport Sign in page.

Support

Visit the HPE Software Support web site at: **https://softwaresupport.hpe.com**

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

HPE Software 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 HPE 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 **https://softwaresupport.hpe.com** and click **Register**.

To find more information about access levels, go to:
https://softwaresupport.hpe.com/web/softwaresupport/access-levels

Contents

Introduction to Network Reporting

The **Network Performance Server** (**NPS**) extends NNMi's functionality by providing a platform for performance data storage, analysis and reporting. The NPS platform is used by the NNM iSPI Performance series of products, which includes the **NNM iSPI Performance for Metrics**, **NNM iSPI Performance for Traffic**, and **NNM iSPI Performance for QA**.

Through NPS, the iSPI Performance products provide a comprehensive set of customizable reports and graphs that support interactive/operational performance management and scheduled reporting. The reports and graphs provide actionable insight into the performance of the network.

At the core of the foundation provided by NPS is the column-based database management system (DBMS), which enables you to effectively store, access, and track network performance data.

Extension Packs provided by the NNM iSPI Performance products add rules and definitions to the NPS deployment to generate reports. These rules and definitions enable the business intelligence component—the BI Server—to generate insightful and actionable reports.

NPS Administration

NPS uses NNMi for all user authentication and logon via Single Sign-On. NNMi users automatically become NPS users the first time they access the NPS Report Menu.

NNMi user roles are applied in NPS. An NNMi administrator is automatically granted administrative privileges on NPS. These privileges include the ability to administer and manipulate content and settings on the BI server.

HPE does not recommend that NPS users make any changes to any settings on the BI Server because changes can disrupt the functionality and performance and render the BI Server unsupportable. The user must become familiar with the BI Server and accept responsibility for the outcome if BI Server settings are modified. The BI Server documents and tutorials are shipped and installed with NPS, and can be accessed from the BI Portal.
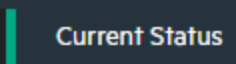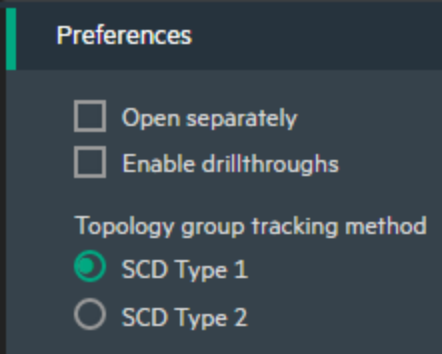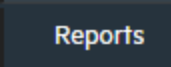
OEM license restrictions stipulate that only one Admin can log in to the BI Server at any one time.

Using the NPS Console

The iSPI Performance interface consists of a workspace navigation panel and a content pane.

Navigation Panel

The navigation panel provides a series of tabs that enable you to create reports. To open a tab, click ⌄ the down arrow. Clicking another tab automatically closes the previously used tab. The following tabs are provided.
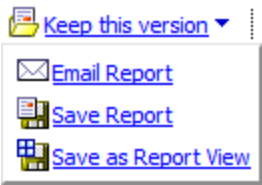
Tab	Description
Current Status	Tracks your current selection and actions and provides the Sign out link. This tab remains open.
Preferences ☐ Open separately ☐ Enable drillthroughs Topology group tracking method ⦿ SCD Type 1 ○ SCD Type 2	Options are: • Open separately • Enable drillthroughs Topology group tracking method • SCD Type 1 • SCD Type 2

Reports	Enables you to set launch reports and change various options of a report.
	
	 Available options are:
	
	 • Quicklaunch ReportViews:
	
	 Lists all the Report Views that are installed by default.
	
	   • Availability & Response for Nodes
	
	   • Capacity Planing on Node CPU%
	
	   • Executive for Interfaces
	
	   • Headline for Interfaces
	
	   • Headline for Nodes
	
	   • Most Changed Exceptions for Nodes
	
	   • Overview for Nodes
	
	   • Overview for Nodes - Business Hrs
	
	   • Response Baseline Sleeve for Nodes

Tab	Description
	• Throughput Calendar for Interfaces • Throughput Heat Chart for Interfaces • Top Interface Util% • Top Interface Util% Chart • Top Node CPU% • Top Node CPU% Chart • Top Problematic Items • Days To Threshold • Managed Inventory • Cross Domain Lists all cross-domain reports that are installed by default. For more information about cross-domain reports, see "Using Cross-Domain Reports" on page 66. • iSPI Metrics: Lists all reports provided by the NNM iSPI Performance for Metrics. • AtmPvc_Health (if this Extension Pack is installed) • Component_Health (if this Extension Pack is installed) • FrameRelayPvc_Health (if this Extension Pack is installed) • Interface_Health (if this Extension Pack is installed) • Self Monitoring (admin users only): Lists reports that help you monitor the health and performance for NPS. • Quicklaunch Diagnostic Reports • PerfSPI_Diagnostics • User Groups: • admin • globalops • guest • level1

Tab	Description
	- level2 - My Folders Additional report groups appear if you install other iSPIs or Custom Poller Extension Packs.
Report History	Lists the reports that have run. To clear the list, click the Clear button.
Topology Filter	Lists the filters that have been used. To clear the list, click the Clear button. To select new filters, click the Launch Topology Selector button. For a list of filters, see "Set Topology Filters" on page 45.
Time Control	Provides options you can set to create custom reports: - Data Available From / Available To - Detail - Hourly - Daily - Relative Start (with options) - Interval (with options) - Grain (with options) - TimeZone (with options) - Auto Refresh (with options) - Hour of Day / Day of Week (with options) To change any of the Hour of Day / Day of Week options, click the Reset button. To accept the selected options, click the Submit button.
BI Server	Provides access to: - Public Folders - My Folders - Administrator Log On - Portal - launches the BI Portal page - Query Studio - launches the custom report feature - Administration - launches the BI Server Administration page
Cross Launching	Provides a direct link to the NNM console.

Content Pane: Menu Items

Selection	Description
Keep this version ▼ ✉ Email Report Save Report Save as Report View	The **Keep this Version** menu item enable you to keep the report version that currently appears in the workspace, email it, save it as a report, or save it as a ReportView (if you are the Administrator or have permission for that option).
▷	Enables you to run the report selected.
View in HTML Format View in PDF Format View in Excel Options ▶	Enables you to view the report in various output formats: HTML, PDF, or Excel.

Content Pane: Links

Link	Description
Time	Opens the Time Control tab in the navigation panel.
Topology	Opens the Topology Filter.
Options	Opens the Report Options pane.
Show/Hide Bookmark	Displays the URL. Once opened, you can close it with Hide Bookmark.
Help	Opens the main page of the online help.

Tip: To log out of the console, click **Sign out** in the Current Status tab.

Location of Program and Runtime Files

NPS consists of static application software files and dynamic runtime files.

The default path for static application software is:

- *Windows:*

    C:\Program Files (x86)\HP\HP BTO Software\NNMPerformanceSPI

- *Linux:*

    /opt/OV/NNMPerformanceSPI

The default path for the dynamic runtime files is:

- *Windows:*

    C:\ProgramData\HP\HP BTO Software\NNMPerformanceSPI

- *Linux:*

  /var/opt/OV/NNMPerformanceSPI

The performance of NPS depends on the I/O performance of the disk subsystem where this directory resides. HPE strongly recommends that this data location be hosted on a high-performance file system. For specifications, see the support matrix.

The NNMPerformanceSPI directory contains the following folders:

- contentstore
- database
- nnmappfailover
- L10N
- logs
- PerfSPI_Diagnostics
- rconfig
- A directory for each installed Extension Pack
- Additional directories for any Custom Collect Report Groups created on NNMi

The folders are placed by the default NPS installation. Additional folders can be present if additional Extension Packs are present.

Data Retention Defaults

The default settings are the following:

Setting	Default Value (Days)	Table
Daily Data	800	Day
Hourly Data	70	Hour
Raw/Detailed Data	14	Raw

You can modify the default settings using the Configuration utility.

Single Sign-On Security

NPS enables you to use a security function known as Single Sign-on (SSO). SSO allows NPS to recognize the same user names and passwords the NNMi console recognizes. When SSO is enabled, a user who is already logged on to NNMi can move from NNMi to a report without having to log on again.

For SSO to operate, the following conditions must be met:

- The NNM system and the NPS system must be configured with fully qualified domain names (FQDN).
- The user must use the fully qualified domain (not "localhost" or the machine name without the domain) when launching the NNMi console.

If the user launches NNMi using an unqualified host name, an error message will appear in the Current Status tab of the navigation panel.

If the address, or the domain, of the NNMi management server or the NPS system changes after you install NPS, re-run the enablement script on NNMi. Re-running the enablement script configures NNMi with updated location information.

Direct access to NPS

You can access NPS directly, without going through NNMi, by pointing your browser at:

http://*<fully-Qualified-serverName>*:9300

When the login window opens, enter the same username and password you use to launch NNMi. After you are authenticated, the NPS home page opens.

You can also use the following methods:

- From the server where NPS is installed, run launchReport.ovpl.
- *Windows only*: Use the menu item in **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Report Menu**.

Monitor NPS Processes

NPS includes reports for admin users only to monitor NPS processes.

To view the available reports:

1. Log in to NPS.
2. On the NPS home page, click the **Reports** tab in the navigation panel.
3. Select **Self Monitoring**(admin users only) > **PerfSPI_Diagnostics** > **DiagnosticMetrics**.

Set Thresholds for Metrics in NNMi

Until you set the thresholds for the performance metrics, the reports do not track exception counts or exception rates.

NNMi does not set thresholds for performance metrics for you automatically. You must set the thresholds for the iSPI Performance metrics manually using the Thresholds Settings form in the NNMi console.

For information about how to set thresholds with the NNMi console, and for a listing of metrics, see "Configure Threshold Monitoring for Interfaces (HPE Network Node Manager iSPI Performance for Metrics Software)" in the *HPE Network Node Manager i Software Online Help for Administrators*.

Although you might want to see exception counts as soon as possible, you should wait until you are better acquainted with the way your network behaves. For example, if the thresholds are set wrong, NNMi might detect too many exceptions. Or, if you configure NNMi to generate incidents in response to threshold conditions, NNMi could generate too many incidents.

NNMi Baseline Settings

You can use configure NNMi and NPS for baseline monitoring in your network environment. For more information, see the "Configure Baseline Settings" topic in the *HPE Network Node Manager i Software Online Help for Administrators*.

Support for NNMi Application Failover

Application Failover for NNMi ensures redundancy. Failover allows a secondary NNMi server to take over immediately after the failure of a primary NNMi server. Failover relies on jboss clustering technology, ongoing file system synchronization, and a Java keystore file that must be copied from NNM to NPS. For details, see the *HPE Network Node Manager iSPI Performance for Metrics Installation Guide*.

NPS supports NNMi Application Failover, which is transparent to users. Aside from an interruption in service lasting about 15 minutes while failover is in progress, users are not aware that a failover took place, and the administrator is not required to perform any special tasks related to failover.

The ability of NPS to support Application Failover depends on files it retrieves from the primary server in the cluster. NPS retrieves these files during startup. As soon as NPS has the files, it begins monitoring the status of the primary server by checking for status changes every 5 minutes. If NPS detects a status change, the following events take place automatically:

- NPS determines which server is the new primary server.
- NPS redirects data collection to a shared directory on the new primary server.
- NPS begins collecting data (metrics and topology files) from the new primary server in the cluster.

> **Note:** After failing over to the standby server, all reportlet dashboards and bookmark links will stop working. You must regenerate all reportlet dashboards and bookmark links after failover.

Immediately after, NPS users are able to link from NPS to NNMi views on the new primary server, just as they could before failover.

NPS reports messages relating to failover (see "Messages Related to Failover" below).

Messages Related to Failover

The following INFO, WARN, and ERROR log file messages pertain to Application Failover.

Condition	Message	Comments
Failover not enabled	DEBUG: NNM Application Failover is not enabled.	Logged every 5 minutes if DEBUG logging is enabled
Failover enabled	INFO: NNM Application Failover is enabled. Checking if active node has changed.	Logged every 5 minutes
Failover enabled	INFO: NNM Hostname currently used by NPS is: $nnm_current	Logged every 5 minutes. Displays the current NNM hostname being used by NPS.
Failover	INFO:	Logged every 5 minutes. Displays the NNM hostname reported as

Condition	Message	Comments
enabled	$getNNMActiveScript reports active nnm node is: $nnm_active	the active node in the cluster.
Failover enabled but the active host has not changed.	INFO: Active NNM host is unchanged.	Logged every 5 minutes.
Failover enabled	INFO: Changing NPS configuration to use $nnm_active	A configuration update is imminent. This message is logged only when the active node is different from the node currently used by NPS.
Failover enabled and config files successfully updated	INFO: Configuration updated. Path to NNM is now $newpath	Logged after the configuration file is successfully updated in response to an Application Failover.
Failover enabled but the NNM hostname is not in the path syntax	WARN: NNM Application Failover is not supported on same-system NNM/NPS installations	If NPS is running on the same machine as NNN, the path does not contain the hostname.
Failover enabled but the NNM hostname is not in the path syntax	WARN: Cannot extract remote nnm hostname from value of PRSPI_NNMDIR ($cfg{PRSPI_ NNMDIR})..	The user might have specified the path in a non-conventional way; for example, by using a hard mount or drive mapping. If the path does not contain the remote NNM hostname, NPS is not able to detect Application Failover.
Failover enabled but the NNM hostname is not in the path syntax	WARN: Expected PRSPI_NNMDIR of form /net/hostname/var... or //hostname/PerfSpi	The user might have specified the path in a non-conventional way; for example, by using a hard mount or drive mapping. If the path does not contain the remote NNM hostname, NPS is not able to detect Application Failover.
Failover enabled but the nnm_ details.xml file is not available.	ERROR: nnm_ details.xml not found in $newpath	This message is logged if there is no copy of nnm_details.xml on the remote NNM machine. This condition occurs if the `nnmenableperfspi` script was not run on that machine, because running the script creates the `nnm_details.xml` file. This condition could also occur if the machine configurations are not identical. For example, if the pathnames to the shared drive are not identical, the new path will be incorrect when NPS substitutes the new hostname into the path. The configurations **must** be identical.

Condition	Message	Comments
Failover enabled but there is a file permissions error	ERROR: Unable to modify timestamp of $newpath/nnm_details.xml: $!.	NPS configuration will not change. A file permissions error is preventing the nnm_detail.xml file from being modified on the NNM server.
Failover enabled but unable to update the config file.	ERROR: Unable to write changes to config file. Will try again next time.	Failure to update the config file could be caused by another process that is locking the config file.

Configure NPS

If you install NPS on the NNMi management server, configuration takes place automatically during the installation. After installation, you can use the configuration utility to make changes.

If you install NPS on a dedicated server, you must use the Configuration utility to specify configuration details. At the end of the installation, before the NPS ETL services are started, the Configuration utility asks you necessary configuration details.

To launch the Configuration utility:

- Windows:

  Select **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Configuration Utility**
- Linux:

  Run runConfigurationGUI.ovpl

Configuration Utility

You can use the Configuration utility to specify the following values.

Configure the iSPI Performance service to log on as:

Windows

- When NPS is installed on a dedicated Windows server, you must select **Remote CIFS Share Access**. For detailed instructions, see the *Installing NPS* section of the *HPE Network Node Manager iSPI Performance for Metrics Interactive Installation Guide*.

  You must provide the same account information that was used with the nnmenableperfspi.ovpl script on the NNMi management server.
- When NPS is installed on the Windows NNMi management server, you must select **Local Path or NFS Share**.

Linux

- When NPS is installed on a dedicated Linux server, you can select **Remote CIFS Share Access** or **Local Path or NFS Share** depending on the configuration provided with the nnmenablementperfspi.ovpl script.

- When NPS is installed on the Linux NNMi management server, you must select **Local Path or NFS Share**.

Specify the path to iSPI Performance data files folder:

- If NPS is installed on the NNMi management server, a local path is displayed in this field. You *must not* change this value.

- If NPS is installed on a dedicated server, the `nnmenableperfspi.ovpl` script creates a shared space on the NNMi server.

   You must specify the location of the shared space in the correct format.

   > **Tip:** Details provided during the last run of the `nnmenableperfspi.ovpl` script are captured in the `nnmenableperfspi.txt` file, which is available in the following directory on the NNMi management server:
   >
   > *On Windows*
   >
   > `%nnmdatadir%\log`
   >
   > *On Linux*
   >
   > `/var/opt/OV/log`

Daily Data Retention Period (days)

The default setting is 800.

Hourly Data Retention Period (days)

The default setting is 70 days.

The maximum retention period is 400 days.

Raw/Detailed Data Retention Period (days)

The default setting is 14 days.

The maximum retention period is 400 days.

Service status:

Click **Stop** or **Start** to stop and start necessary NPS processes.

If you change any of the default settings, click **Apply**, click **Stop**, click **Start**, and then click **Exit**.

To retain the originally configured settings, click **Exit**.

Master Runtime Configuration File NNMPerformanceSPI.cfg

The NNMPerformanceSPI.cfg file controls the operation of the iSPI Performance for Metrics.

The file contains values written by the Configuration Utility, as well as many other options with their standard and recommended settings. Users should NOT modify the contents of this file directly. Doing so can affect the functionality and performance of NPS and render it unsupportable.

User Groups and Security Groups in NPS

Both NNMi and NPS restrict access to data based on user/UserGroup/SecurityGroup/node mappings.

User Groups and Security Groups are created by NNMi administrators. For each user, NPS filters the data queries for all reports and ad hoc queries, so that the user has access only to data for managed nodes whose Security Groups are associated with the User Groups to which the user belongs.

NPS data security restricts all data access according to the following policies:

- Filter the fact table queries for all report graphs and tables, including reports created using Query Studio.

- Filter the dimension data queries used in topology filter prompts or in Query Studio.

- Filter group member queries used in topology filter prompts or in Query Studio. NPS does not display data related to the groups, where none of the members have access to the report data.

The NNMi built-in User Groups "NNMi Administrators (admin)" and "NNMi Global Operators (globalops)" are automatically granted visibility to all performance data regardless of security group assignments. All users who are assigned to either the NNMi admin or globalops User Group have full access to all data in reports and queries.

Members of the "NNMi Administrators (admin)" User Group are granted full administrative privileges to the NPS BI Server, and have unfettered rights to save, move,delete content (including content created or owned by other users), and perform other administrative tasks.

Members of the "NNMi Global Operators (globalops)" User Group are granted read-only access to all data, including reports and queries saved by other users, but do not have administrative capabilities on the NPS BI Server and cannot modify the content of public folders or other User Group folders.

Example

Say that NPS data is divided among the following locations:

- Site A

- Site B

And each location includes the following local users:

- Operator 1: This user can access data for only Site A.

- Operator 2: This user can access data for only Site B.

The following global users can access data from all locations:

- Admin: The NNMi and NPS Administrator administers both locations and accesses NPS data from all locations.

- GlobalOperator: The user for both Site A and Site B must access NPS data in both locations.

Once the administrator (Admin in this example) creates these User Groups in NNMi, NPS ensures that access to NPS features and data is restricted according to the privileges defined for each User Group.

User Account	NPS User Group	NPS Access Privileges
Admin	NNMi and NPS Administrator	- Can access all NPS data irrespective of location - Can access all user group data in NPS - Can add, modify, or delete NPS users and user groups using NNMi console - Can access all report templates, schedule reports, and define global report settings
GlobalOperator	Global Operator	Can generate reports for all sites, including Site A and Site B

Operator 1	Level 1 Operator	Can generate reports for only Site A
Operator 2	Level 1 Operator	Can generate reports for only Site B

For more information on configuring User Groups, see "Configuring Security" in the *HPE Network Node Manager i Software Online Help for Administrators*.

Check the Validity of NNM iSPI Performance Configurations

The Configuration Checker enables you to verify the validity of the configuration file entries. It displays the following message if all the entries are valid:

`INFO: Configuration File Validated OK`

To launch the Configuration Checker:

- Windows:

  Select **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Configuration Checker**
- Linux:

  Run runCheckConfig.ovpl

Collect Diagnostic Data for Support

The Diagnostics Collector utility enables you to collect diagnostic data for Support. The data will be used to to diagnose the problem.

To launch the Diagnostics Collector:

- Windows:

  Click **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Diagnostics Collector**
- Linux:

  Run collectDiagnostics.ovpl

The Diagnostics Collector saves the output, in the form of a compressed tar.gz file, in the following location:

- Windows:

  %NPSDataDir%\collectDiag
- Linux:

  /var/opt/OV/HPNNMPerformanceSPI/collectDiag

Forward the tar.gz file from the collectDiag folder to Support.

Reset NPS

The reset utility enables you to delete all data from the database with the following options:

- Keep all archived data.
- Restore archived data for re-processing.
- Delete all archived metrics/topology.

- Delete all archived data and log files and restart all services.
- Delete all archived data and recreate the database with the default size, location, and options.

To launch the NPS Reset utility:

Windows:

Select **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Reset Utility**

Linux:

Run `resetSPI.ovpl`

The Reset Utility window displays the following options:

```
1. Delete all data from database, & keep all archived data

2. Delete all data from database, & restore archived data for re-processing

3. Delete all data from database, & delete all archived metrics/topology

4. Delete all data from database, delete all archived data and logfiles, & restart all
services

5. Delete and recreate entire database using default size, location, and options

Enter Option (1-5), or Q to quit:
```

Analyze Log Files

The Log File Analyzer enables you to view the following:

- A daily summary of errors, warnings, and other notifications produced by each process within each Extension Pack
- Timing data for selected processes belonging to each Extension Pack
- Timing data for generating reports for different Extension Packs

The Log File Analyzer analyzes the `prspi.log` and `PerfSPI_Diagnostic_Metrics.log` files.

To launch the Log File Analyzer:

Windows:

Select **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Log File Analyzer**

Linux:

Run `log_analyzer.ovpl`

Warnings, Errors, and Other Notifications in iSPI for Performance Log Files

Summary data for warnings, errors, and other notifications covers the previous two weeks. The last summary covers today. The summary for today is incomplete until midnight. Summary data indicates the following:

- Date
- Number of errors per process, if any

- Number of warnings per process, if any
- Number of processes with fatal errors, if any

Here are examples:

```
prspi.20080729.log:
WARN  - 240 in ETL.Component_Health
WARN  - 268 in ETL.Interface_Health
WARN  - 285 in MainServiceLoop
ERROR - 2 in MainServiceLoop
prspi.20090915.log:
WARN  - 27 in ETL.Interface_Health
WARN  - 714 in MainServiceLoop
FATAL - 1 in Utility.startBI
FATAL - 152 in ETL.Interface_Health
FATAL - 1 in Utility.installExtensionPack
```

A warning normally indicates a transient condition, usually a temporary mismatch, that will self-correct. If you see a warning message or a fatal error message, you can examine it in more detail by viewing the associated logfile in a text editor.

The last log in the list is today's partially complete log. This partially complete log does not have a date attached to it. The file name for today's log is `prspi.log`

View Timing Data

Under the summary of all warnings and errors, if you press **Enter**, you can continue to see the timing data.

The timing data shows the following:

- Total number of times a process executed over the previous two weeks
- Average execution time per process over the previous two weeks and standard deviation
- Maximum execution time per process over the previous two weeks
- Average number of records processed per execution
- Average number of records processed per second

Monitor Log Files

The log file viewer is Chainsaw 2.0. With Chainsaw, you can monitor DEBUG, INFO, WARN, ERROR, and FATAL messages as they reach the `prspi.log` file. The file contains every message generated since the previous midnight. The path to the file is:

- Windows:

  %NPSDataDir%`\logs`

- Linux:

  `/var/opt/OV/NNMPerformanceSPI/logs`

To launch the Log File monitor:

- Windows:

  Select **Start** > **All Programs** > **HP** > **NNM iSPI Performance** > **Log File Monitor**

- Linux:

  Run `runChainsaw.ovpl`

Use the Chainsaw Message Interface for the NNM iSPI Performance

The welcome page includes several tabs. To open the message interface, click the **NNMPerformanceSPI/logs/prspi.log** tab.

The Chainsaw message interface displays the following panes:

- Event pane
- Detail event pane
- Tree logger pane

The event pane changes constantly, displaying the most recent message in prspi.log. Additional information about that message appears in the detail event pane. When you select an older message in the event pane, the detail event pane refreshes with additional information about that message.

The detail event pane indicates the following:

- Level
- Logger
- Time
- Thread
- Message
- NDC [null]
- Class [not used]
- Method [not used]
- Line
- File
- Properties
- Throwable [not used]

The logger tree pane shows the **Root Logger** folder.

The ETL folder, which appears under the Root Logger folder, contains three logger folders, one for each standard Extension Pack. An additiional logger folder appears for any additional Extension Packs that are installed.

New messages generated by the logger folders appear once every few seconds in the event pane. To see messages for one logger only, type that logger name in the **Refine focus on** field; for example:

`Refine focus on:Interface_Health`

You can expand the event pane by closing the logger tree pane and the detail event pane.

Chainsaw version 2.0 features enable you to:

- Fully customize each tab.
- Control the frequency of updates.
- Specify your own rules for color-coding events (the defaults are yellow for WARN, red for ERROR).

- Apply simple filters, or advanced expression-based filters.

- Save memory by setting a cyclic-based view that looks at X events only.

For more information about Chainsaw, run the tutorial under the Help menu or go to
`http://logging.apache.org/chainsaw/quicktour.html`

Set Thresholds for Metrics in NNMi

Until you set the thresholds, you cannot track exception counts with the help of reports.

NNMi does not set thresholds for performance metrics for you automatically. You must set the thresholds for the NNM iSPI Performance for Metrics manually using the Thresholds Settings form in the NNMi console. You can set thresholds for the following metrics:

- Discard rate in

- Discard rate out

- Error rate in

- Error rate out

- Utilization in

- Utilization out

- CPU 5 second utilization

- CPU 1 minute utilization

- CPU 5 minute utilization

- Memory utilization

- Buffer utilization

- Buffer miss rate

- Buffer failure rate

For information about how to set thresholds using the NNMi console, see *Configure Threshold Monitoring for Interfaces (HPE Network Node Manager iSPI Performance for Metrics Software)* in the *HPE Network Node Manager i Software Online Help for Administrators*.

Although you might want to see exception counts as soon as possible, setting thresholds right away might not be feasible. You might want to wait until you are better acquainted with the way your network behaves. If you set thresholds prematurely, you might experience problems. For example, if the thresholds are wrong, NNMi might detect too many exceptions. Or, if you configure NNMi to generate incidents in response to threshold conditions, NNMi might generate too many incidents.

Definition of Threshold Terms

- **High Value (HV)**: Any value above this value is in the high range; any value in the high range triggers a High status.

- **High Value Rearm (HVR)**: Any value below High Value; indicates that the High status has ended.

- **High Trigger Count (HTC)**: The number of consecutive polls a high value must remain the same before the status changes to High. A trigger count of 1 is usually appropriate.

- **Low Value (LV)**: Any value below this value is in the low range; any value in the low range triggers a Low status.

- **Low Value Rearm (LVR)**: Any value above Low Value;indicates that the Low status has ended.

- **Low Trigger Count (LTC)**: The number of consecutive polls a low value must remain the same before the status changes to Low. A trigger count of 1 is usually appropriate.

- **Duration**: The time limit before which the baseline metric state cannot change to Abnormal Range or generate an incident. This time period specifies the minimum time within which the metric value must remain out of the configured Baseline Range.

- **Duration Window**: The window of time in which the Upper Baseline Limt or Lower Baseline Limit criteria must be met.

Custom Poller

NNMi's Custom Poller feature uses SNMP MIB variables to gather additional device information that is useful in managing your network. The Custom Poller enables you to:

- View the list of topology nodes for which additional information is being polled as well as the name of the associated Policy.

- View the polling results. The first time the results for a specified MIB variable are discovered, the results appear in a Polled Instance object. The Polled Instance object is updated whenever a change in State is detected and includes the most current polled value that caused the State to change.

- View incidents that are generated by Custom Poller for a MIB Poll Variable.

For example, as an administrator, you could have NNMi monitor the Status of COM (communication) ports on all of your Windows servers. You could also configure NNMi to send incidents when certain MIB values are detected. Or to map returned MIB values to States and include those States as part of the node Status calculations.

As an operator, the Custom Poller enables you to:

- View the list of topology nodes for which additional information is being polled as well as the name of the associated Policy.

- View the polling results. The first time the results for a specified MIB variable are discovered, the results appear in a Polled Instance object. The Polled Instance object is updated whenever a change in State is detected and includes the most current polled value that caused the State to change.

- View incidents that are generated by Custom Poller for a MIB Poll Variable.

You can use the data collected by Custom Poller collections to construct and view reports using the iSPI Performance for Metrics.

Overview of Custom Poller Collections and Their Associated Policies
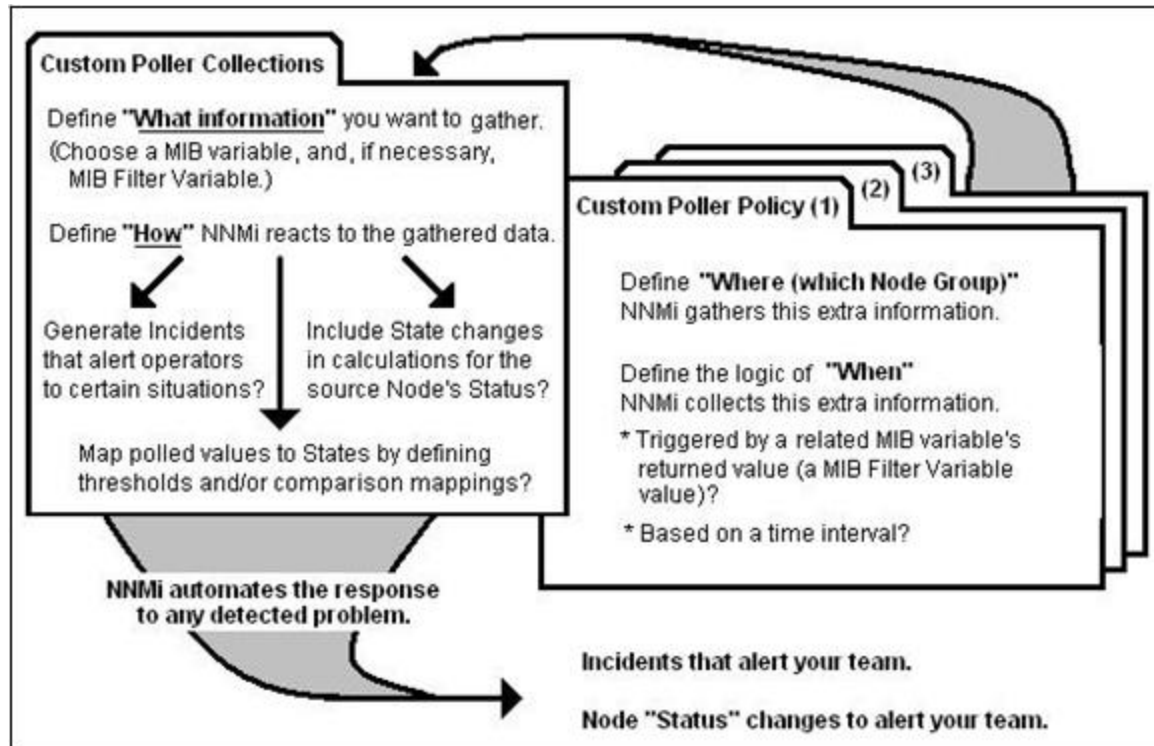
See the NNMi online help for more information on the Custom Poller feature.

Reports for Custom Poller Collections

You can configure the NNM iSPI Performance for Metrics to create reports based on the data collected by NNMi's Custom Poller. The Report Groups tab in the Custom Poller Configuration form enables you to generate reports from the metrics obtained from Custom Poller collections.

The NNM iSPI Performance for Metrics creates new Extension Packs to display reports generated from the data collected by Custom Poller collections.

To create Custom Poller collection and report groups, follow the instructions in the *Create Custom Polling Configurations* sections in *NNMi Help for Administrators*.

Guidelines to Creating a Custom Poller Extension Pack

- You can associate a Custom Poller Collection with only one report group.
- If you associate more than one Custom Poller Collection with the same Report Group, make sure the combination of Collections generates a meaningful report. Select Collections whose MIB Variables represent the same index value across Custom Poller Collections.
- Avoid using keyboard symbols while naming the Custom Poller collection and customPoll OID. The symbols `<,>,",',&,\` are replaced by a **!** character on the report.
- Reports are not created if the name of the Custom Poller includes the **#** symbol.
- You cannot rename a Custom Poller Extension Pack.

Remove Custom Poller Extension Packs

You can use command line tools to remove Custom Poller Extension Packs.

To remove Custom Poller Extension Packs:

> **Note:** In a distributed deployment of NPS, follow these steps only the system with the ETL Server role.
>
> In an environment with multiple ETL Servers, follow these steps only on the ETL Server where you installed the Custom Poller Extension Pack.
>
> For more information about these environments, see the *NNM iSPI Performance for Metrics Deployment Reference*.

1. Log in with administrative or root privileges to the system where NPS is installed.
2. To view the list of Extension Packs, run:

   `uninstallExtensionPack.ovpl`

3. To delete a specific Extension Pack, run:

   `uninstallExtensionPack.ovpl -p <extension_pack_name>`

You cannot recover the data associated with an Extension Pack if you inadvertently delete it. Make sure to type the correct Extension Pack name with the command. Do not delete the following default Extension Packs unless you want to remove iSPI Performance for Metrics reports: Interface_Health, Component_Health. If you delete the PerfSPI_Diagnostics Extension Pack, you will not be able to monitor the health of NPS processes.

Modify Topology Attribute Values and Group Memberships from NNMi

The topology attributes and groups help you filter the data and build the report with data polled from the nodes and interfaces you select. NPS retrieves the topology attribute details from NNMi and other iSPIs. If you make modifications to topology attributes and group memberships from NNMi, the changes are reflected in NPS within about an hour.

Changes are not applied retroactively to the existing data. Changes are always forward-looking. NPS preserves the historical accuracy of the data as it was when it was collected.

For example:

- If a new node is added to a group, it is only added for the new data collected going forward. NPS cannot assume that the node was a member of the group in the past.
- If an interface attribute changes—for example, if speed is changed from an old to a new value—the change is applied going forward. NPS does not assume that the new speed applies to the historical data.

If you delete or change a member of a topology attribute, the old member continues to appear in the list of topology filters in NPS. The old attribute member continues to appear in the list of topology filters. If you select the old member, the data collected before making the change gets sampled in the report.

Example

Say you modify node group name A to node group name B in the NNMi console. The new node name will appear in NPS after an hour or so. Later, the old node group name will continue to appear in the list of node groups in the Topology Filters section. But, if you select A, you will not be able to retrieve data samples associated with A that are polled after the node group name was changed to B.

Suppose you had an interface group called Oracle Server Ports that identifies switch interfaces assigned to your Oracle Server cluster, but now you have a new server cluster available online. You decide to transfer

your Oracle applications to the new one and re-purpose the old servers to host the Microsoft Exchange Server.

You therefore rename Oracle Server Ports to Exchange Server Ports and create a new interface group for Oracle Server Ports identifying the ports associated with the new servers. If you run a report at the end of the month to see the health of your Oracle Server Ports and your new Exchange Server Ports, NPS will define an Oracle Server Ports group to include the data collected for the original ports up to the time you made the change and for the new ports after that.

Likewise, NPS will define the Exchange Server Ports group to only include the ports from the time the group was renamed. It will not include the ports while they were still being used for Oracle.

Set Polling Characteristics in NNMi

Performance polling impacts NNMi and your network. If you enable additional performance polling, NNMi must allocate a portion of its resources to handle the additional polling. Your network must accommodate an increase in SNMP traffic. To use the system and network resources wisely, limit performance polling to the devices you require the most.

By default, NNMi enables performance polling for the following Node Groups:

- Network Infrastructure Devices
- Routers

The following are some of the ways you can modify the performance polling defaults that will be in place right after the iSPI is installed:

- Set thresholds for a node group you are already polling for performance data (no thresholds are set by default).
- Modify the frequency of performance polling.
- Add nodes to the Important Nodes group, and enable performance polling for this group.
- Enable performance polling for the Microsoft Windows Systems group.

All these processes involve monitoring configuration tasks. Monitoring configuration tasks are related to polling.

For information about any of these processes, see the *NNMi Help for Administrators*.

Using Reports

NPS generates reports on operational metrics that are collected by different NNM iSPIs and stored into the NPS database.

You can perform the following tasks with NPS reports:

- View the operational performance of metrics over a period of time.
- Modify the default settings of a report to view additional attributes related to the monitored element.
- Customize dashboards with predefined reportlets.
- Customize the schedule and format for report delivery.
- Create customized reports using Query Studio.

Avoid using Cognos BI Report Studio to create or customize NNM iSPI Performance reports. HPE does not provide support for such reports.

Report Types

The NPS generates different types of reports with the help of the data available in its repository (column-based DBMS) and the rules defined in extension packs. Reports help you view the operational performance of metrics over a period of time. The NPS also enables you to modify the default settings of a report to view additional attributes related to the monitored element.

Calendar Report

The Calendar report uses a traditional, calendar-style layout to show hourly statistics of two metrics. When you open a report, you can view the statistics of two default metrics. You can use the options link to view the report for non-default metrics.

Chart Details

The Chart Detail report represents the actual values of two metrics at every display grain over a period of time. You can use the options link to view the report for non-default metrics.

Heat Chart

The Heat Chart (or the normalized Heat Chart) report represents the *normalized* values of a metric, calculated at every hour of the day, in a tabular format. The NPS uses color codes to indicate the performance characteristic of the metric in this report.

Dashboard Reports

The Dashboard Report provides an overview of the network elements available on your network by using a combination of gauges, heat charts, and ordered Top 5 lists. The focus of the report is element-based threshold breaches, which are classed as exceptions. Threshold breaches can point to a network-related problem or may indicate that your thresholds within NNMi need to be changed.

Headline Reports

The Headline Report provides a broad view of performance by presenting eight different graphs on a single page. This report enables you to view every aspect of performance at once. Use the graphs to view trends, spot-isolated spikes, exceptions, and verify that performance is meeting expectations.

Top N Reports

The Top N report displays the top ten values of the metric over the selected time range.

Top N Chart Reports

The Top N report displays the top ten values of the metric over the selected time range. In addition, the report enables you to view the graphical representation of the top ten values.

Top N Frequency Distribution Reports

The Top N Frequency Distribution report ranks network elements by the metrics you select and shows the distribution of metric values for each collected sample during the selected time range.

You can use this report as an alternative to the Top N report. In addition to providing all the details that are available on the Top N report, this report helps you see how much time the metric value has spent in each bucket.

Top N Table Reports

This report ranks network elements by the metrics you select. Unlike the Top N report, this report does not show any bar charts or time series graphs. In a large environment, NPS can generate the Top N Table report faster than it can generate the Top N report. If you want to view Top N elements in the least possible time, choose the Top N Table report instead of the Top N report.

Threshold Sleeve Reports

The Threshold Sleeve report analyzes how close the metric value is to the threshold levels. In other words, this report enables you to analyze when the metric will breach the threshold.

Baseline Sleeve Reports

The Baseline Sleeve report shows the baseline sleeve determined by NPS for the selected time range.For more information about baselines, see "Using Baselines" on page 69.

Peak Period Reports

The Peak Period report ranks the utilization and performance of an element during the busiest time of the selected time range.

Managed Inventory Reports

The Managed Inventory report enables you to view the number of instances of each topology attribute seen in the collected data within the selected time range. For example, if the report is run for "Last 1 Hour," you can see the number of unique values for each attribute seen in the data collected over the last hour. This report can be useful for checking the number of elements polled over any given time range.

Most Changed Reports

The Most Changed report compares performance the elements for two different (consecutive) time periods and ranks them by the amount of change. The sort order is most-changed to least-changed.

Launch Reports

You can launch reports created by NPS from the NNMi console, the NPS console, or the iSPI Performance BI Server Portal.

To launch a report from the NPS console:

1. Launch the NPS console.

   To launch the NPS console, click **Actions > HPE NNM iSPI Performance > Reporting - Report Menu** in the NNMi console.

   Alternatively, to launch the NPS console directly without logging on to the NNMi console, launch the following URL:

   http://*<NPS_system>*:9300

   In this instance, *<NPS_system>* is the fully qualified domain name of the NPS system.

2. In the NPS console, select **Reports** in the navigation panel.

   NPS presents the default, ready-to-use Self Diagnostic Extension Pack (to generate reports on the performance of NPS itself). Depending on the number of the NNM iSPI products or Extension Packs available on the NPS system, you could see many other sets of reports.

3. Select a report type; for example, Chart Detail or Heat Chart.

4. Modify the default report settings, if required.

To launch a report from the NNMi console with a context:

1. Launch the NPS console:

   - *From an Incident view.* Right-click an incident, and then click **HPE NNM iSPI Performance > Reporting - Report Menu**.

   - *From a Topology Map.* Right-click an object, and then click **HPE NNM iSPI Performance > Reporting - Report Menu**.

   - *From an Inventory view.* Right-click an inventory object, and then click **HPE NNM iSPI Performance > Reporting - Report Menu**.

   The NPS console opens with the specific context (in other words, the topology attribute and Extension Pack are already selected).

2. In the NPS console, select the report type from the navigation pane. The report opens with the preselected topology attribute.

3. To navigate back to the NNMi console:

   - In the NPS home page navigation panel, click the **Cross Launching** tab.

   - Click **NNMi Console**.

The NNMi Console link enables you to go back to NNMi workspaces based on the topology objects currently selected in NPS.

After logging off from the NNMi console, if you want to log on again as a different user, you must clear browser cookies before launching any NPS reports.

Launch Reportlets

You can launch reportlets created by NPS or NNM iSPI products from NNMi, the NPS home page, or the iSPI Performance BI Server Portal.

To launch a reportlet from the NPS home page:

1. Log in to NPS using either the NNMi console or the NPS home page.
2. On the NPS home page, select an Extension Pack in the navigation panel, and then expand the **Reportlets** folder.
3. Select a reportlet type (for example, Exceptions Pie Chart or Simple Bar and Line Chart).
4. Modify the default reportlet settings, if required.

Reportlet Controls Toolbar

The Reportlet page enables you to modify the reportlet properties using the following tools in the Reportlet Controls toolbar.

Reportlet Control	Description
	Displays or hides the reportlet control toolbar
	Displays or hides the following information: - Start date for the data displayed - Start time for the data displayed - Display grain for the data displayed
	Displays or hides the Report Options panel.
	Displays or hides the Topology Filter panel.
	Displays a bookmark to the reportlet in a new panel.
?	Launches the Online Help topic for the reportlet.

Change Reportlet Settings

You can change the default settings like metrics, grouping by attribute, and topology attribute.

To change the reportlet settings, follow these steps:

1. Launch the reportlet from the NPS navigation panel.
2. Click **Show Toolbar**.

3. To change reportlet options, click ⚙️**Toggle Options**. The Report Options panel opens.

   In the Report Options panel, select metrics, grouping by options, and top/bottom N values.

4. To use topology filters, click 🔲. The Topology Filter panel opens.

   In the Topology Filter panel, select topology attributes of your choice.

NPS preserves your choice of metrics for a reportlet. When you come back to the reportlet in the future, the metrics you selected the last time are automatically used as the default metrics.

Launch Related Reports with Contextual Navigation

For easy access to related and relevant reports, NPS provides links, or contextual navigation points, on the report content to launch reports to focus on a particular element or time span. By clicking a hyperlink on the report, you can navigate to related reports with the context of the original report.

The table lists contextual navigation points for each report type.

| Report Type	Contextual Navigation Point
Calendar	**X-axis**: Opens report links for reports with the same Date/time context and the same primary metric.
Chart Detail	**X-axis**: Opens report links for reports with the same Date/time context and the same primary metric.
Heat Chart	**Hour of Day Row Header**: Opens report links for reports with the same Hour of Day filter of the selected row.
	Day of Month Header: Opens report links for reports with the same Date/Time filter of the selected row.
	Data cell value: Opens report links for reports with the same Hour of Day filter of the selected row and same Date/Time filter of the selected date.
Headline	**X-axis**: Opens report links for reports with the same Date/Time context and same metric.
Dashboard	**Interface name from the Top 10 utilization, Response Time, or Discard charts**: Opens report links for reports with the same topology filter to selected interface and the same metric.
	Needles of the Exception Count gauges: Opens report links for reports with the same metric.
Top N	**Element name** (for example, interface name or node name): Opens Top N report for the selected element with the same topology filter.
Most Changed	**Element name** (for example, interface name or node name): Opens Most Changed report for the selected element with the same topology filter.

You can launch a report from the same Extension Pack tab, or from a different one. If you launch a report from a different Extension Pack tab, only the relevant topology filters are preserved. Because metrics vary between Extension Packs, the selected metric cannot be transferred to the target report.
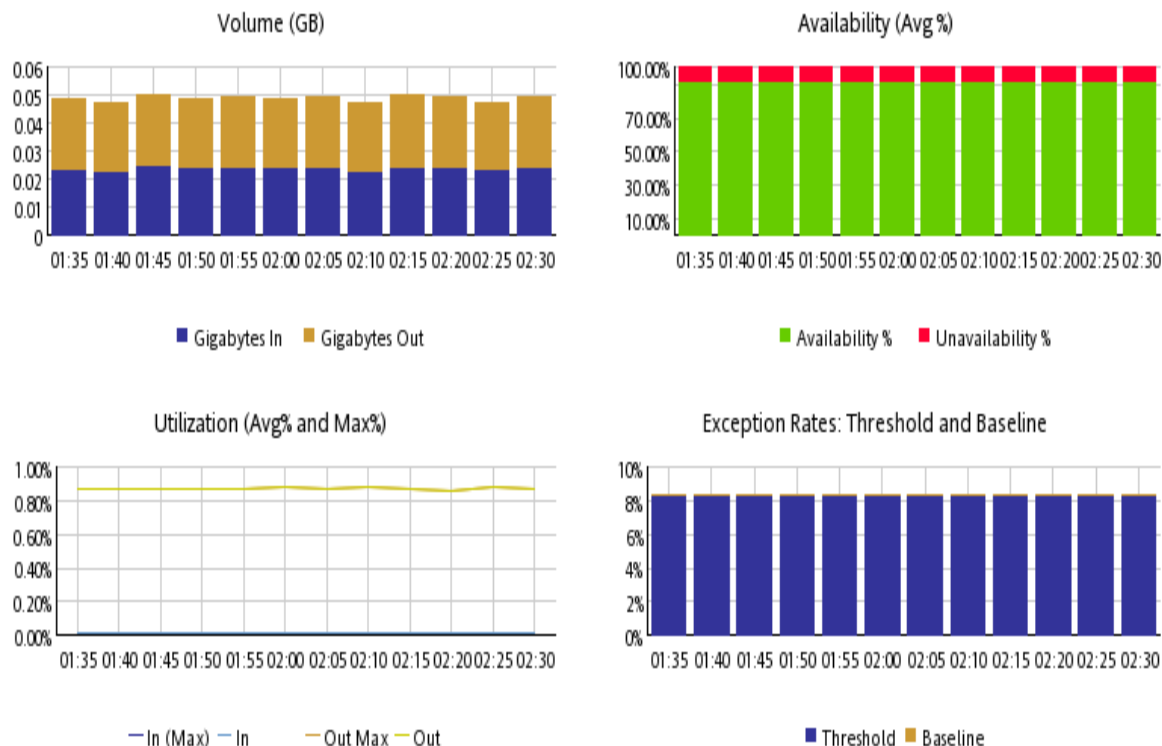
Enable Contextual Navigation

You can drill through reports to see more detail on elements you choose.

The Preferences tab of the navigation panel provides the option to automatically enable drillthrough on reports you plan to run.
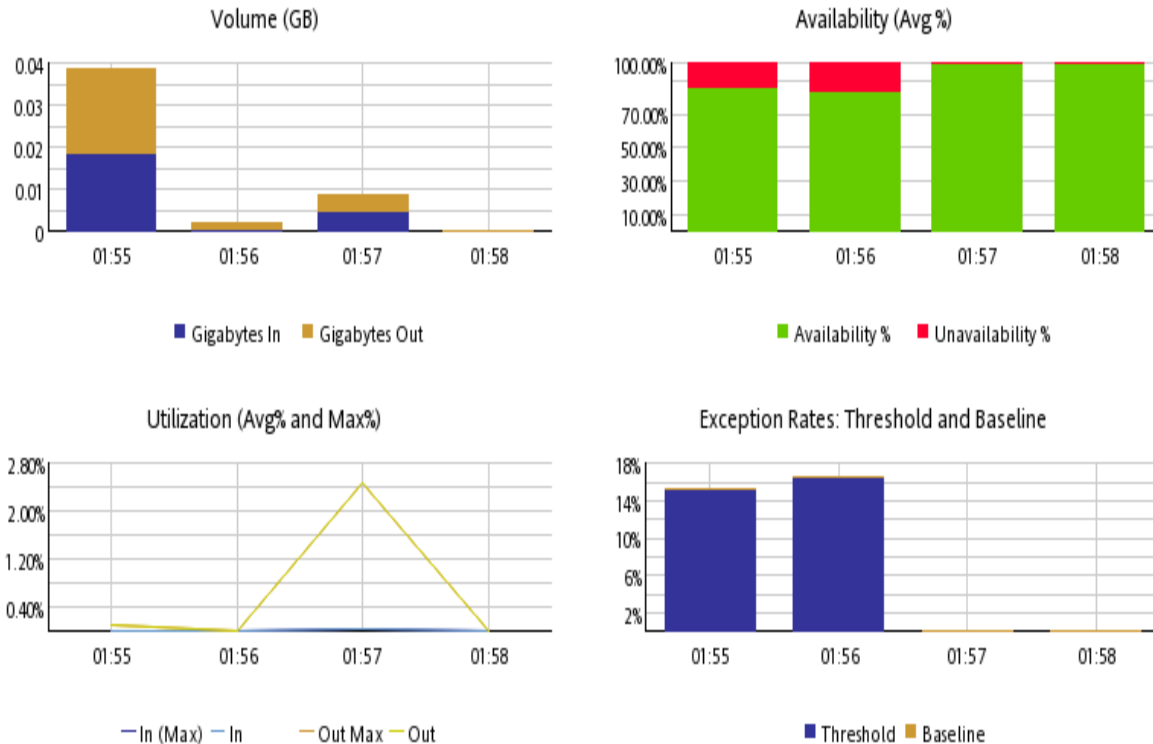
Preferences

☐ Open separately
☑ Enable drillthroughs
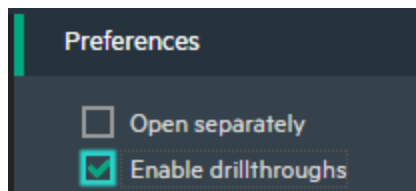
Or you can click the ⊞ drillthrough icon for a report you already ran.

When drillthrough is enabled, you can click an area of interest within the report to get more details. For example, say that you clicked inside the following Volume graph at time 01:55.

The following detail would appear for that time period.

Volume (GB)

Availability (Avg %)

Utilization (Avg% and Max%)

Exception Rates: Threshold and Baseline

Change Default Settings Before Launching Reports

You can change default settings like metrics, grouping by option, and topology filter before you launch the report in the NPS console.

To change report options (like metrics and grouping by options), right-click the report in the navigation pane, and then select **Run with Options**. The Report Options panel opens in the console.You can make necessary changes, click **Confirm Selection**, and then click **Finish** to launch the report. For more information, see "Change Report Options" on the next page.

To apply topology filter (or to change the existing topology filter), right-click the report in the navigation pane, and then select **Topology**. The Topology Filter panel opens in the console.You can make necessary changes, and then click **Finish** to launch the report. For more information, see "Set Topology Filters" on page 45.

To change the time setting, go to the Time Control panel in the navigation pane, select the time span, start time, display grain, and related settings, and then click **Submit**. For more information, see "Change Time Settings with the Time Control Tab" on page 49.

Change Default Settings

NPS generates reports with a set of parameters, which are set when you launch a report from the NPS home page or the NNMi console. However, you can use the NPS console to modify the default settings. Table: Changing Default Settings presents all the parameters that you can change on a report.

Table: Changing Default Settings

Parameter	How to Change
Metric	Reports and reportlets display graphical and statistical representation of different operational characteristics of the managed network. These characteristics are represented with **metrics** (parameters that indicate the health and operational performance of the network). Although every report and reportlet opens with the default metric set, the Report Options panel helps you generate the report or reportlet with the metrics of your choice. **To change report options:** 1. Launch a report. 2. Click **Options**. The Report Options panel opens.
Time controls	Time controls help you set the time range for the report or reportlet you want to view. The NPS database stores the data collected from different nodes over a period of time. Using the time controls parameters, you can view a report or reportlet created for the range of time of your choice. You can also select the display grain of a report or reportlet. When a report or reportlet is rendered, the time range is always shown in the Filter Summary section at the top of the report or reportlet.
Topology filter	The topology filters enable you to create a report that is filtered, or constrained in scope, by specific topology attributes. By default, when you launch a report or reportlet for the first time from the NPS home page, NPS generates the report or reportlet that is not filtered. However, you can select a specific node or a set of nodes from the NNMi console (from the Topology Maps workspace) and then launch the NPS home page to view reports or reportlets only on the selected nodes. The topology filters extend this capability by offering a range of filters on the NPS home page that help you select the data to be included or excluded prior to viewing the report. When a report or reportlet is rendered, the selected filters are always shown in the Filter Summary section at the top of the report or reportlet.

Change the Default Settings of Reportlets

Reportlets display quick impressions of larger reports using the NNM iSPI Performance for Metrics dashboards.

NPS generates reportlets with a set of parameters, which are set when you launch a reportlet from the NPS console.

To select a new metric for the reportlet, click **Show Toolbar** and **Toggle Options**.

To apply topology filter, click .

Change Report Options

The NPS console enables you to change report options in multiple ways. You can click the Options link on the report or right-click the report in the navigation pane and run the report with prompts.

To change report options:

1. Launch the report.
2. Click **Options** on the report. The Report Options pane appears.

Alternatively:

Right-click the report in the navigation pane, and then click **Run with Options**. The report opens with the Report Options pane.

After the Report Option pane appears, you can modify report options such as the metric and Grouping By parameters. The Report Options pane presents a list of all metrics provided by the Extension Pack in a well-structured table from which you can select a metric of your choice and build a report with it.

Report Type

Calendar

To modify report options, first open the Report Options pane by clicking the **Options** menu on the report.

To select a metric:

1. Right-click the metric.

   > **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
   >
   > You can also use **tag**[1]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select only one metric.
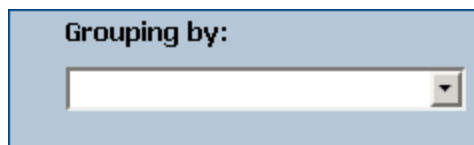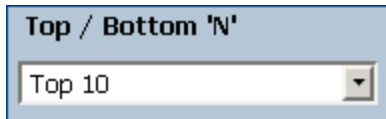
   > **Tip:** Click **Show Description** to see the metric definition.

Chart Detail

To modify report options, first open the Report Options pane by clicking the **Options** menu on the report.

Metrics

To select a metric:

1. Right-click the metric.

   > **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
   >
   > You can also use **tag**[2]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select up to six metrics.

   > **Tip:** Click **Show Description** to see the metric definition.

To remove a metric from an axis:

---

[1]A tag is a keyword that represents an aggregation type or a domain-specific group.
[2]A tag is a keyword that represents an aggregation type or a domain-specific group.

Right-click the metric on the axis, and then click **Delete**. The metric disappears from the axis.

To change the order of the metric:

Right-click the metric on the axis, and then click **Move Up** or **Move Down**. The metric changes its position on the axis and, subsequently, changes its position on the report when the report is launched.

To modify the report layout:

Chart or Table
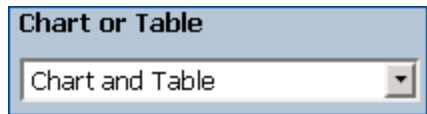
Chart and Table ▼

This drop-down list enables you to choose to view a chart or table or both on the report. Select Chart or Table or Chart and Table depending on your requirement.

Heat Chart

To modify report options, first open the Report Options pane by clicking the **Options** menu on the report.

To select a metric:

1. Right-click the metric.

    > **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
    >
    > You can also use **tag**[1]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric. On the axis, you can select only one metric.

    > **Tip:** Click **Show Description** to see the metric definition.

Dashboard

To modify report options, first open the Report Options pane by clicking the **Options** menu on the report.

Grouping By:

To change the Grouping By option:

Grouping by:

▼

In the Grouping By list, select an attribute to define the grouping mechanism for the tabular data on the report.

Top N

To modify report options, first open the Report Options pane by clicking the **Options** menu on the report.

Top/Bottom N:

---

[1]A tag is a keyword that represents an aggregation type or a domain-specific group.

The Top/Bottom N list helps you specify the number of ranks that you want to see on this report.

To select data points on the report:

Top / Bottom 'N'

Top 10

In the Top/Bottom N list, select a value of your choice.

To change the Grouping By option:

Grouping by:

In the Grouping By list, select an attribute to define the grouping mechanism for the tabular data on the report.

To select more than one attribute, click **Add New Grouping**. A new Grouping By box with the list of attributes appears.

To remove a grouping attribute, click **Remove Grouping**.

Metrics:

To select a metric:

1. Right-click the metric.

   > **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
   >
   > You can also use tag[1]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select up to six metrics.

   > **Tip:** Click **Show Description** to see the metric definition.

You can limit the number of rows on the report by specifying a range of values in the From and To boxes on each axis. After typing a value in the From or To box, you must press **Enter**. For the percentage type of metrics, specify values in decimal fractions. For example, to specify 9%, type 0.09.

To remove a metric from an axis:

Right-click the metric on the axis, and then click **Delete**. The metric disappears from the axis.

To change the order of the metric:

Right-click the metric on the axis, and then click **Move Up** or **Move Down**. The metric changes its position on the axis and, subsequently, changes its position on the report when the report is launched.

---

[1]A tag is a keyword that represents an aggregation type or a domain-specific group.

Note: Right-clicking a metric on the Y1 or Y2 axis does not bring up any context menus when no other metrics are present on the axis.

Display Time Series Chart:
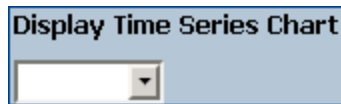
This drop-down list enables you to choose to view a time series line graph of .the Y1 metric (only the first metric on the Y1 axis) for the top or bottom N elements.
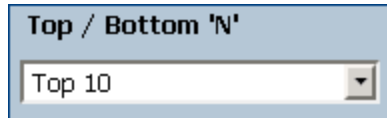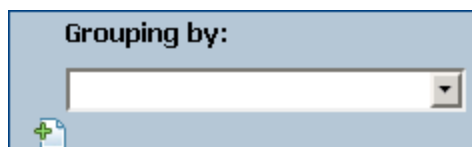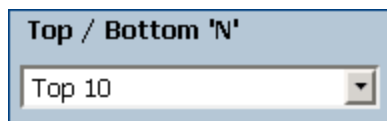
Display Table Extras:
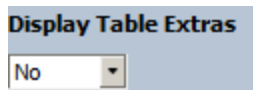
This drop-down list enables you to choose to view additional bar charts of .the Y1 metric for the top or bottom N elements.

Top N Chart

Top/Bottom N:

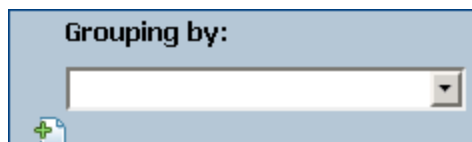The Top/Bottom N list helps you specify the number of ranks that you want to see on this report.

To select data points on the report:

In the Top/Bottom N list, select a value of your choice.

To change the Grouping By option:

In the Grouping By list, select an attribute to define the grouping mechanism for the tabular data on the report.

To select more than one attribute, click **Add New Grouping**. A new Grouping By box with the list of attributes appears.

To remove a grouping attribute, click **Remove Grouping**.

Metrics:

To select a metric:

1. Right-click the metric.

    Tip: You can use the search feature to find the metric of your choice. To search a metric, type the

metric name in the Search box.

You can also use **tag**[1]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select up to six metrics.

Tip: Click **Show Description** to see the metric definition.

You can limit the number of rows on the report by specifying a range of values in the From and To boxes on each axis. After typing a value in the From or To box, you must press **Enter**. For the percentage type of metrics, specify values in decimal fractions. For example, to specify 9%, type 0.09.

To remove a metric from an axis:

Right-click the metric on the axis, and then click **Delete**. The metric disappears from the axis.

To change the order of the metric:

Right-click the metric on the axis, and then click **Move Up** or **Move Down**. The metric changes its position on the axis and, subsequently, changes its position on the report when the report is launched.

Note: Right-clicking a metric on the Y1 or Y2 axis does not bring up any context menus when no other metrics are present on the axis.

Top N Table

Top/Bottom N:

The Top/Bottom N list helps you specify the number of ranks that you want to see on this report.

To select data points on the report:

Top / Bottom 'N'

Top 10

In the Top/Bottom N list, select a value of your choice.

To change the Grouping By option:

Grouping by:

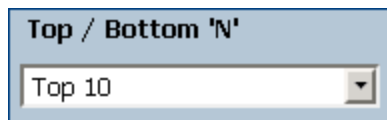In the Grouping By list, select an attribute to define the grouping mechanism for the tabular data on the report.

To select more than one attribute, click **Add New Grouping**. A new Grouping By box with the list of attributes appears.

[1]A tag is a keyword that represents an aggregation type or a domain-specific group.

To remove a grouping attribute, click ✕**Remove Grouping**.

Metrics:

To select a metric:

1. Right-click the metric.

> **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
>
> You can also use **tag**[1]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select up to six metrics.

> **Tip:** Click **Show Description** to see the metric definition.

You can limit the number of rows on the report by specifying a range of values in the From and To boxes on each axis. After typing a value in the From or To box, you must press **Enter**. For the percentage type of metrics, specify values in decimal fractions. For example, to specify 9%, type 0.09.

To remove a metric from an axis:

Right-click the metric on the axis, and then click **Delete**. The metric disappears from the axis.

To change the order of the metric:

Right-click the metric on the axis, and then click **Move Up** or **Move Down**. The metric changes its position on the axis and, subsequently, changes its position on the report when the report is launched.

> **Note:** Right-clicking a metric on the Y1 or Y2 axis does not bring up any context menus when no other metrics are present on the axis.

Top N Frequency Distribution

Top/Bottom N:

The Top/Bottom N list helps you specify the number of ranks that you want to see on this report.
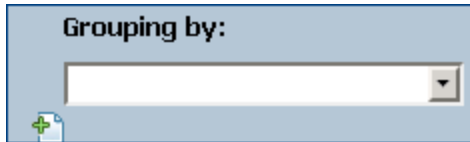
To select data points on the report:

In the Top/Bottom N list, select a value of your choice.

To change the Grouping By option:

---

[1]A tag is a keyword that represents an aggregation type or a domain-specific group.

In the Grouping By list, select an attribute to define the grouping mechanism for the tabular data on the report.

To select more than one attribute, click **Add New Grouping**. A new Grouping By box with the list of attributes appears.

To remove a grouping attribute, click **Remove Grouping**.

Metrics:

To select a metric:

1. Right-click the metric.

   > **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
   >
   > You can also use **tag**[1]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select up to six metrics.

   > **Tip:** Click **Show Description** to see the metric definition.

You can limit the number of rows on the report by specifying a range of values in the From and To boxes on each axis. After typing a value in the From or To box, you must press **Enter**. For the percentage type of metrics, specify values in decimal fractions. For example, to specify 9%, type 0.09.

To remove a metric from an axis:

Right-click the metric on the axis, and then click **Delete**. The metric disappears from the axis.

To change the order of the metric:

Right-click the metric on the axis, and then click **Move Up** or **Move Down**. The metric changes its position on the axis and, subsequently, changes its position on the report when the report is launched.

> **Note:** Right-clicking a metric on the Y1 or Y2 axis does not bring up any context menus when no other metrics are present on the axis.

By default, the Y2 series of metrics is disabled on this report. To use the Y2 series, select the Show Y2 Series check box.
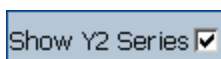
[1]A tag is a keyword that represents an aggregation type or a domain-specific group.

Bucket values:

The Top N Frequency Distribution report uses **buckets** to indicate different ranges of metric values.

About Buckets

A bucket represents a range of values. The Top N Frequency Distribution report enables you to define 12 different buckets (or ranges of values) for a metric series—from bucket 0 (**Bkt 0**) through bucket 11 (**Bkt 11**). Each bucket is associated with a color code (in-built, not configurable). The generated report represents the distribution of different value ranges with the color codes that are associated with buckets.

You must define the ranges by typing an upper limit in the text box provided for each bucket. The first bucket (Bkt 1) represents all values below the specified upper limit.

If you want to create evenly distributed buckets, specify the upper limit for Bkt 11 and Bkt 0, and then click **Spread Values**.
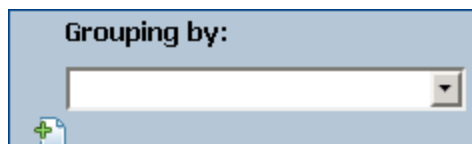
Tip: Since you can define buckets (value ranges) only for each metric series and not for individual metrics, you must select only related metrics in a metric series. You can use the same series and bucket ranges for percentage metrics, but do not place counter metrics (such as Volume) and percentage metrics together on the same series.

Most Changed

Top N:

Top N

Top 10

The Top N option helps you specify the number of data points that you want to see on the report.

To change the Grouping By option:

Grouping by:

In the Grouping By list, select an attribute to define the grouping mechanism for the tabular data on the report.

To select more than one attribute, click **Add New Grouping**. A new Grouping By box with the list of attributes appears.

To remove a grouping attribute, click **Remove Grouping**.

Metrics:

Select metrics of your choice.

To select a metric:

1. Right-click the metric.
2. In the context menu, click **Send to Y1**. You can select only one metric.

> **Tip:** Click **Show Description** to see the metric definition.

You can use the search feature in the Name section to find the metric of your choice. You can use **tag**[1]s to narrow down your search.

Baseline Sleeve

You can select the baseline metric from the Baseline Metric drop-down menu.
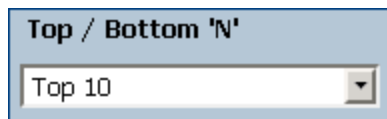
Threshold Sleeve

You can select the threshold metric from the Threshold Metric drop-down menu.

Peak Period

Top/Bottom N:

The Top/Bottom N list helps you specify the number of ranks that you want to see on this report.
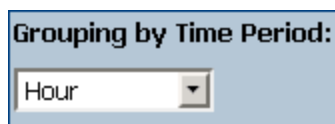
To select data points on the report:

Top / Bottom 'N'

Top 10

In the Top/Bottom N list, select a value of your choice.

Grouping by Time Period:

Grouping by Time Period:

Hour

In the Grouping by Time Period box, select the time period by which you want to group the peak values.

Metrics:

To select a metric:

1. Right-click the metric.

   > **Tip:** You can use the search feature to find the metric of your choice. To search a metric, type the metric name in the Search box.
   >
   > You can also use **tag**[2]s to narrow down your search. To use tags, click inside the Tags box, and then select tags of your choice.

2. In the context menu, click **Send to Y1** if you want to select this metric for the Y1 axis; click **Send to Y2** if you want to select this metric for the Y2 axis. On each axis, you can select up to six metrics.

   > **Tip:** Click **Show Description** to see the metric definition.

---

[1]A tag is a keyword that represents an aggregation type or a domain-specific group.
[2]A tag is a keyword that represents an aggregation type or a domain-specific group.

You can limit the number of rows on the report by specifying a range of values in the From and To boxes on each axis. After typing a value in the From or To box, you must press **Enter**. For the percentage type of metrics, specify values in decimal fractions. For example, to specify 9%, type 0.09.

To remove a metric from an axis:

Right-click the metric on the axis, and then click **Delete**. The metric disappears from the axis.

To change the order of the metric:

Right-click the metric on the axis, and then click **Move Up** or **Move Down**. The metric changes its position on the axis and, subsequently, changes its position on the report when the report is launched.

> **Note:** Right-clicking a metric on the Y1 or Y2 axis does not bring up any context menus when no other metrics are present on the axis.

Set Topology Filters

NPS generates reports created from different metrics that are polled from different nodes, interfaces, and other network elements. You can use the Topology Filter feature to modify the report to contain only the data polled from network elements of your choice. Topology attributes of an element show the topographical characteristics of the element on the network such as its location, group, type, and so on. The Topology Filter feature enables you to filter the complete set of discovered elements and generate reports on elements with specific topology attribute values.

All topology attributes are obtained from NNMi or the collectors and pollers of the iSPIs. You cannot add new attributes or delete existing attributes.

By default, NPS preserves the historical integrity of topology data. Any changes made to topology attributes (Interface Speed, ifAlias values, and so on) are not retroactively applied to the existing historical data, but only to new data.
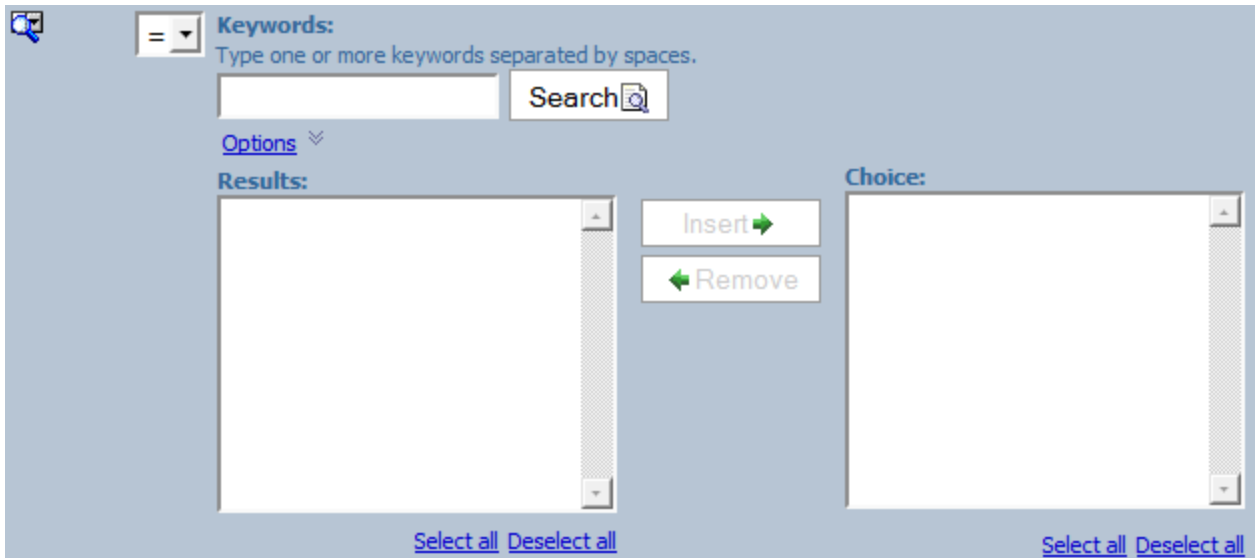
To use topology filters:

1. Log in to NPS from either the NNMi console or the NPS home page.
2. In the navigation panel, open the **Reports** tab and select the report you want to view.

   To modify the topology filters for one of the most used reports, select a report from the **My Folders** location or from the Extension Pack folders.
3. In the navigation panel, open the **Topology Filter** tab and click **Launch Topology Selector**.

Topology Filter

Attribute		Selection
Topology group tracking method ⓘ		SCD Type 1 ▼

Attribute	
InterfaceGroup Name	▭ ▤ ▤ 🔍
NodeGroup Name	▭ ▤ ▤ 🔍
MPLS L3VPN Name	▭ ▤ ▤ 🔍
MPLS VRF Name	▭ ▤ ▤ 🔍
Qualified Interface Name	▭ ▤ ▤ 🔍
Interface Name	▭ ▤ ▤ 🔍
Interface Alias	▭ ▤ ▤ 🔍
Interface Descr	▭ ▤ ▤ 🔍
Interface Index	▭ ▤ ▤ 🔍
Interface Type	▭ ▤ ▤ 🔍
Interface Physical Address	▭ ▤ ▤ 🔍
Interface Speed (In:Out)	▭ ▤ ▤ 🔍
Interface ID	▭ ▤ ▤ 🔍
Interface UUID	▭ ▤ ▤ 🔍
Interface ODBID	▭ ▤ ▤ 🔍
Node Name	▭ ▤ ▤ 🔍
Node Short Name	▭ ▤ ▤ 🔍
Node Contact	▭ ▤ ▤ 🔍
Node Location	▭ ▤ ▤ 🔍
Node Family	▭ ▤ ▤ 🔍
Node Vendor	▭ ▤ ▤ 🔍
Node ID	▭ ▤ ▤ 🔍
Node UUID	▭ ▤ ▤ 🔍
Node ODBID	▭ ▤ ▤ 🔍
Tenant Name	▭ ▤ ▤ 🔍
Tenant UUID	▭ ▤ ▤ 🔍
SecGroup Name	▭ ▤ ▤ 🔍
SecGroup UUID	▭ ▤ ▤ 🔍
Node Annotation	▭ ▤ ▤ 🔍
Interface Annotation	▭ ▤ ▤ 🔍

[Apply] [Reset]

4. Do one of the following for the topology attributes you want to filter:

- To select just one value for a topology attribute, click ▦**Single value select**, and select a value from the drop-down list.

- To select multiple values, click ▦**Multi value select**, and select the values from the drop-down list. You can click, Shift+click, Ctrl+click, click Select all, or click Deselect all.

- To search and select a value from a list of values, click 🔍**Search and select**. The search dialog opens.

5. To apply different types of topology filters, click **Apply** (which will suppress the values that are irrelevant).

   To discard your selection, click **Reset**.

6. Click **Finish** to generate your report.

Cookies are used to preserve your preferred method of selecting each topology filter. For example, if you open a single value selector for an attribute, the next time you run a report, the single value selector will already be open. Be aware that opening more selectors adds extra time to running the report.

The Self Diagnostic tab displays reports on different NPS processes with the data collected from NPS. Therefore, you cannot see any node or interface-related attributes in the Topology Filters section of any Self Diagnostic reports.

You can also exclude selected objects from the report.

Select a Topology Group Tracking Method

NPS generates reports based on topology filters and metrics. Topology filters generate reports on network elements with specific topology attributes, and display data polled from selected network elements. Metrics determine the health and performance of the selected topology elements.

The Preferences pane provides the following topology group tracking methods:

- SCD Type1, which applies to data for which historical changes are not saved. For example, you could create a new group and use it to filter data you collected previously (prior to forming the group).

- SCD Type 2, which applies to data for which the historical changes are saved. This allows you to keep track of the interfaces and nodes that are a member of group at a certain point in time and see the reported data.

SCD (Slowly Changing Dimensions) is a term borrowed from data warehousing concepts and refers to dimensions that have data that changes slowly rather than on a regular time-based schedule.

You can define a group of topology elements and be able to look at data that was collected even before the group existed. The data for this group would be based on current group contents. In such case, you would select Type 1 as your filter.

Or you can monitor the statistics for a group of topology elements such that when the elements of the group change over time, the collected historical data is not updated to reflect the current contents of the group. You would select Type 2 as your filter.

Example

The connectivity between the routers in Site A has been suffering from a high data loss percentage for the past week. You decide to check the routers with the highest volume of incoming and outgoing traffic in this location. You create a node group for the routers in Site A and generate the Top N Report to see the list of routers with the highest volume of data. However, if you use the traditional NodeGroup Name filter to apply the node group you created today, you cannot filter older data. This means, using the default NodeGroup Name filter, you cannot display data older than the node group.

To resolve this, you can use the SCD Type 1 filter. Type 1 groups apply the most recent group definition to all the historical data. In other words, reports generated based on Type 1 topology groups display all data that matches the most recent definition of the group irrespective of the age of the data. In this example, the report, when generated based on Type 1 node group, would display past one week's data.

To select the Type 1 topology group:

1. In the navigation panel, click **Reports** and select the report you want to view.

2. In the navigation panel, click **Preferences**.

3. Select a Type 1 group. For example, to select a Type 1 interface group, select an interface group for the InterfaceGroup_Type1 Name topology filter. Do one of the following for the topology attributes you want to filter:

   - To select just one value for a topology attribute, click ▦**Single Value Select**, and select a topology filter.

   - To select multiple values, click ▦**Multi value select** and the topology filters (by using **Ctrl-click** or **Ctrl+Shift**).

   - To search and select a value from a list of values, click ▦**Search and select**.

   When launching reports from NNMi with a NodeGroup or InterfaceGroup selected, the default behavior is to continue to use the traditional (Type2) "NodeGroup Name" or "InterfaceGroup Name" filters in NPS. If, however, you use the Type1 filters, you can use the Topology Filters page to clear the existing filters and apply the "NodeGroup_Type1 name" or "InterfaceGroup_Type1 Name" filters instead. The Type1 groups are also available in the Top-N reports as an element group-by option.

4. To apply different types of topology filters, click **Apply** and repeat step 3.

   Clicking the Apply button suppresses the values from the other topology attributes that are irrelevant for your selection.

To discard your selection, click **Reset**.

5. Click **Confirm Selection**. NPS generates the report based on the specified topology attributes.

Cookies are used to preserve your preferred method of selecting each topology filter. For example, if you open a single value selector for an attribute, the single-valued selector will already be open the next time you run a report. Avoid opening more selectors than you routinely need because they add overhead (extra time) to running the report.

Change Time Settings with the Time Control Tab

To specify a preferred time zone:

1. Log in to NPS as Administrator.

2. In the navigation pane, click **Reports** workspace, select an Extension Pack, and select a report to apply the time zone.

3. In the navigation panel, click the **Time Control** tab.

4. In the Time Zone field, select a time zone.

5. Click **Submit**.

> **Note:** You must click **Submit** for the changes to take effect.

Keep in mind the following before changing the time zone.

- Semi-Annual Clock Changes

  If you turn clocks back 1 hour, NPS records the metrics for that hour twice, once before the time change, and once immediately after the time change. Collecting statistics twice for the same hour doubles the number of counts (volume in bytes, volume in packets) and the number of samples aggregated into an average.

  But if you turn clocks ahead 1 hour, a gap of an hour occurs in the records because NPS does not collect any data sample for the missing hour. Counts for that hour will be zero, and counts for the whole day will be artificially low.

- Time Zones

  NNMi marks performance data with a universal time stamp known as "seconds since epoch." NPS converts NNMi's universal time stamp into local time, as defined by the time zone selected for NPS. From the perspective of a user in a different time zone, NPS data will be out of sync with the local time.

  If NPS and the NNMi management servers are running on separate systems, the system clocks must match.

Change Time Controls

When you launch a report from the NPS home page, the report shows the data collected over the selected time range relative to the most recent data sample.

To change time controls:

1. Log on to NPS from either the NNMi console or the NPS home page.

2. In the navigation panel, click the **Reports** tab, select the Extension Pack and the report you want to

view.

To modify the time controls for one of the most used reports, click **My reports** or select a report from one of the Extension Pack folders.

3. On the navigation panel, open the **Time Control**pane.

4. You can change the following parameters:

Time Control	Function
Relative Start	Sets the start date for the reports based on the available data
Adjust time	Adjusts the time
Interval	Changes the time range for the report
Grain	Changes the display grain for the report
TimeZone	Enables you to select a preferred time zone
Auto Refresh	Refreshes report data automatically
Hour of Day	Enables you to: • Select the appropriate hour of the day from the list. • For multiple selections, select the appropriate hour from the list, and then hold down CTRL to select each additional hour.
Day of Week	Enables you to: • Select the appropriate day of the week from the list. • For multiple selections, select the appropriate day from the list, hold down CTRL to select each additional day.

5. Click **Submit**.

> **Tip:** If you want to cancel the Hour of Day or Day of Week selection, click **Reset**.

Change the Display Grain for Reports

The Time Grain defines how samples of data are grouped across the Time Range.

A 1 hour chart with a Time Grain of 5 minutes displays 12 samples of data. Adjusting the time range automatically sets the Time Grain to a default value. You can adjust the Time Grain manually using the navigation panel on the NPS home page.

To avoid partial samples of data being grouped together, NPS adjusts the Start Date/Time value to the boundary of the default time range.

To change the default time grain for the reports:

1. Log in to NPS from either the NNMi console or the NPS home page.

2. In the navigation panel, click **Reports** and select the report you want to view.

3. In the navigation panel, select **Time Control**.

4. In the Grain section, select a time grain for your report.

5. Click **Submit**.

Your selection of Time Grain defines the data source for the reports. NPS uses the following types of data source:

- Internal Aggregate tables

  Using a coarse time grain (1 hour or more) ensures that your reports will source their data from the internal Aggregate tables.

- Internal Raw tables

  Using a fine time grain (less than 1 hour) requires NPS to analyze the volume of live data in the database and set the table type accordingly.

  If the start time requested is before the earliest live data, NPS uses the internal aggregate table as the data source. Otherwise, NPS uses the internal raw table as the data source.

Time Ranges with Default Display Grains

When you set a time range for a report, a predefined display grain is applied to the report. The following table shows the mapping between the time ranges with the default display grains.

Time Range	Display Grain
1 Month	1 Day
1 Week	1 Day
1 Day	1 Hour
12 Hours	15 Minutes
2 Hours	5 Minutes
1 Hour	5 Minutes
30 Minutes	5 Minutes
15 Minutes	1 Minute
5 Minutes	1 Minute
1 Minute	As Polled

Your choice of display grain defines the data source for the reports.

The following types of data sources exist within NPS:

- Raw tables (data for up to 14 days)
- Hourly tables (data for up to 70 days)
- Daily tables (data for up to 800 days)

Change the Start Date or Time for Reports

The report shows the data collected from the time and date specified in the Time Control section.

To change the start date or time for a report:

1. Log in to NPS from either the NNMi console or the NPS home page.

2. In the navigation panel, click **Reports** and select the report you want to view.

3. In the navigation panel, select **Time Control**.

4. In the Relative Start field, select **No**.

5. In the Start Date/Time section, click on the **Date** box, and select the appropriate date.

6. In the Start Date/Time section, click on the **Time** box, and select the appropriate time.

7. Click **Submit**

NPS adjusts the Start Date/Time automatically under the following circumstances:

- **Start Date/Time and Time Range extend beyond the available data**.

  NPS moves back the Start Date/Time and selects the requested *Time Range* from the currently available data set.

- **Start Date/Time requested is prior to the oldest data**.

  If your request for a Start Date/Time is prior to the oldest data in the system, NPS advances the Start Date/Time to coincide with the oldest data available. The Time Control advises you of this when your data set begins.

- **Start Date/Time is not on the boundary of the Time Grain**.

  NPS reports always show data grouped into sample sets. The sample size is defined by the Time Grain. If the Start Date/Time is not on the boundary of a Time Grain sized sample, results can look unusual when displayed. To prevent this, the Time Control always attempts to round down the Start Date/Time to the beginning of the nearest Time Grain sized sample.

  For example, a Start Time of 12:15 a.m. and an Interval of 1 day has a default Time Grain of 1 Hour, resulting in 24 samples. However, the samples for the hour representing 12:00 only have data from 45 minutes and will appear unusually low. To correct this, the Time Control moves the Start Date/Time back to begin at 12:00 a.m. automatically.

Change the End Date or Time for Reports

The report shows the data collected from the start time until the date specified in the Time Control tab.

To change the end date or time for a report:

1. Log in to NPS from either the NNMi console or the NPS home page.

2. In the navigation panel, click **Reports** and select the report you want to view.

3. In the navigation panel, select **Time Control**.

4. In the Relative Start field, select **No**.

5. In the End Date/Time section, click on the **Date** box, and select the appropriate date.

6. In the End Date/Time section, click on the **Time** box, and select the appropriate time.

7. Click **Submit** to generate the report for the time range ending on the date and time specified in the End Date/Time.

NPS adjusts the Start Date/Time automatically under the following circumstances:

- **End Date/Time and Time Range extends beyond the available data**.

  NPS moves back the End Date/Time and selects the requested Time Range from the currently available data set.

- **Time Range requested is prior to the oldest data**.

  If you request a time range (time between the Start Date/Time and End Date/Time) that is prior to the oldest data in the system, NPS advances the Start Date/Time and End Date/Time to coincide with the oldest data available.

  The Time Control advises you of this when your data set begins.

Change the Time Range (Interval) for Reports

The time ranges fall into the following categories:

- **Fixed Period Length relative to the most recent sample stored on the system until now**

  To select this type of time range:

  a. Log in to NPS from either the NNMi console or the NPS home page.

  b. In the navigation panel, click **Reports** and select the report you want to view.

  c. In the navigation panel, select **Time Control**.

  d. In the Relative Start section, select **Yes**. If you select a Time Range that is relative to the most recent data (for example, Last 1 Hour), NPS modifies the Start Date/Time accordingly.

- **Fixed Period Length starting at the given Start Date/Time**

  You can select a time range applicable to a specific start and end date/time.

  If you select a time range that would push the start date earlier than the available data sample, NPS automatically adjusts the time range to set the start date/time at the earliest data record.

  - To select the start date or time, see Change the Start Date or Time for the Report.

  - To select the end date or time, see Change the End Date or Time for the Report.

Selecting a Time Range automatically sets the Time Grain to an appropriate value. You can change the display grain manually.

Too many data points on a chart can make the results difficult to read.

Select a Relative Start for Reports

The Start Date/Time defines the point of origin of the report. The report shows the data collected from the time and date specified in the Start Date/Time field.

To set relative start for your report:

1. Log in to NPS using either the NNMi console or the NPS home page.

2. In the navigation panel, click **Reports** and select the report you want to view.

3. Open the **Time Control** pane.

4. In the Relative Start field, select **Yes**. By default, relative start is set to Yes.

5. Select a time range for the report from the drop-down box. For example, to generate a report based on the last month's data, select 1 Month from the drop-down box.

6. Select a different time grain if required.

7. Click **Submit**.

If you request a time range that is relative to the most recent sample of data, NPS moves the Start Date/Time appropriately. In some cases, especially with a very recent installation, this can result in a Start Date/Time that appears to be older than the start of the first data collection interval.

Adjust Time

The adjust time buttons have the following function:

Icon	Description
⊕	**Adjust time - zoom in**. Enables you to view data points in a smaller time period.
⊖	**Adjust time - zoom out**. Enables you to reset and view the original data.
←	**Adjust time - back**. Enables you to view a previous time period.
→	**Adjust time - forward**. Enables you to select a future time period.

Change the Default Date and Time for Reports

To change the default date and time for reports:

1. Log in to NPS from either the NNMi console or the NPS home page.
2. In the navigation panel, click **Reports** and select the report you want to view.
3. In the navigation panel, click **Time Control**.
4. Select the following:

   - In the Hour of Day list, select a default time for the reports.

   - In the Day of Week list, select a default date for the reports.

5. Click **Submit**.

To cancel the selection, click **Reset**.

Change the Default Time Zone for Reports

By default, reports are generated according to NPS's time zone. You can select a different time zone while viewing or scheduling a report in the Time Zone field.

NPS uses the selected time zone to generate all subsequent reports.

The selected time zone overrides the time zone preferences associated with your user account. If you set the time zone back to the default value, NPS uses the following time zones:

- **The time zone set in your user account preferences**

  This overrides the NPS time zone. To change the default time zone, see Change the Report Time Zone with the Time Control Tab

- **The time zone used by NPS**

  If you did not set any preference for time zone in your user account, NPS uses this option.

The Filter Summary (🔽) of each report displays the selected time zone unless the effective time zone is the default NPS server time zone, known as "Server Time."

Select the default time zone if you generate reports in a large scale environment. Selecting a different time zone than Server Time requires NPS to convert the time stamps for the selected time zone while generating the report. This can consume a lot of time.

Time Changes and Time Zones

Semi-annual clock changes

If you turn clocks back one hour, statistics for one hour are recorded twice, once before the time change, and once immediately after. Collecting statistics twice for the same hour doubles the number of counts (volume in bytes, volume in packets) and the number of samples aggregated into an average. If you turn clocks ahead one hour, a gap occurs because no samples are collectable for the missing hour. Counts for that hour are zero, and counts for the whole day are artificially low.

Time zones

NNM marks performance data with a universal time stamp known as **seconds since epoch**. NPS converts NNM's universal time stamp into local time, as defined by the location of NPS. From the perspective of a user in a different time zone, NPS data is out of sync with the local time.

If NPS and the NNMi management server are running on separate systems, make sure that the system clocks match. Otherwise, confusion will arise as you move from one application to the other.

Special Scenarios

- *Start Date/Time and Time Range extend beyond the available data*: The Start Date and Time automatically shift to allow the requested Time Range to be selected from the available data set.

- *Start Date/Time requested is prior to the oldest data*: If you request a Start Date/Time that is prior to the oldest data in the NPS, the Start Date and Time shift to coincide with the oldest data available. The Time Controls section notifies you of this after you make the selection.

- *Start Date/Time is not on the boundary of the display grain*: Graphs and charts always show data grouped into sample sets, with the sample size defined by the display grain. The display grain defines how samples of data are grouped across the time range. For example, a 1-hour chart with a display grain of 5 minutes displays 12 samples of data. If the Start Date/Time is not on the boundary of a display grain-sized sample, results can look unusual. The NPS always attempts to round down the Start Date/Time to the beginning of the nearest display grain-sized sample.

For example, a Start Time of 12:15 a.m. and a Time Range of 1 Day will have a default Time Grain of 1 Hour, resulting in 24 samples. However, the samples for the hour representing 12:00 will only have data from 45 minutes and will appear unusually low. To correct this, the Time Control will move the Start Date/Time back to begin at 12:00 a.m. automatically.

Related Topics

Change the Default Settings for Reports

Time Ranges with Default Display Grains

Refresh Report Data Automatically

The report is automatically refreshed with the fresh data samples if you select the auto refresh rate with a relative start date and time.

For example, if you selected Last 1 Hour as the Time Range, and 1 Minute in the Auto Refresh field, the report displays the data samples collected in the past hour and refreshes the data every minute. As a result, if you generated the report at 12:00 p.m., the report shows data from 11:00 a.m. to 12:00 p.m., and refreshes the report data at 12:01 p.m.

To automatically refresh the report data:

1. Log in to NPS from either the NNMi console or the NPS home page.

2. In the navigation panel, click **Reports** and select the report you want to view.

3. In the navigation panel, select **Time Control**.

4. In the Relative Start field, select **Yes**. By default, relative start is set to Yes.

5. From the **Auto Refresh** drop-down box, select one of the following values:

    - 30 Seconds

    - 1 Minute

    - 5 Minutes

    - I Hour

    By default, Auto Refresh is set to Off.

6. Click **Submit**

Cross Launch a Report

Where appropriate, the report or reportlet presents links to navigate to another related report or reportlet. This feature enables you to refine your selection based on time, topology, or metrics displayed on the report.

To cross launch a report or reportlet:

1. Click the data point in the report or reportlet that launches another report or reportlet.

    For example, click a point of time in the X axis of the Chart Detail report to drill down to the Heat Chart report displaying the metric performance for the hour selected. Or, open the Top N report to see the list of devices contributing most to the data at that interval.

    You can set your browser to open either a new window or a new tab when you drill down. This enables you to revisit the initial report and examine another data point at a later time.

2. Select the metrics for the drill-through report or reportlet.

    NPS saves your selection for the next session.

3. Select a report or reportlet from the navigation panel. The new report or reportlet displays the selected metric.

    For more information on the contextual navigation available for each report type, see Contextual Navigation.

Examples

- The Dashboard report for Interface Health shows a high availability exception count. Clicking the needle of the gauge launches the Report Menu with Availability Exceptions as the primary metric. You can then launch the Top N report to find which set of nodes is causing the availability exceptions.

- Looking at the Top N interfaces and their network discard figures, one interface seems more problematic than usual. Clicking the interface name launches the Report Menu with that interface as a topology filter. You can then launch the Chart Detail report to examine discard figures over the selected time interval for just that interface and determine when the problem started.

- Using the Heat Chart report for Interface Health shows a peak in traffic on one particular day. Clicking the cell with the high value launches the Report Menu with that day already selected. You can then launch the Interface Health Headline report and get an overview of multiple metrics for each hour during that day.

Create Bookmarks in Reports

NPS enables you to create bookmarks for reports in the NPS home page.

These bookmarks encapsulate all the selected settings (Time, filters, Options) and enable you to quickly visit the report again from your browser.

Bookmarks are convenient for emailing and sharing reports.

To have the bookmark open the report within the same time range as the original report, select a non-relative time range; for example, Start Date/Time with "1 Hour" and not "Last 1 hour" as the time range. Otherwise, the report will show the most recent data.

To create a bookmark for a report:

1. Launch a report.
2. Click **Show Bookmark**. A URL appears.

   Do the following:

   - *Only for Microsoft Internet Explorer*: Click **Add Bookmark**. A bookmark to the report is automatically added to the browser with a new entry to the Favorites list.

   - *Only for Mozilla Firefox*: To create a new bookmark by using the displayed URL, right-click the URL and select **Bookmark This Link**.

The bookmarked URL retains all the properties of the report.

Create Bookmarks in Reportlets

NPS enables you to create bookmarks for reportlets in the NPS home page.

These bookmarks encapsulate all the selected settings (Time, filters, Options) and enable you to quickly visit the reportlet again from your browser.

Bookmarks are convenient for emailing and sharing reportlets.

If you want the bookmark to open the reportlet within the same time range as the original reportlet, select a non-relative time range; for example, Start Date/Time with '1 Hour' and not 'Last 1 hour' as the time range. Otherwise, the reportlet will show the most recent data.

To create a bookmark for a reportlet:

1. Launch a reportlet.

2. Click ⬦**Show Toolbar** to display the reportlet toolbar.

3. Click 📖**Show Bookmark**. A URL appears.

   Do the following:

   - *Only for Microsoft Internet Explorer*: Click **Add Bookmark**. A bookmark to the reportlet is automatically added to the browser with a new entry to the Favorites list.

   - *Only for Mozilla Firefox*: Create a new bookmark by using the displayed URL.

The bookmarked URL retains all the properties of the reportlet.

Add URL Actions in the NNMi Console

After you generate a bookmark, you can create a new Actions menu item with the newly generated bookmark, which enables you to launch a report or reportlet directly from the NNMi console.

To add URL actions in the NNMi console to launch a report or reportlet:

1. Generate the bookmark for the required report or reportlet.

2. In the NNMi console, select the Configuration workspace from the workspace navigation panel.

3. Select **User Interface**.

4. Select **Menu Items**.

5. Click ✳ **New**.

   The Menu Item form opens.

6. In the basics section of the Menu Item form, specify the following details:

   - Menu Item Label: The name of the report or reportlet you want to launch.

   - Unique Key: A unique identifier for the menu item. Type an identifier.

     You can use alpha-numeric and period (**.**) characters. Do not exceed 80 characters. Do not use spaces and underline characters. To make the value unique, use the Java name space convention while specifying this value. Include the Menu Item Label value as part of the unique key; for example:

     com.*<company_name>*.nnm.menu.item.*<menu_item_label>*

   - Author: Specify the default value Customer, or click [icon] ▼ to add a new author name; for example, NPS.

   - Parent Menu: Retain the default value—Actions.

   - Ordering: Specify an integer between 1 and 100. The lowest number appears at the top of the Actions menu. The highest number appears at the bottom of the Actions menu.

   - Prepend Separator: Select this option to add a horizontal line before the report or reportlet menu item in the Actions menu.

   - Enable: Select this option to use the report or reportlet bookmark.

7. Retain all other default settings in the Basics section.

8. In the Menu Item Context tab, click ✳**New**.

   The Menu Item Context form opens.

9. In the Basics section, click [icon] ▾ and **New Launch Action**.

   The Launch Actions form opens.

10. In the Launch Actions form, specify the following:

    - Name in the Name box

    - Bookmark URL in the Full URL section

11. Click [icon] **Save and Close**.

12. Click [icon] **Save and Close** in the Menu Item Context form.

13. Click [icon] **Save and Close** in the Menu Item form.

14. Click [icon] **Save and Close** in the User Interface Configuration form.

Now you can launch the report or reportlet from the Actions menu.

View a Report using Different Report Formats

NPS and the BI Server use web-based HTML as the default report format.

You can also view a report in Adobe PDF or Microsoft Excel formats.

To view a report in non-HTML format:

1. In the NNMi console, click **Actions** > **HPE NNM iSPI Performance** > **Reporting - Report Menu**. The NPS home page opens.

2. On the navigation panel, select the **Reports** tab. NPS presents the Self Diagnostic reports for the default, ready-to-use Self Diagnostic Extension Pack. Depending on the number of the NNM iSPI products or Extension Packs available on the NPS system, you could see other report folders.

3. Under an Extension Pack, select a report.

4. In the Time Controls tab, select a time range for the report. For information on selecting the time range for a report, see Change Time Controls.

5. Click [icon] ▾ **View in HTML Format**.

6. Select one of the following options:

   - **View in PDF Format**. A PDF version of the report appears in the same browser window.

   - **View in Excel Options** and **View in Excel 2007 Format**, **View in Excel 2007 Data**, or **View in CSV Format**.

Run a Report in Background

Some reports can take a long time to generate. You can run such reports in the background and have the output saved to a folder on the BI Server. You can also have it emailed.

If the report takes more than a few seconds to generate, a prompt will appear offering an alternate delivery method for your report.
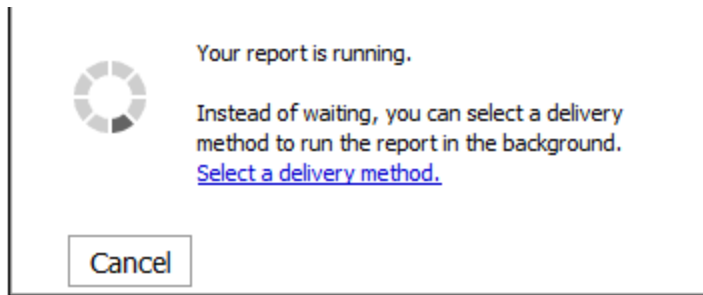
If you click **Select a Delivery Method**, the following screen appears:
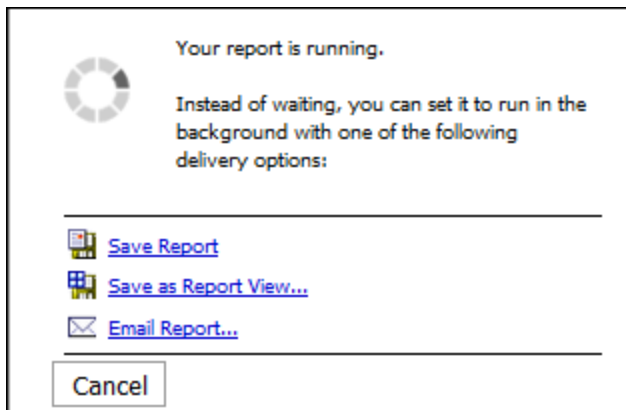
If you choose Save as ReportView, enter a name for the report and save to MyFolders or to another folder that allows you to save reports. You can then view the new ReportView at your convenience from the left-hand navigation panel.

If you choose email, enter the email fields.

Scheduling Reports

You can use the BI Server Portal of the NPS console to schedule delivery of reports.

> **Note:** To be able to schedule reports in the NPS console, you must log on to NNMi as an administrator or Level 2 operator.

1. In the navigation pane of the NPS console, click **BI Server > Portal**. The HPE NNM iSPI Performance BI Portal opens.
2. In the Public Folders tab, navigate to the Extension Pack of your choice.
3. On the page where all the reports are listed, click **More** against the report of your choice. The Perform and Action page opens.
4. Click **New schedule**.
5. Select the frequency (under the Frequency section) and delivery options (under the Options section). To configure delivery options, you must select **Override the default values** under the Options section.

> **Note:** Do not select the XML format. Delivery of reports in the XML format is not supported. Do not select **Delimited text (csv)** for any reports other than Top N and Chart Detail.

6. Click **OK**.

Set Up an Email Server

To use the NPS console to email reports, you must run the `configureBIEmailServer.ovpl` script one time to configure the details for your email server.

When you run the script, any existing email server information will be overwritten. Always restart the BI Server after running the `configureBIEmailServer.ovpl` script.

To restart the BI Server, run the following commands:

1. **stopBI.ovpl**
2. **startBI.ovpl**

Limit Access to Reports

The Administrator can manage access to reports that are run when a ReportView is created.

Access to such reports can be controlled in any of the following ways:

- Log in as the tenant before creating the ReportView.
- Create the ReportView as Administrator but delete the saved output.
- Apply tenant name filtering before saving the ReportView.

Create Custom Reports with Query Studio

You can create custom reports with metrics from all available Extension Packs.

HPE does not offer any support in using Query Studio while creating custom reports.

To create a new report:

1. Log in to NPS as Administrator.
2. On the NPS home page, click the **BI Server** tab in the navigation panel.
3. Select **Query Studio**. The Select a package (Navigate) window opens.
4. Select a Public Folder.

   The options are:
   - Cross Domain

   - iSPI Metrics

   - Quicklaunch ReportViews

- Self Monitoring (admin users only)

- User Groups

For example, select iSPI Metrics.

5. Select an Extension Pack; for example, Interface_Health.

The HPE NNM iSPI Performance BI Query Studio window opens.

The Query Studio window consists of a Menu panel; two toolbars (Standard and Style), either of which can be hidden; and the work area where you will see your report being built.

6. Click **Insert Data** in the Menu panel to list the query items for the Extension Pack you selected.

For our Interface Health example, the options are:

- d_Interface_Health_Time

- d_InterfaceTopology

- g_InterfaceGroup Type1_InterfaceTopology

- g_InterfaceGroup Type2_InterfaceTopology

- g_NodeGroup Type1_InterfaceTopology

- g_NodeGroup Type2_InterfaceTopology

- g_MPLS L3VPN Type1_InterfaceTopology

- g_MPLS L3VPN Type2_InterfaceTopology

- g_MPLS VRF Type1_InterfaceTopology

- g_MPLS_VRF Type2_InterfaceTopology

- f_Virtual_InterfaceMetrics

- nnmdetails_Interface_Health

7. Select a time range as follows:

   a. Expand the time query item; for example, d_Interface_Health_Time.

   b. Right-click a time range—for example, Day (of Week)—and click **Insert** or drag it into the work area.

   The time range appears in a column in the work area. You can add more time ranges.

   You can also add time dimension columns that will return relative time queries. For more instructions, see "Adding Relative Time Queries to Custom Reports with Query Studio" on the next page.

   To delete a time range from the work area, right-click the column for the time range, and click **Delete**.

8. Select a topology group as follows:

   a. Expand the InterfaceTopology query item; for example, d_InterfaceTopology.

   NPS displays only the topology groups associated with your security group.

   b. Right-click a topology group and click **Insert**.

   The attribute appears in a column in the work area.

To display meaningful data in the report, select a topology group relevant to the metrics you plan to select.

To delete an attribute from the work area, right-click the column for the attribute and click **Delete**.

Node groups and interface groups are separate items in the menu. To select a node or interface group, repeat the previous steps.

9. Add a topology context as follows:

    a. Expand an InterfaceGroup or NodeGroup query item.

    b. Right-click a group and click **Insert**.

    The group will be added to your report.

    To delete an attribute from the work area, right-click the column for the attribute and click **Delete**.

    NPS displays all topology attributes from all installed Extension Packs. However, not all topology attributes are applicable to metrics from all Extension Packs. If you point the mouse arrow over a topology attribute, a tool tip shows the Extension Packs that support that attribute.

10. Select a metric for the report as follows:

    a. Expand the f_Virtual_InterfaceMetrics query item.

    b. Right-click a metric click **Insert**. The metric will appear in your report.

    To delete a metric from the work area, right-click the column for the metric and click **Delete**.

11. Format the report with the style toolbar.

12. To give your report a title:

    a. Click **Title**. The Edit title area window opens.

    b. Enter a title and, if you want, a subtitle.

    c. Click **OK**.

13. To run the report, click **Run Report** in the Menu navigation panel. You will be prompted to enter a TimeZone. Enter the same value you selected when you performed the Insert Data steps.

14. To save the report, click **Manage File** in the Menu navigation panel. Select the **Save As...** option.

    Make sure to save it in My Folder or another location. Do not save any reports made using Query Studio in any of the Extension Pack folders.

15. You can view the report when you log in to the iSPI Performance BI Server Portal.

Do not use Cognos BI Report Studio to create or customize reports. HPE does not offer support for reports created with Report Studio.

For more information on using Query Studio, see the following BI Server documents:

- *Query Studio Quick Tour*
- *Query Studio User Guide*

Adding Relative Time Queries to Custom Reports with Query Studio

You can add time dimension columns to your custom report that will return relative time queries. The Date/Time dimension column provides the following time grain options:

- Month
- Week
- Day (of Month)
- Day (of Week)
- Day
- Hour (of Day)
- Hour
- Hour (epoch)
- Half Hour (of Day)
- Half Hour
- Qtr Hour (of Day)
- Qtr Hour
- 10 Minute (of Day)
- 10 Minute
- 5 Minute (of Day)
- 5 Minute
- Minute
- Minute (epoch)
- This Week
- This Day
- This Hour
- This Minute

For example, you could build a report that calculates the difference between Week and This Week.

1. Launch Query Studio, select the Interface Health ExtensionPack and expand the f_Virtual_InterfaceMetrics folder.

2. Select the options Week and This Week and drag them into the work area.

3. To calculate the data, press **Ctrl** and click both column headings.

4. Click the Calculate button. The Calculate dialog opens.

5. In the Operation box, select - (difference in days). The options are:

    - - (subtraction)

    - - (difference in days)

    - - (difference in months)

    - - (difference in years)

    - Add days

    - Add months

    - Add years

- Earliest

- Latest

- Year

- Month (number)

- Week

- Day

6. Under **Order**, select **Days between(Week, This Week)** and click **Insert**. A new column, Days between (Week, This Week) is added to the report.

7. Click the new column and click the ▽ Filter button. The Filter dialog opens.

8. Under **From**, enter 0, which represents "this week."

   Under **To**, enter 7, which represents "last week."

   A range can also be entered; for example, 0-7.

   Click **OK**.

   The value is calculated.

You can add more columns to this report. For example, you could add a measure such as Utilization (pctile95). And you can delete columns until the report reflects the information you need.

Contextual Drill-Through

NNM iSPI Performance for Metrics enables you to analyze the performance of one metric using multiple reports.

The ▽ icon and adjacent text tell you how the element and time range filters are set. Unless you modify the defaults, NPS enables the default elements and time range for the report. With the default elements and time range, the report displays every component that is delivering performance data to NNMi.

Carrying the Filter Settings Forward

The metric values that appear in reports are clickable links. When you click a metric value on a report, NPS saves the following information and opens the NPS home page:

- Selected topology filter
- Selected metrics (if available)
- Metric value that you clicked on
- Selected time range

From the navigation panel, select a report. When the report opens, the content of the new report is limited to the network elements related to the metric value you clicked in the previous report.

Preserving the context (preserving the filter and option settings when you move from one report to another) is known as contextual drill-through.

Contextual drill-through is not available if the metric value displayed on the report is NULL. If you select a component type (for example, CPU) that does not have metric values available for the selected metric (for example, Memory Utilization), the contextual drill-through link might not be enabled.

If you navigate from a summary report to a live report, only the applicable filters carry forward.

For example, a metric selection carries forward, but a Day of Week or Hour of Day setting does not carry forward to the live report because these settings do not apply to live reports. Instead, the live report opens with the time period set to the default, which is Last 15 Minutes in most cases. If you navigate from a live report to a summary report, the summary report ignores the time period you selected and instead shows data for the default time period, either Last Day or Month-to-Date.

Using Cross-Domain Reports

At the time of installing a new Extension Pack, NPS generates a cross-domain Extension Pack that combines the metrics and topology attributes of all Extension Packs that are available on the NPS system.

The cross-domain Extension Pack helps you view reports built with the data provided by all Extension Packs that are installed on the system (across all iSPIs). The cross-domain Extension Pack includes all types of standard reports and reportlets. For a list of standard reports, see "Report Types" on page 28.

The cross-domain Extension Pack is created and a new item (CrossDomain Reporting) appears in the navigation pane in the NPS console under **Cross Domain > AllExtensionPacks > AllExtensionPacksMetrics**.
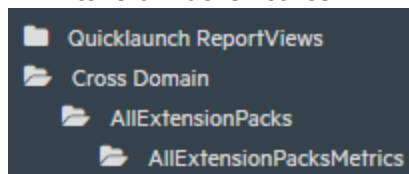
In addition, the cross-domain Extension pack adds the following three Report Views in the list of Report Views:

- Top Problematic Items
- Days To Threshold
- Managed Inventory

For more information about the cross-domain Report Views, see "Cross-Domain Report Views" on page 82.

The cross-domain extension pack shows empty reports when the NPS database contains no data for the Component Health extension pack. Make sure the NPS database contains some data for the Component Health extension pack.

To check that the database contains data for all extension packs, run the following command on the NPS system:

dbsize.ovpl -s

> **Note:** NPS requires a longer time to generate cross-domain reports after you configure NPS to use a non-default time zone (using the Time Control panel).

Using metrics

When you launch a report from the cross-domain extension pack, you can use all metrics that are supported by installed extension packs. On a single report, you can use metrics that originally belonged to different extension packs. For example, you can choose the `Utilization In (avg)` metric from the Interface Health extension pack and the `Number of Packets (Outgoing)` metric from the Interface Traffic extension pack (provided by the NNM iSPI Performance for Traffic).

Using topology attributes

With a cross-domain report, you can use all topology attributes that are supported by installed extension packs.

The `ObjectType` attribute enables you to perform a filtering action on the report to view entities of a similar type (such as interfaces or components).

The `ObjectName` attribute enables you to perform a filtering action on the report to view specific monitored element (such as a specific node, interface, or CPU).

Adjusting for Past Outages

You can use NNMi's Schedule Node Outage feature or the `nnmscheduledoutage.ovpl` command to schedule planned outages. After you schedule outages, NNMi changes the node status to Out-of-Service and suspends any discovery or monitoring of that node during the specified outage time period.

This feature enables you to make a record of a scheduled outage in the past. To adjust the data in the NPS database to reflect the effect of planned outages that occurred in the past, you must perform the following tasks on the NNMi management server:

1. Log on to the NNMi management server.

2. Make sure that the following directory exists:

   On Windows

   `%nnmdatadir%\shared\perfSpi\datafiles\outage`

   On Linux

   `/var/opt/OV/shared/perfSpi/datafiles/outages`

   If necessary, create the directory manually.

3. Create a cron job (on a Linux system) or a scheduled task (on a Windows system) to run the following command:

   > **Note:** The cron job or scheduled task must run at least once a day.

   On Windows

   %nnminstalldir%\bin\nnmscheduledoutage.ovpl dump -retroactive -start "$(date +%Y-%m)-01T00:00" -end "$(date --date='next month' +%Y-%m)-01T00:00" -format CSV > %nnmdatadir%/shared/perfSpi/datafiles/outages/$(date +%Y-%m-%d).csv

   On Linux

   /opt/OV/bin/nnmscheduledoutage.ovpl dump -retroactive -start "$(date +%Y-%m)-01T00:00" -end "$(date --date='next month' +%Y-%m)-01T00:00" -format CSV > /var/opt/OV/shared/perfSpi/datafiles/outages/$(date +%Y-%m-%d).csv

> **Note:** The outage must be fully contained within the retention period of the raw/detailed data. You can check the retention period of the raw/detailed data in the Configuration Utility.

Using NNM iSPI Performance for Metrics Reports to Monitor the Network

Spot the nodes that are down.

The Node Availability and Node Reachability metrics indicate if a node is down and if it is unable to communicate with NNMi. A report can show if a node has been down for a significant period of time.

To spot the nodes that are down:

1. In the NPS console, click **iSPI Metrics > Component_Health > ComponentMetrics > Top N**.
2. On the Top N report, click **Options**, and make the following changes in the Report Options pane:

   - Select **Sort All in Ascending** in the Top/Bottom 'N ' box.

   - Make sure Grouping By is set to Node Name.

   - Send the Node Availability (avg) metric to the Y1 pick list.

   - Send the Node Reachability (avg) metric to the Y2 pick list.

     > **Tip:** Delete the other (default) metrics from the pick lists.

3. Click **Confirm Selection**.
4. The Top N report lists all nodes with ascending values of Node Availability (avg). Nodes that are down for at least past one hour should appear in the first few rows of the table with the Node Availability (avg) value marked as 0%. The Node Reachaility (avg) value for such nodes should be close to 0%.

   > **Tip:** You can change the Time Control settings to spot nodes that were down in the past.

Identify the nodes and interfaces that violated thresholds set in NNMi

1. Make a note of metrics (monitored attributes) for which thresholds were configured.

   You can see the configuration details in the Node Settings (for nodes) and Interface Settings (for interfaces) forms. Configured metrics are listed in the Threshold Settings tab.

2. In the NPS console, click **iSPI Metrics > Component_Health > ComponentMetrics > Top N** (for components) or **iSPI Metrics > Interface_Health > InterfaceMetrics > Top N** for interfaces.
3. On the Top N report, click **Options**, and make the following changes in the Report Options pane:
4. Make sure the Group By option is set to Qualified Component Name or Qualified Interface Name.
5. In the Y1 pick list, add a threshold exception count metric.

   A metric value with the `Threshold Exception Count (sum)` extension represents the total count of data samples that breached the threshold value.

6. In the Range column of the Y1 pick list, select Excluded; in the From column, type 0; in the To column, type 1.

This configuration ensures the report is going to discard all components or interfaces with no threshold breaches.

7. Click **Confirm Selection**. The Top N report lists the components or interfaces that experienced threshold breaches in the past one hour. You can change the time range from the Time Control workspace to view threshold exceptions detected in a different time period.

Identify the node and the interface on the node that experienced the maximum packet loss in the last one day

1. In the NPS console, click **iSPI Metrics > Interface_Health > InterfacetMetrics > Top N**.

2. On the Top N report, click **Options**, and make the following changes in the Report Options pane:

   - Select **Sort All in Ascending** in the Top/Bottom 'N ' box.

   - Make sure Grouping By is set to Node Name.

   - Send the Node Availability (avg) metric to the Y1 pick list.

   - Send the Node Reachability (avg) metric to the Y2 pick list.

     Tip: Delete the other (default) metrics from the pick lists.

3. Click **Confirm Selection**.

4. The Top N report lists all nodes with ascending values of Node Availability (avg). Nodes that are down for at least past one hour should appear in the first few rows of the table with the Node Availability (avg) value marked as 0%. The Node Reachaility (avg) value for such nodes should be close to 0%.

   Tip: You can change the Time Control settings to spot nodes that were down in the past.

Using Baselines

The NNM iSPI Performance for Metrics is capable of monitoring network elements by using a calculated baseline for use in reporting, forecasting, and thresholds. You can use deviation from historical baseline as a threshold (instead of a fixed threshold value). Baselines are calculated based on previously collected values for each network element, taking into account variations based on the time of the day and the day of the week. The NNM iSPI Performance for Metrics creates exceptions and sends incidents to the NNMi Incident View when the monitored metric value violates the baseline threshold.

Baselines provide effective monitoring of network elements for which the pattern of usage varies over a period of time. For example, network usage during peak business hours is much higher than the network usage on weekends.

While using a baseline, the NNM iSPI Performance for Metrics uses the following workflow:

1. Analyzes the historical data (all the data samples polled over the monitoring cycle)

2. Computes a baseline for the metric by using the Triple Exponential Smoothing technique:

   Key features of the smoothing technique used by the NNM iSPI Performance for Metrics:

   - Uses an initialization period of three weeks from the start of the monitoring; during this period, no incidents are generated

- Uses a seasonality cycle of one week (with 5-minute granularity)

- Uses the number of deviations (above and below) from the values specified in the Baseline Settings form to compute the range of normal values or the **baseline sleeve**

  For more information, see "How the NNM iSPI Performance for Metrics Calculates Baselines" on page 78.

3. Determines exceptions; you can view these exceptions by choosing the special *baseline exception count* and *baseline exception rate* metrics on reports.

4. Sends incidents to the Incident View of the NNMi console.

For an example, see "Example" on page 76.

List of Metrics for Baselines

Irrespective of which network element (node or interface) you choose, all of the following metrics appear in the Baseline Settings form. You must carefully choose only the relevant metrics.

Metric	Interface/Node Monitoring	Description
Input Utilization	Interface	The total number of incoming octets traversing the interface as a percentage of the total possible number of octets (based on the ifSpeed value). From Interface to Interface, the exact MIB variables queried vary based on interface speed and whether the system supports the high speed counters for interfaces. Each interface in an Interface Groups has its utilization calculated by taking the total traffic on all administratively up interfaces in the group and dividing that by the total possible bandwidth.
Output Utilization	Interface	The total number of outbound octets traversing the interface as a percentage of the total possible number of octets (based on the ifSpeed value). From Interface to Interface, the exact MIB variables queried vary based on interface speed and whether the system supports the high speed counters for interfaces. Each interface in an Interface Group has its utilization calculated by taking the total traffic on all administratively up interfaces in the group and dividing that by the total possible bandwidth.
Backplane Utilization	Node	Percentage of backplane usage in relation to the total amount of backplane resources available.
Buffer Utilization	Node	Percentage of buffer usage in relation to the number of buffers available.
CPU 1Min Utilization	Node	Percentage of CPU usage in relation to the total amount of CPU available. This percentage is measured at 1-minute intervals.
CPU 5Min Utilization	Node	Percentage of CPU usage in relation to the total amount of CPU available. This percentage is measured at 5-minute intervals.
CPU 5Sec Utilization	Node	Percentage of CPU usage in relation to the total amount of CPU available. This percentage is measure at 5-second intervals.

Metric	Interface/Node Monitoring	Description
Disk Space utilization	Node	Percentage of disk space usage in relation to the total amount of disk space available.
Memory Utilization	Node	Percentage of memory usage in relation to the total amount of memory available.
Management Address ICMP Response Time	Node	Indicates the Internet Control Message Protocol (ICMP) response time (in milliseconds) from the management station to the target node.

Step-by-Step Procedure to Configure a Baseline

1. Log on to the NNMi console as an administrator.
2. Click **Configuration > Monitoring Configuration**. The Monitoring Configuration form opens.
3. Go to the Interface Settings or Node Settings tab, select the interface or node group for which you want to configure a baseline, and then click 📁 **Open**. The Interface Settings or Node Settings form opens.

   > **Tip:** If monitoring settings are not already configured for the group, fill out the forms in the left pane.
   >
   > If you want to create a baseline for input or output utilization, make sure to select the Enable Interface Performance Polling check box. To create baselines for any other metrics (see "List of Metrics for Baselines" on the previous page), make sure to select the Node Performance Polling check box.

4. In the Interface Settings or Node Settings form, go to the Baseline Setting tab.
5. Click ✳ **New**.The Baseline Setting form opens.
6. In the Baseline Threshold Settings section, specify the details with the help of Table: Baseline Settings.
7. In the Baseline Deviation Settings tab, specify the details with the help of Table: Baseline Deviation Settings.

Baseline Settings

Attribute	Description
Monitor Attribute	Select a metric from the Monitored Attribute list (one of the metrics listed in "List of Metrics for Baselines" on the previous page).
Threshold Enabled	Select this option if you want to send incidents to the NNMi's Incident View. **Note:** If you do not select this check box, the NNM iSPI Performance for Metrics does not send incidents in NNMi's Incident View. However, you can still monitor exceptions with the help of reports.

Baseline Settings, continued

Attribute	Description
Duration[a]	Designate the minimum time within which the value must remain out of the configured Baseline Range before the state changes to Abnormal Range and (optionally) an incident is generated. Use the Baseline Deviation Settings tab to set the upper and lower limits of the baseline range. Note the following: • If you do not configure a Baseline Range, the NNM iSPI Performance for Metrics uses the default value of 3 deviations. • The Polling Interval should be less than or equal to the Duration.
Duration Window[a]	Designate the window of time in which the Upper Baseline Limit or Lower Baseline Limit criteria must be met. **Note:** The value must be greater than 0 (zero) and can be the same as the Duration value.

[a]The Duration and Duration Window values are used for detecting exceptions and sending incidents to the NNMi Incident View. If you select a duration of A seconds and a duration window of B seconds (A<B), the NNM iSPI Performance for Metrics registers an exception and triggers an incident when the metric value lies outside of the calculated baseline sleeve for A seconds (or longer) during the last B seconds. For more details, see .

Baseline Deviation Settings

Attribute	Description
Upper Baseline Limit Enabled	If enabled, the NNM iSPI Performance for Metrics uses the Upper Baseline Limit attribute value to determine the number of deviations above the average that defines the upper baseline limit. If disabled, NNMi does not define the upper baseline limit.
Upper Baseline Limit - Deviations above average	Enter the number of deviations above the average values that the NNM iSPI Performance for Metrics should use to determine the upper baseline limit.
Lower Baseline Limit Enabled	If enabled, the NNM iSPI Performance for Metrics uses the Lower Baseline Limit attribute value to determine the number of deviations below the average that defines the lower baseline limit. If disabled, the NNM iSPI Performance for Metrics does not define the lower baseline limit.
Lower Baseline Limit - Deviations below average	Enter the number of deviations below the average values that the NNM iSPI Performance for Metrics should use to determine the lower baseline limit.

After configuration is complete, the NNM iSPI Performance for Metrics starts computing the baselines by using the collected data. During the initialization period (the first three weeks), you will not see any exceptions.

Using Reports to Monitor Exceptions

The NNM iSPI Performance for Metrics generates exceptions when a metric value is found to exceed the baseline sleeve. You can view those incidents in the Incident View in the NNMi console. In addition, you can use NPS reports to monitor exceptions detected by the NNM iSPI Performance for Metrics. Ideal reports to monitor exceptions are:

- **Top N** or **Top N Chart:** Shows Top N elements that reported most exceptions
- **Chart Details:** Shows a time-series chart to indicate total count of exceptions

To monitor exceptions:

1. Log on to the NPS console.
2. From the navigation pane, launch a report of your choice.
3. Click **Options**.
4. Select one of the metrics listed in Table: Metrics for Monitoring Exceptions.
5. Click **Confirm Selection**.

Metrics for Monitoring Exceptions

Metric	Description
Baseline Exception Rate	Percentage of samples that has the exception flag set.
Baseline Exception Count	Total number of sample counts marked for an exception.
	NPS raises an exception for the metric if a certain number of samples within a sample window is out of normal range.
	For example, NPS raises a baseline exception if 7 samples out of a sample window of 10 samples are beyond the upper normal range. You must configure NNMi for the upper and lower normal ranges to enable NPS to raise the exceptions.
	If the topology filter utilization is between Upper Normal and Lower Normal, the exception count is set to 0 (Zero) for the selected sample.
	If you disable the upper or lower normal thresholds in NNMi, NPS does not raise exceptions for the upper and lower normal ranges.
	The iSPI Performance products can analyze the following using this metric:
	- Analyze the performance of the selected topology filter for a specific time range.
	- Identify network performance problem if the exception count is persistently high for the selected time range.

Using Reports to Monitor Baselines

The NNM iSPI Performance for Metrics enables you to monitor the baseline with the help of the Baseline Sleeve report to monitor baselines. This report highlights the baseline sleeve on a line graph and shows a table that lists the lower and upper limits at each display grain.

To monitor a baseline sleeve:

1. Log on to the NPS console.
2. From the navigation pane, launch the Baseline Sleeve report.
3. Click **Options**. This prompt enables you to choose one of the metrics that was configured for baselining during interface or component monitoring configuration.
4. Click **Confirm Selection**. The report shows a line graph of the selected metric and highlights a range of values that is computed as the baseline sleeve for the selected period of time.

You can also launch other reports like Top N or Chart Details to monitor different characteristics of the baseline for the monitored metric. Table: Metrics to Monitor the Baseline provides a list of metrics that help you monitor baseline characteristics.

Metrics to Monitor the Baseline

Metric	Description
Baseline Average	Average value presenting the normal (expected) range of utilization for the selected topology filter based on the historical data
	Can analyze the following using this metric:
	• The actual utilization of the topology element as compared to the normal utilization ranges
	• Forecast the topology element utilization for a selected time range.
	• Identify the over-utilized and under-utilized topology elements before they violate the thresholds.
Baseline Deviation	Deviation of the samples spread over a specific time range.
	Represents the extent of the range of values normally seen for the selected time range in the seasonal cycle, based on historical data.
	Can analyze the following using this metric:
	• Analyze the utilization pattern for the selected topology elements over a selected time range.
	• Identify the utilization values that are lower or higher than the mean utilization. This in turn enables you to identify the over-utilized and under-utilized topology elements.
Lower Threshold	The lower deviation configured in the Baseline Settings form.
Upper Threshold	The upper deviation configured in the Baseline Settings form.
Upper Normal	Upper limit of the normal range for the sample. The normal range for the sample may vary for every iSPI Performance product.
	NNMi supplies the upper normal value based on values that you enter in the Threshold Configuration form. You can disable the upper normal value if you do not require to set the upper threshold for the metric.
	See the *HPE Network Node Manager i Software Online Help for Administrators* for information about the Threshold Configuration form.

Metrics to Monitor the Baseline, continued

Metric	Description
	The value is calculated using the following formula: Upper Normal = Baseline Average + (N1 * Baseline Deviation) where N1 is the upper sleeve width
Lower Normal	Lower limit of the normal range for the sample. The normal range for the sample can vary for every iSPI Performance product. NNMi supplies the lower normal value based on values you enter in the Threshold Configuration form. You can disable the lower normal value if you do not set the lower threshold for the metric. See the *HPE Network Node Manager i Software Online Help for Administrators* for information about the Threshold Configuration form. The value is calculated using the following formula: Lower Normal = Baseline Average - (N2 * Baseline Deviation) where N2 is the lower sleeve width
Days to Threshold	Forecast for the topology element utilization to reach the upper or lower threshold limit. Represents the number of days left for the utilization to reach the threshold value. NPS calculates this metric based on the following: • Baseline Slope • The static upper and lower thresholds defined in NNMi Threshold Configuration form

Example

1. Configure a baseline threshold for nodes in the NNMi Baseline Setting configuration form.

Incidents are created if detected abnormal values persist for 15 minutes (or longer) in the last 30-minute window

   a. Select CPU 1Min Utilization as the monitored attribute.

   b. Select Threshold Enabled.

   This selection enables the NNM iSPI Performance for Metrics to send incidents to the NNMi Incident View.

   c. Type a duration of 15 minutes.

   d. Type a duration window of 30 minutes.

   Because of this configuration, the NNM iSPI Performance for Metrics will send incidents to the NNMi Incident View if an exception persists for 15 minutes of the last 30-minute window.

    e. In the Baseline Deviation Settings tab:

        i. Select the Upper Baseline Limit Enabled check box.

        ii. Type 3 for the deviation above average.

        iii. Select the Lower Baseline Limit Enabled check box.

        iv. Type 3 for the deviation below average.

> **Tip:** You can choose either the upper baseline limit or lower baseline limit or both.

2. After the initialization period (three weeks), launch the Baseline Sleeve Component Health report and view the baseline.

Time	CPU 1min Utilization (avg)	Exceptions (%)	Lower Normal	Upper Normal
12:00	25.00%	0.00%	23.64%	24.64%
12:05	24.00%	0.00%	23.68%	24.68%
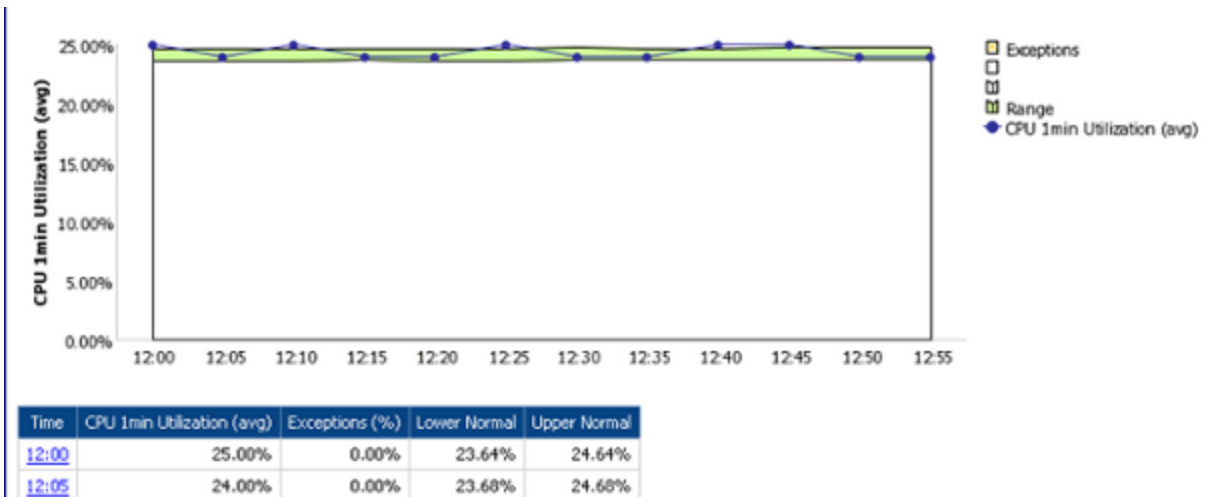
3. When the NNM iSPI Performance for Metrics detects a metric value that lies outside of the baseline sleeve for 15 minutes or longer within the last 30-minute window, the NNM iSPI Performance for Metrics raises an exception and sends an incident to the NNMi Incident View. You can view these exceptions on a Chart Detail report by selecting the `CPU 1Min Utilization - Baseline Exception Count` metric.

Display Grain : 15 Minutes

Legend:
Metric in Y1 (solid line)
CPU 1min Utilization (avg)
CPU 1min Utilization - Upper Normal (max)
CPU 1min Utilization - Lower Normal (min)

Metric in Y2 (dashed line)
CPU 1min Utilization - Baseline Exception Count (sum)
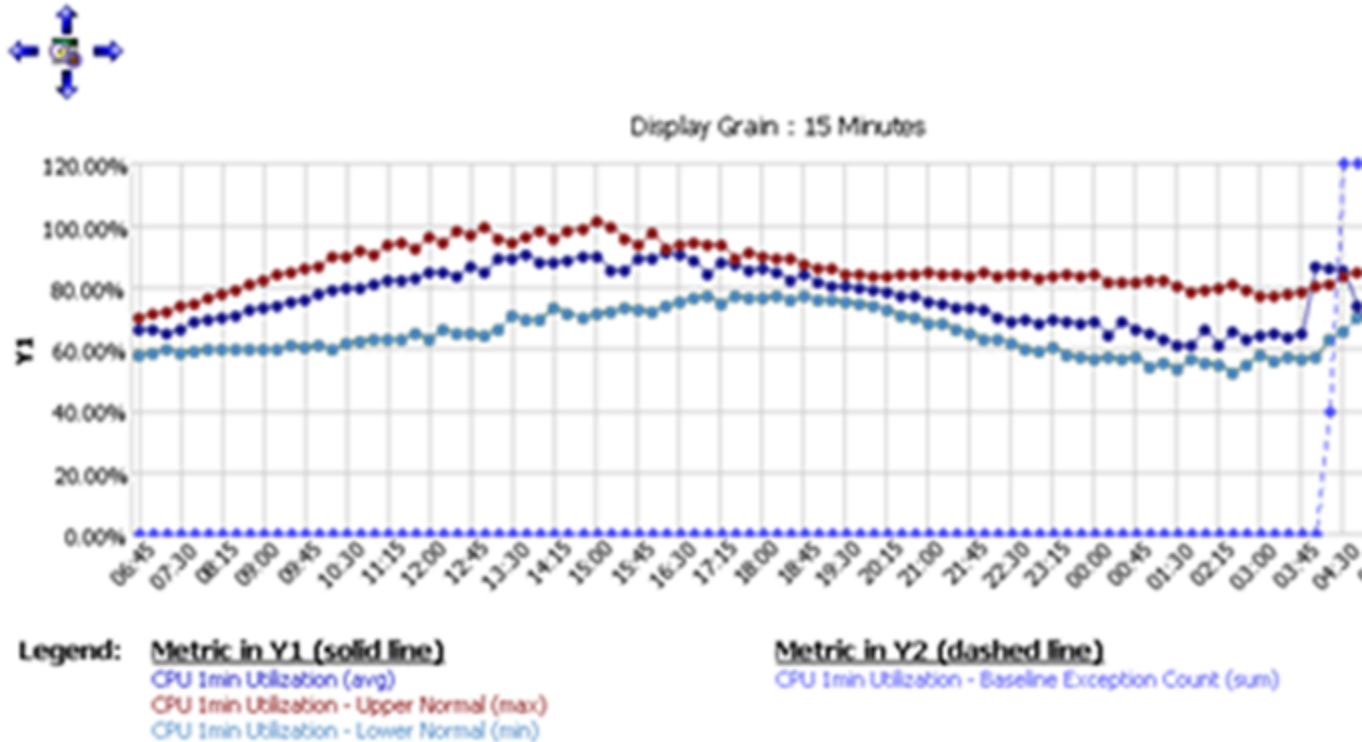
How the NNM iSPI Performance for Metrics Calculates Baselines

A baseline characterizes the normal pattern of behavior for the time series of the monitored metric instance. The upper and lower limits that define the normal operating range are derived from the baseline and deviation values. The normal operating range is also referred to as the **baseline sleeve**.When the detected metric value lies outside of the baseline sleeve, an exception is detected.

The width of the baseline sleeve is a multiplier function of the historical deviation observed in the data; that is, if historically the values do not vary much, the sleeve remains narrow; if the values vary greatly, the NNM iSPI Performance for Metricscreates a wider sleeve. The multiplier chosen can be used to tune the selectivity of the range–a multiplier of 2 captures about 95% of samples; a multiplier of 3 captures 99% of samples.

Because the baseline sleeve defines the range of values that are considered normal, the process of exception detection must define abnormal samples as the samples that lie outside of the sleeve. To reduce the number of false alarms, this process does not classify every abnormal sample as an exception. The sleeve width of 3 deviations (or, in other words, a sleeve calculated with a multiplier of 3) captures 99% of samples, which means, 1% of the samples lie outside of the sleeve.

When abnormal values persist for a long time, a new "normal " value is determined as the baseline and the baseline sleeve changes.
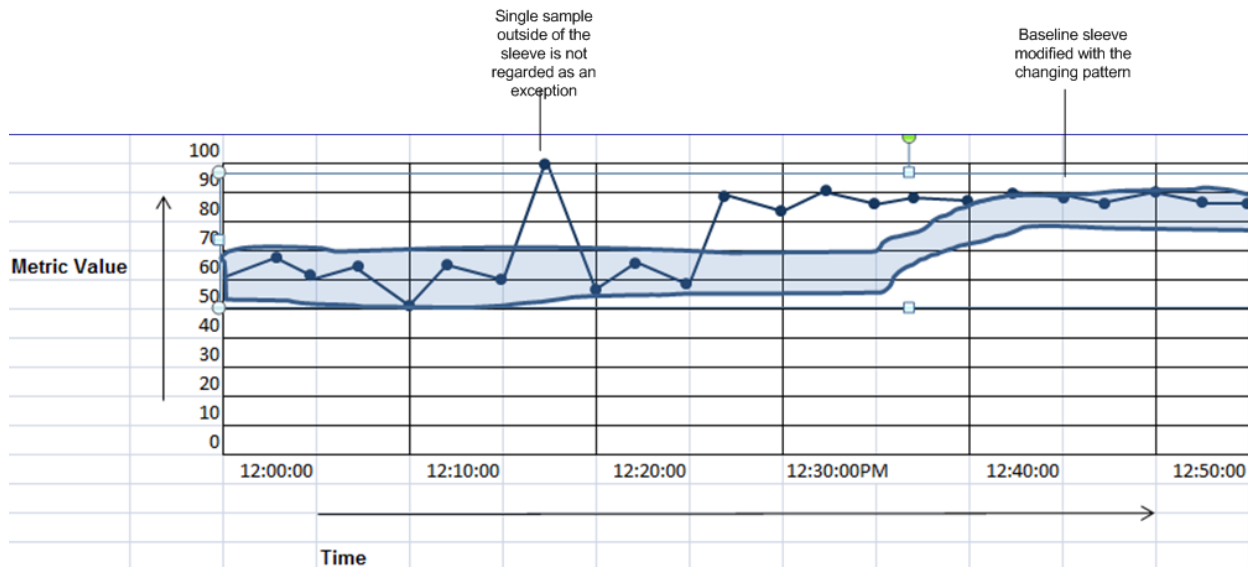
Triple Exponential Smoothing Technique

The NNM iSPI Performance for Metrics uses the Triple Exponential Smoothing technique to calculate baseline. This technique is used to predict the expected current value in a time series with the help of the previous value and the previous prediction.

$y_1, ..., y_{t-1}, y_t$ is a time series.

$\acute{y}_t$ is the predicted value at time t.

$\acute{y}_t = \alpha y_{t-1} + (1-\alpha)\acute{y}_{t-1}$

$\acute{y}_t$ is the weighted average of the previous sample in the time series and the previous prediction. The calculation of $\acute{y}_t$ is based on the principle that the most recent sample is the most relevant sample for making predictions. A sample becomes less relevant with age. The weight of a sample in calculating the prediction decays exponentially over time and eventually becomes insignificant.

The Triple Exponential Smoothing technique assumes that the time series can be broken down into three components:

- The overall mean smoothing component or **intercept** (a_t)
- The **trend** smoothing component (b_t)
- The **seasonal** smoothing component for changing seasonal patterns (c_t)

The prediction $\acute{y}_t$ is the sum of the above three components:

$\acute{y}_t = a_t + b_t + c_t$

Each component (a_t, b_t, and c_t) is calculated with the exponential smoothing technique:

- Intercept component = weighted average. of de-seasonalized sample and the previous intercept + previous trend

  $a_t = \alpha(y_t - c_{t-s}) + (1-\alpha)(a_{t-1} - b_{t-1})$

- Trend component = weighted average of the difference between the current and the previous intercepts and

the previous trend.

$$b_t = \beta(a_t - a_{t-1}) + (1 - \beta)b_{t-1}$$

- Seasonal component = weighted average of the difference between the current sample and the current intercept and the seasonal component from the previous season

$$c_t = \gamma (y_t - a_t) + (1-\gamma)c_{t-s}$$

All three parameters (α, β, and γ) control the rate at which the baseline is modified with the changing pattern.

Lower and Upper Limits of the Baseline Sleeve

The upper (u) and lower (l) limits of the baseline sleeve are calculated by adding or subtracting a defined multiple of the historic baseline deviation to or from the baseline average.

$u_t = ý_t + (\delta_+ * d_{t-S}); l_t = t - (\delta_- * d_{t-S})$

The values of δ_+ and δ_- (configured in the Baseline Deviation Settings tab) influence the width of the baseline sleeve. Smaller values of δ create a narrow sleeve; larger values of δ create a wide sleeve.

Seasonality and Grain

A short season length, for example, of 24 hours, fails to capture weekly patterns such as lighter workload on weekends. This can result in:

- Very wide baseline sleeve
- A baseline slanted toward the weekday average at the start of the weekend and a baseline slanted toward to weekend average at the start of the week

For calculating baselines, the NNM iSPI Performance for Metrics uses a seasonality of 1 week with a grain of 5 minutes, and therefore, daily, as well as weekly, patterns are captured while calculating the sleeve.

Initialization

The NNM iSPI Performance for Metrics initializes the intercept component with the value obtained from the first sample. The deviation (δ), the trend component (b), and the seasonal component (c) are initialized to zero.

The initial values influence the baseline sleeve for some time. Gradually, with the arrival of new data samples, the baseline sleeve begins to stabilize. One of the contributing factors of the seasonal component—γ—continues to change through the first three seasons, and then finally stabilizes and remain constant from the beginning of the fourth season.

Detecting Exceptions

After defining the baseline sleeve as the range of normal values, the NNM iSPI Performance for Metrics can detect abnormal samples as the ones that lie outside of the sleeve. By defining a moving window of Y samples, the NNM iSPI Performance for Metrics detects an exception if the window contains X samples (or more). The NNM iSPI Performance for Metrics uses a sample interval of 5 minutes (that is, for computing baselines, samples are read by the NNM iSPI Performance for Metrics at an interval of 5 minutes).

For example, if Y is 6 (30 minutes) and X is 3 (15 minutes), the NNM iSPI Performance for Metrics detects an exception when the abnormal behavior persists for at least 15 out of 30 minutes.

Using Report Views

A ReportView is a snapshot of a report. You can modify the default settings (like topology attribute and time range) of a report and save the modified report as a ReportView. Instead of launching the original report and modifying the settings all over again, you can simply launch the ReportView.

Create Report Views when you want to:

- View reports that take a long time to load
- Frequently view reports that require significant customization

You can also enable the commenting option on Report Views.

By default, the NNM iSPI Performance for Metrics provides a set of Report Views that are available for use after the NNM iSPI Performance for Metrics installation is complete.

To create a ReportView:

1. Log on to the NPS console.
2. Launch a report of your choice.
3. Make necessary changes to the report, if you like.
4. *Optional.* Enable the commenting option for the ReportView that you are going to create.

   To enable the commenting option on the ReportView:

   a. Right-click the report in the Reports navigation pane, and then click **Properties**.
   b. Go to the Report tab.
   c. Click **Advanced Options**.
   d. Select Enable Comments in Saved Output Versions.
   e. Click **OK**.

5. To save the report as a ReportView, click **Keep this Version > Save as Report View**. The Save as Report View dialog box opens.
6. In the Save as Report View dialog box, do one of the following:

   a. Click **My Folders** to save the ReportView in My Folders, and then click **OK**.
   b. Click Select Another Location, select a pre-existing folder, and then click **OK**.

      You can click **New Folder** to create a new folder if no folders exist.

7. Click **OK**.

You can now launch the saved ReportView from the NPS navigation panel.

You cannot change the topology filter and time settings of an existing ReportView. You can, however, launch a saved ReportView first, and then use the NPS navigation panel to modify the topology filter or time setting of the report.

Create Shortcuts to Commonly Used Reports

You can quickly launch reports from the navigation panel using shortcuts.

You can create a custom shortcut folder that contains a set of Report Views with options you select.

To create a custom ReportView:

1. Create a folder and name it anything you want; for example, *My Report Views*.
2. Launch a report with filters by clicking **Run with prompt**.
3. Create a ReportView by clicking **Keep this version** and **Save as Report View**, selecting **Select another location**, and selecting the folder you created.
4. Provide a name for the ReportView; for example, Top 10 Nodes by Volume in Last Day.
5. Repeat to create another ReportView.

Cross-Domain Report Views

While creating a cross-domain Extension Pack, NPS creates a set of predefined Report Views. To access these Report Views, click **Quicklaunch ReportViews** in the navigation pane. Cross-domain Report Views are listed with all other Report Views packaged with the NNM iSPI Performance for Metrics Extension Packs under **Quicklaunch ReportViews**.
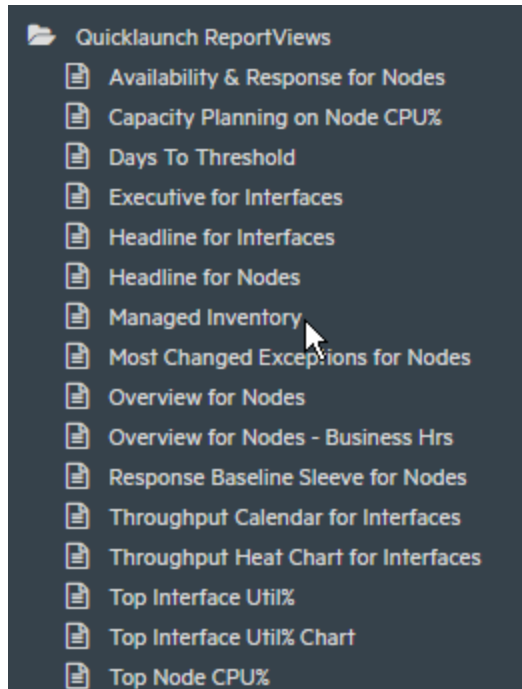
NPS creates the following Report Views when you generate a cross-domain extension pack:

- Top Problematic Items
- Days To Threshold
- Managed Inventory

Top Problematic Items

The Top Problematic Items ReportView is essentially a Top N report that shows top 10 objects discovered on the network with highest overall exception rates. The ReportView also shows the threshold exception and baseline exception rates for each object.

Additionally, the report shows the following details:

- The object type of each object
- The extension pack where the object belongs

> **Threshold Exception Rate:** The Threshold Exception Rate metric indicates the percentage of time for which the object experienced threshold exceptions.
>
> **Baseline Exception Rate:** The Baseline Exception Rate metric indicates the percentage of time for which the object experienced baseline exceptions.
>
> **Overall Exception Rate:** The Overall Exception Rate indicates the percentage of time for which the object experienced threshold exceptions or baseline exceptions or both.

For an effective analysis, use the drill-through feature on the ReportView to launch the Baseline Sleeve or Headline report to monitor baseline or threshold exceptions for an object.

Example

1. The following ReportView shows a list of objects that experienced threshold and baseline exceptions.

Rank	Object Name	Object Type	ExtensionPack Name	Overall Exception Rate (avg)	Threshold Exception Rate (avg)	Baseline Excep Rate (
1	Fa0/0 on (mplsce71.fc.usa.hp.com)	Interface	Interface_Health	100.00%	100.00%	0
2	Fa0/45 on (wanrouter-1.fc.usa.hp.com)	Interface	Interface_Health	100.00%	100.00%	0
3	Fa0/0/3 on (wanrouter-1.fc.usa.hp.com)	Interface	Interface_Health	100.00%	100.00%	0
4	Fa0/0/13 on (wanrouter-1.fc.usa.hp.com)	Interface	Interface_Health	100.00%	100.00%	0

2. To analyze threshold exceptions experienced by the object `Fa0/0 on (mplsec71)`, follow these steps:
   a. Click the object. A new NPS console opens in a new browser tab.
   b. In the new NPS console, open the Headline report of the Interface Health extension pack.
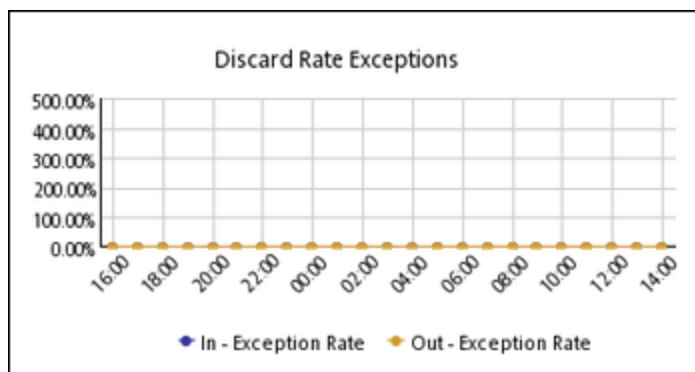
On this report, you can monitor threshold exceptions experienced by the `Fa0/0 on (mplsec71)` object.

3. To analyze baseline exceptions experienced by the object `Fa0/0 on (mplsec71)`, follow these steps:
   a. Click the object. A new NPS console opens in a new browser tab.
   b. In the new NPS console, open the Baseline Sleeve report of the Interface Health extension pack.
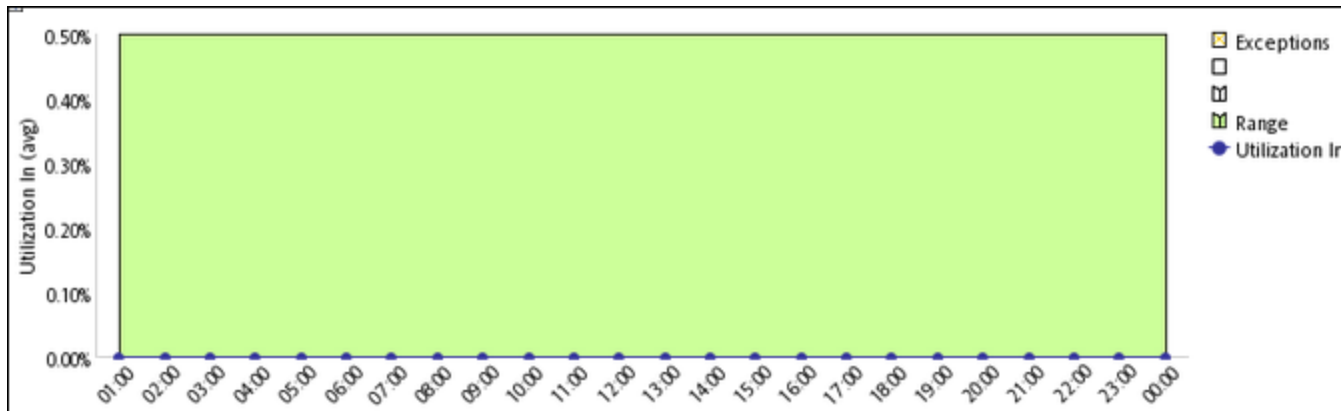
On this report, you can monitor baseline exceptions experienced by the `Fa0/0 on (mplsec71)` object.

Days To Threshold

The Days To Threshold ReportView provides a forecast analysis of threshold exceptions for all objects across all installed extension packs. This is a Top N report that shows all discovered objects on the network sorted by ascending values of the Overall Days To Threshold metric. The ReportView also shows the overall exception rate for each object.

Overall Days To Threshold: The Overall Days To Threshold metric indicates the forecasted number of days by which the object will experience threshold exception.

For an effective analysis, use the drill-through feature on the ReportView to launch the Chart Details report and select one of the forecast-specific metrics on the Chart Details report.

Managed Inventory

The Managed Inventory ReportView shows the number of instances of each topology attribute found in the collected data for all installed extension packs. This report provides a quick way of finding the count of a specific network element on the network. For example, the count of the Qualified Interface Name topology attribute shows total number of interfaces on the network.

Performance Graphing

The Performance Graphing window displays graphs that provide data about connections, nodes, and interfaces on your system.

To use this feature:

- Select a performance-polled node from the NNMi node inventory view. Click **Actions** > **HPE NNM iSPI for Performance** > **Performance Troubleshooting**. The Performance Graphing window is launched for only the selected node.

  All the interfaces of the node are also available.

- Select a performance-polled interface from the NNMi node inventory view. Click **Actions** > **HPE NNM iSPI for Performance** > **Performance Troubleshooting**. The Performance Graphing window is launched for the selected interface.

  The interface appears under the parent node in the tree. All other interfaces of the node are also visible.

If the interface is not being polled, the panel is blank.

Performance Tab in the NNMi Analysis Pane

In addition to providing the Performance Graphing console, NPS enables you to view the performance graph of a selected network element in the Analysis pane in an inventory view. The Analysis pane shows the performance graph in the Performance tab.

> **Note:** The time scale displayed on a graph represents the time on the client system where you launched the graph with a web browser and not the time on the NPS system or the NNMi management server.

Launch Performance Troubleshooting from the Incident View

In the NNMi Incident View, you can launch the Performance Troubleshooting window for the performance and fault incidents.

To launch the Performance Troubleshooting window:

1. Log on to the NNMi console as an administrator or operator.
2. Go to the Incident View.
3. Right-click an incident of your interest, and then click **HPE NNM iSPI Performance > Performance Troubleshooting**.

   Alternatively, click **Actions > HPE NNM iSPI Performance > Performance Troubleshooting**.

   The Performance Troubleshooting window opens and shows a graph that is plotted with the data from two hours before and until two hours after the incident occurred.

   For performance incidents, the Performance Troubleshooting window shows graphs of metric that is the cause of the incident.

   For fault incidents (such as node down), the Performance Troubleshooting window shows graphs of CPU and memory utilizations of the node that is the cause of the incident.

For the performance and fault incidents, NNMi adds a new Performance tab in the Analysis pane and shows performance graphs.

Access the Performance Troubleshooting Console

For operational performance management, NNMi supports an external launch to the performance analysis panel, which creates graphs that enable you to analyze the topology objects (nodes and interfaces) on your system.

To access this feature, go to the NNM console. Select either a node or an interface, open the Actions menu, select **Performance** > **Troubleshooting**.

The Performance Troubleshooting window is divided into two panes. Pane 1 shows a tree view with the root node called Network. When expanded, all the nodes on the network that were performance polled are shown. The interfaces for these nodes are also shown. Pane 2 appears blank. However, if you unselect the Quick view option, a third pane appears showing two graphs, Pre-defined and Metrics.

You can drag and drop metrics to produce graphs for a selected time period.

In this console, you can create graphs with all the metrics from all installed Extension Packs (including Custom Poller Extension Packs).

How to Find Data for a Connection in the NNMi Analysis Panel

You can find data for NNMi L2 connections through the performance analysis panel.

To use this feature, select an interface in a view in the NNMi console.

Graphs show data captured for the interfaces that comprise the endpoints of the connection over the past day.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

The following graphs are provided:

Graph 1:

- Availability (max)
- Availability (min)
- Availability (avg)

Graph 2:

- Utilization In (avg)
- Utilization Out (avg)

Graph 3:

- Error Rate (avg)
- Discard Rate (avg)

How to Find Data for an Interface in the NNMi Analysis Panel

You can find data for NNMi interfaces through the performance analysis panel.

To select an interface, start in a view in the NNMi console. The performance panel opens once you select an interface.

Graphs show data captured over the past day.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

The following graphs are provided:

Graph 1:

- Availability (avg)

Graph 2:

- Utilization In (avg)
- Utilization Out (avg)
- Utilization (avg)

Graph 3:

- Discard Rate In (avg)
- Discard Rate Out (avg)
- Discard Rate (avg)

Graph 4:

- Throughput In (avg)
- Throughput Out (avg)
- Throughput (avg)

How to Find Data for a Node in the NNMi Analysis Panel

You can find data for NNMi nodes through the performance analysis panel.

To select a node, start from the Flow Enabled Nodes view in the NNMi console. The performance panel opens once you select a node.

Graphs show data captured over the past day.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

If you select a node for which no polling is being done, the performance panel will be blank.

The following graphs are provided:

Graph 1:

- Memory Utilization (avg)
- CPU Utilization (avg)

Graph 2:

- Memory Utilization - Threshold Exception Rate
- CPU Utilization - Threshold Exception Rate

How to Find Data for a Node Fault in the NNMi Analysis Panel

You can find data for NNMi nodes through the performance analysis panel.

When NNMi detects that a node is down (not responding), a NodeDown event is generated.

To select a NodeDown incident, start from the Incident view in the NNMi console. The performance panel opens and displays several graphs.

The graphs report data for the following time period:

- Start Time: Time period of incident minus 1 hour
- End Time: Time period of incident plus 1 hour

If the plus 1 hour time period is not possible, data for the past 1 hour is shown.

To navigate forward in the graph, click >. To navigate backward, click <.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

The following graphs are provided:

Graph 1:

- Memory Utilization (avg)
- CPU Utilization (avg)

Graph 2:

- CPU Utilization - Threshold Except Rate
- Memory Utilization - Threshold Exception Rate

How to Find Data for a Node Performance Event in the NNMi Analysis Panel

You can find data for NNMi nodes through the performance analysis panel.

To select a performance incident for a performance-polled node, start from the NNMi Incident view in the NNMi console. If NNMi detects that a node is violating a configured performance metric threshold—for example, CPU utilization—a CPUOutofRangeOrMalfunctioning event is generated and made visible in the Incident view. When you click this incident, a performance tab appears in the analysis pane. When you click the tab, the graphs described below appear.

Supported events are:

- CPUOutofRangeOrMalfunctioning
- BufferOutOfRangeOrMalfunctioning
- MemoryOutOfRangeOrMalfunctioning

The graphs report data for the following time period:

- Start Time: Time period of incident minus 1 hour
- End Time: Time period of incident plus 1 hour

If the plus 1 hour time period is not possible, data for the past 1 hour is shown.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

If you select a node for which no polling is being done, the performance panel will be blank.

The following graphs are provided:

Graph 1:

- Memory Utilization (avg)
- CPU Utilization (avg)

Graph 2:

- CPU Utilization - Threshold Exception Rate
- Memory Utilization - Threshold Exception Rate

Graph 3:

This graph is shown only if there is a Buffer incident.

- Buffer Utilization
- Lower and Upper Thresholds for Buffer Utilization

Graph 4:

This graph is shown only if there is a Buffer incident.

- Buffer Hits and Misses

How to Find Data for an SNMP Agent in the NNMi Analysis Panel

You can find data for NNMi nodes that host an IP address through the performance analysis panel.

To select an SNMP agent, start from the SNMP Agent Inventory view in the NNMi console. The performance panel opens.

Graphs show data captured over the past day.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

You can refresh the graphs with the Refresh button.

The following graphs are provided:

Graph 1:

- ICMP response time
- SNMP response time

Graph 2:

- CPU Utilization (avg)
- Memory Utilization (avg)

Graph 3:

- CPU Utilization - Threshold Exception Rate
- Memory Utilization - Threshold Exception Rate

Graph 4:

- Node Availability (avg)
- Node Reachability (avg)

How to Find Connection Data for a Node in the NNMi Analysis Panel

You can find connection data for NNMi nodes through the performance analysis panel.

To select a node, start from the Flow Enabled Nodes view in the NNMi console. The performance panel opens once you select a node.

Graphs show data captured over the past day.

To navigate forward in the graph, click >. To navigate backward, click <.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

If you select a node for which no polling is being done, the performance panel will be blank.

The following graphs are provided:

Graph 1:

- Memory Utilization (avg)
- Memory Utilization - Threshold Exception Rate

Graph 2:

- CPU Utilization (avg)
- CPU Utilization - Threshold Exception Rate

Graph 3:

- Node Backplane Utilization (avg)
- Node Backplane Utilization - Threshold Exception Rate

Graph 4:

This graph is shown only if the node is capable of exporting Traffic data (netflowv5, Netflowv9, SFlowv5, IPFIX).

- Volume - In Bytes (sum)
- Volume - Out Bytes (sum)

How to Find Interface Data for NNMi Ports in the NNMi Analysis Panel

You can find data for NNMi ports through the performance analysis panel.

To select an interface, start in a view in the NNMi console. The performance panel opens once you select a port from the Ports inventory.

Graphs show data captured over the past day.

You can adjust the time period shown by clicking the arrow on the top right-hand side. This opens the time navigation panel. Any change you make is reflected in all of the graphs.

The following graphs are provided:

Graph 1:

- Availability (max)
- Availability (min)
- Availability (avg)

Graph 2:

- Utilization Out (avg)
- Utilization In (avg)

Graph 3:

- Error Rate (avg)
- Discard Rate (avg)

Graph 4:

- Throughput In (bps) (avg)
- Throughput Out (bps) (avg)

How to Find Data for an Interface Performance Event in the NNMi Analysis Panel

You can find data for interfaces through the performance analysis panel.

When NNMi detects that an interface is violating a configured performance metric threshold (for example, utilization), an InterfaceUtilization event is generated and shown in the Incident view. When you click this incident, the performance analysis panel opens.

The graphs report data for the following time period:

- Start Time: Time period of incident minus 1 hour
- End Time: Time period of incident plus 1 hour

If the plus 1 hour time period is not possible, data for the past 1 hour is shown.

The following types of incidents are reported:

- InterfaceFCSWLANErrorRateHigh
- InterfaceFCSLANErrorRateHigh
- InterfaceInputDiscardRateHigh
- InterfaceOutputDiscardRateHigh
- InterfaceInputUtilizationHigh
- InterfaceOutputUtilizationHigh
- InterfaceOutputErrorRateHigh
- InterfaceInputErrorRateHigh

The following graphs are provided for each incident:

- Interface Utilization
- Interface Availability

- Interface Throughput
- Interface Discard and Error Rate

Custom Reportlet Dashboards

You can create reportlet dashboards based on multiple reportlets from the same or different Extension Packs.

Say that you needed to monitor the following:

- Hourly utilization of Interface A on Component X
- Average normal CPU utilization for Component X
- Number of days left for memory utilization to reach threshold

You could create a custom dashboard that combines the following reportlets:

- Interface Health Heat Chart Report for the metric Utilization (avg)
- Component Health Simple Chart Report for the following metrics:
  - Primary metric: Memory Utilization - Baseline Average (avg)
  - Secondary metric:CPU 5min Utilization - Days to Threshold (min)

Create a Custom Reportlet Dashboard with the iSPI Performance BI Server Portal

Task 1: Log in to NPS

Task 2: Copy bookmarks for the reportlets to include in the dashboard

1. On the NPS home page, click the **Reports** tab.
2. Select an Extension Pack; for example, **iSPI Metrics** > **Interface_Health** > **InterfaceMetrics > Reportlets**.
3. Select a reportlet; for example, Heat Chart reportlet.
4. Click **Show Bookmark**.
5. Copy the bookmark and paste it in a text editor.

Repeat these steps for all the reportlets you want included in the reportlet dashboard.

Task 3: Create the reportlet dashboard using the BI Portal

1. On the NPS home page, click the **BI Server** tab.
2. Click **Portal**.
3. Select a location for your reportlet dashboard; for example, My Folders.
4. Click ▣**New Page** in the top right corner.
5. In the Specify a name and description-New page wizard form, specify the required information about the dashboard.
6. In the Set columns and layout-New page wizard form, specify:

- Number of columns you require in the dashboard

- Width of the columns. For a multi-column dashboard, the columns must be the same width.

7. In the Set columns and layout-New page wizard form, click **Add**.

8. Select **Cognos Utility**.

9. In the Available entries column, select a utility; for example, to display multiple reportlets in the dashboard, select **HTML Viewer**.

10. Click [➡] **Add** to move the selected utilities in the Selected entries column.

11. Click **OK**.

12. Click **Next**.

13. In the Set page style-New page wizard, specify the title, style, and any instructions you want the reportlet dashboard to show.

    Do not select the following:

    - Hide title bars

    - Hide Edit buttons on the title bar

14. Click **Next**.

15. In the Select an action-New page wizard form, select the required options and click **Finish**.

The new dashboard opens in Edit mode.

Task 4: Edit the dashboard to display the selected reportlets

1. Select any of the following tools:

    - [icon] **Edit** in the title bar

    - [icon] **Edit** in the Actions column

2. Specify the content for the reportlet dashboard using the reportlet bookmark. For example, if you selected HTML Viewer to view the reportlets, copy the bookmark you saved in the text editor and paste it in the HTML content field.

3. Specify the height for the panel where the BI Portal displays the reportlet.

4. Click **OK**.

> **Note:** Reportlets Dashboards will not work if you make any changes to the NNMi FQDN, NPS FQDN, NPS port, or the communication protocol (HTTP and HTTPS). After upgrading NPS to a new version, or restoring data from a different NPS system, reportlet dashboards stop working.
>
> Also, after NNMi fails over to the standby server in an Application Failover environment, all reportlet dashboards stop working.
>
> You must regenerate all reportlet dashboards in the above cases.

HPE iSPI Performance BI Server Portal

The HPE iSPI Performance BI Server Portal enables you to perform the following tasks:

- Open **Public Folders** and access reports from **BI Server Connection**.
- Open **My Reports** and view reports with the filter preferences you saved.
- Set up the following preferences:
  - Set up the display parameters for the reports.
  - Select the report format
  - Set up a schedule for running summary reports.

    Note: To be able to schedule reports in the NPS console, you must log on to NNMi as an administrator or Level 2 operator.

  - Set up email report delivery.
  - Modify BI Server tuning parameters for large multi-user installations.
  - Select language for your reports
- Create custom reports using Query Studio.

Launch the iSPI Performance BI Server Portal

To launch the iSPI Performance BI Server Portal:

1. Log in to NPS as Administrator.
2. On the navigation panel on the NPS home page, click **BI Server**.
3. Click **Portal**.

Optionally, you can start the iSPI Performance BI Server Portal as follows:

1. Go to http://*<localhost>*:9300/p2pd/NPS.html.
2. On the login page, enter your ErsAdmin user name and password.

Log on to the iSPI Performance BI Server Portal as ErsAdmin

If you are unable to launch NNMi, you can reach the iSPI Performance BI Server Portal by going to `http://<localhost>:9300/p2pd`.

When the login page opens, type your ErsAdmin username and password.

HPE recommends changing the default password after you install NPS. To change the password, run the `changeBIpwd.ovpl` command.

If you cannot launch NNMi, contact Support.

If you can open the iSPI Performance BI Server Portal, you can confirm that the BI Server is running.

Change Report Format Preferences

The default report format used by NPS and the BI Server is HTML.

You can set the iSPI Performance BI Server Portal preferences and enable NPS to generate reports in any of the following formats:

- PDF
- Microsoft Excel 2007
- Microsoft Excel 2007 data
- Microsoft Excel 2002
- Delimited text (CSV) *[Comma Separated Value (CSV) Files: displayed as Delimited Text (CSV) on the iSPI Performance BI Server Portal]*
- XML

To change the format of the reports generated by NPS:

1. Log in to NPS home page with administrator privileges.
2. In the navigation panel, click the **BI Server** tab.
3. Click the **Portal** link.
4. In the iSPI Performance BI Portal form, click **My Area Options**.
5. Select **My Preferences**.
6. In the General tab, select a Report Format. (The default format is HTML.)
7. Click **OK**.

The new report format becomes the default report format for your user account. The BI Portal applies the new report format to all reports generated from your user account.

Set the Same Schedule for Multiple Reports

To set the same schedule for multiple reports, you must create a job. A job identifies a collection of reports, report views, or queries that are scheduled together and share the same settings.

To set the same schedule for multiple reports:

1. Log on to the BI Server portal (in the NPS console, click **BI Server**, and then click **Portal**).
2. Click the **New Job** icon on the toolbar.
   The New Job wizard opens.
3. Specify a name and location for the job.
4. In the Select the steps–New Job wizard form, click **Add** to add the job steps.
   You can submit the steps in sequence or all at once. To modify the sequence of steps, click **Modify the sequence** and rearrange the sequence of steps.

5. In the Select an action-job form, specify whether you want to save and run, schedule, or only save the new job.

6. Click **Finish**.

Creating a New Folder in the BI Portal

To create a new folder in the BI Portal:

1. In the HPE NNM iSPI Performance BI Portal, click **New Folder**.

2. In the New Folder Wizard form, specify the name, description, and location for the new folder.

3. Click **Finish**.

BI Server Documentation

To access the BI Server documentation, click **Help** on the BI Server Portal.

Auditing

By default, the NNM iSPI Performance for Metrics tracks a history of the following:

- Sign-in and sign-out activity for each NNM iSPI Performance for Metrics user.
- Starting and stopping of processes. Examples include starting the BI server, stopping Sybase ASE database.

The NNM iSPI Performance for Metrics auditing is enabled by default. Audit information is written to a new audit log file every day. The audit log files reside in the following directory:

`<OVDataDir>/nmsas/perfspi/log/audit-<yyyy-MM-dd>.log`

To see the most recent audit report:

In the NNMi console menu bar, select **Tools** → **NPS Audit Log**.

> **Note:** The audit log is available from **Tools** → **NPS Audit Log** only if NPS is installed on the NNMi management server.

The log provides a variety of information about the recent account activity.

As an administrator, you can configure the following:

- "Disabling Auditing" below
- "Specifying the Retention Period of Audit Logs" on the next page

Disabling Auditing

To disable auditing:

1. Open the following file:

   - *On Windows:* `%ovdatadir%\NNMPerformaceSPI\rconfig\auditlog.prop`

   - *On Linux:* `/var/opt/OV/NNMPerformaceSPI/rconfig/auditlog.prop`

2. Locate the `audit.log.enabled` property, and then set the property to `false`.

3. Save your changes.

4. Restart the BI Server by running the following commands:

   a. **stopBI.ovpl**

   b. **startBI.ovpl**

Specifying the Retention Period of Audit Logs

By default, the NNM iSPI Performance for Metrics retains each archived audit log file, one per day, for 14 days.

To change the retention period of the archived audit log file:

> **Note:** This number does not affect the current day's audit log file.

1. Open the following file with a text editor:

   a. *On Windows:* `%OVDataDir%\NNMPerformanceSPI\rconfig\auditlog.prop`

   b. *On Linux:* `/var/opt/OV/NNMPerformanceSPI/rconfig/auditlog.prop`

2. Locate the `audit.log.maxretention` property:
   `audit.log.maxretention=14`

3. Modify the number to include the number of days the NNM iSPI Performance for Metrics must retain each audit log file. For example, to change the number of days to one week, enter:
   `audit.log.maxretention=7`

4. Save your changes.

5. Restart the BI Server by running the following commands:

   a. **stopBI.ovpl**

   b. **startBI.ovpl**

Troubleshooting

This section helps you troubleshoot problems in launching or viewing reports and provides you with information to help you avoid problems from occurring. See "Errors and Warnings Glossary" below for a list of error messages and warnings. To troubleshoot specific problems, see all other topics in this section.

Errors and Warnings Glossary

You can open a reference list of commands by going to the command index.

A	B	C	D	E	F	S	General

Message	Type	Cause/Suggested Actions
A		**Top**
Attempting to extend allocated space for DBSPACE {DB_SPACE} using dbsize.ovpl	WARN	The database requires additional disk space. The ETL process will allocate more.
Automatic NNM HA add-on configuration failed: Exit Code {EXIT_CODE}	ERROR	High availability configuration failed.
C		**Top**
Cannot extract remote nnm hostname from value of PRSPI_NNMDIR ({PRSPI_NNMDIR})	ERROR	On a dedicated server installation, the configuration entry for the NNMi shared directory should contain the name of the NNMi server. Failure to determine this hostname affects application failover configurations.
Cannot find the OV Install Dir <{DIR}>.	ERROR	*%nnminstalldir%* does not exist. Check the NNMi installation.
Cannot find the nnm.ports.properties file to establish port for jboss http server.	ERROR	If the jobs port cannot be determined, the iSPI cannot be enabled. Check NNMi installation.
Check NNM HA environment, and then run nnmenableperfspi.ovpl.	ERROR	Enabling the iSPI in an HA environment must be done in a specific order. See the HA-specific topic in the *NNMi Deployment Guide* for more information.
Check your environment for NNM environment variables and run	ERROR	Some or all required

Message	Type	Cause/Suggested Actions
nnmenableperfspi.ovpl.		NNMi variables cannot be determined. Check NNMi installation.
Command Failed: {PERL} getNNMActiveScript: {CHILD_ERROR}	ERROR	An application failover specific script failed. Check your application failover environment.
Command Failed: {UNZIP} {ZIPFILE}: {CHILD_ERROR}	ERROR	An error occurred unzipping the identified file. The CHILD_ERROR might give more information (corrupt zip file, permissions issue, etc.).
Command: {OVDIR}/misc/nnm/ha/nnmhaclusterinfo.ovpl -config NNM -get HA_PERFSPI_ADAPTER_DIR	ERROR	A high availability specific script failed. Check your high availability environment.
Command: {OVDIR}/misc/nnm/ha/nnmhaclusterinfo.ovpl -config NNM -get NNM_INTERFACE	ERROR	A high availability specific script failed. Check your high availability environment.
Command: {OVDIR}/misc/nnm/ha/nnmhaconfigure.ovpl NNM -addon PerfSPIHA	ERROR	A high availability specific script failed. Check your high availability environment.
Configuration file contains errors.	ERROR	The iSPI configuration file contains errors. Check the associated error messages for specifics.
Configuration: Unable to access directory {PRSPI_METRICS}	ERROR	The shared directory for metrics is inaccessible. Check that NFS or Samba is

Message	Type	Cause/Suggested Actions
		operating properly.
Configuration: Unable to access directory {PRSPI_NNMDIR}	ERROR	The main shared directory is inaccessible. Check that NFS or Samba is operating properly.
Configuration: Unable to access directory {PRSPI_TOPODUMP}	ERROR	The shared directory for topology information is inaccessible. Check that NFS or Samba is operating properly.
Configuration: Directory {PRSPI_METRICS} must have write access	ERROR	The shared directory for metrics is not writeable. Check your NFS or Samba configuration and file permissions.
Configuration: Directory {PRSPI_TOPODUMP} must have write access.	ERROR	The shared directory for topology information is not writeable. Check your NFS or Samba configuration and file permissions.
Configuration: java not found in 'java_dir' ({java_dir})	ERROR	The HPE Java installation appears to be corrupted.
Configuration: perl not found in 'perl_dir' ({perl_dir})	ERROR	The HPE Perl installation appears to be corrupted.
Configuration: {NAME} not defined.	ERROR	A required configuration parameter is not defined.
Configuration: {NAME} value must be a valid path. Invalid value: {VALUE}	ERROR	A required configuration parameter must be a valid path and it is not.

Message	Type	Cause/Suggested Actions
Configuration: {NAME} value must be numeric. Invalid value: {VALUE}	ERROR	A required configuration parameter must be numeric and it is not.
D		**Top**
DB Processes still not running...restart attempt failed.	ERROR	The database server process stopped while ETL was running. ETL attempted to restart the process but failed to do so. ETL cannot continue. Try starting the database server process manually using startDB.ovpl
DBSPACE {DBSPACE}, FILE {FILE}: Cannot extend because there is no reserve space left.	WARN	The database needs to be extended but cannot because there is no disk space available. You should free existing disk space or add new disks before your database runs out of space.
DBSPACE {DBSPACE}, FILE {FILE}: Cannot extend {PATH} because there is insufficient disk space left on the file system.	WARN	The database needs to be extended but cannot because there is no disk space available. You should free existing disk space or add new disks before your database runs out of space.
Database Processes not running!	WARN	The database server stopped while ETL was running.
Dropping connection_id {CONNECTION_ID}	WARN	The specified connection is being dropped. This is

Message	Type	Cause/Suggested Actions
		normally a transparent operation for report users.
E		**Top**
Error - Unable to remove the nnm actions for iSPI for Performance	WARN	Check that NNM is running. You can also remove actions from the NNMi console.
Error running dbisql: {DB_ERROR} ({SQL})	WARN	The described SQL error occurred. Check for other messages associated with this one.
Expected PRSPI_NNMDIR of form /net/hostname/var... or //hostname/PerfSpi	WARN	On a dedicated server install, the NNM directory path should be one of the forms shown.
Extension Pack {EXTENSION_PACK}: process not running. Restarting	WARN	An ETL process had to be restarted.
F		**Top**
Failed archiving logfile PerfSPI_Diagnostic_Metrics.log	WARN	The log archive process failed.
Failed archiving logfile prspi.log	WARN	The log archive process failed.
Failed building Extension Pack for new or updated NNM metadata: {METADATA}	ERROR	This should not be encountered by customers.
Failed chdir to {DIR}: {OS_ERROR}	ERROR	Failed changing directory to the specified directory. The OS error might give more information (directory does not exist, permissions, etc.).
Failed command:copy({FILE_LIST}):{OS_ERROR}	WARN	Failed copying the

Message	Type	Cause/Suggested Actions
		specified files. The OS error might give more information (directory does not exist, permissions, etc.).
Failed command:copy({FILE_LIST}):{OS_ERROR}	ERROR	Failed copying the specified files. The OS error might give more information (directory does not exist, permissions, etc.). The program immediately exits.
Failed command:fork():{OS_ERROR}	ERROR	Failed forking (creating a new process). The OS error might give more information (lack of system resource, etc.). The program immediately exits.
Failed command:mkpath({PATH_LIST}):{OS_ERROR}	ERROR	Failed creating the specified directories. The OS error might give more information (directory does not exist, permissions, etc.). The program immediately exits.
Failed command:mkpath({PATH_LIST}):{OS_ERROR}	WARN	Failed creating the specified directories. The OS error might give more information (directory does not exist, permissions, etc.).
Failed command:move({FILE_LIST}):{OS_ERROR}	ERROR	Failed moving the specified files. The OS error might give more information (directory does not exist, permissions,

Message	Type	Cause/Suggested Actions
		etc.). The program immediately exits.
Failed command:move({FILE_LIST}):{OS_ERROR}	WARN	Failed moving the specified files. The OS error might give more information (directory does not exist, permissions, etc.).
Failed command:open({HANDLE},{ARGS}):{OS_ERROR}	ERROR	Failed opening the specified file. The arguments and OS error might give more information (directory does not exist, permissions, etc.). The program immediately exits.
Failed command:open({HANDLE},{FLAG},{ARGS}):{OS_ERROR}	ERROR	Failed opening the specified file. The arguments and OS error might give more information (directory does not exist, permissions, etc.). The program immediately exits.
Failed command:opendir(DIR,{DIR Configuration: {NAME} not defined.NAME}):{OS_ERROR}	WARN	Failed opening the specified directory. The OS error might give more information (directory does not exist, permissions, etc.).
Failed command:system({COMMAND}:{CHILD_ERROR})	WARN	Failed running the specified system command. The OS error might give more information (lack of system resources etc.).
Failed command:system({COMMAND}):{CHILD_ERROR}	ERROR	Failed running the

Message	Type	Cause/Suggested Actions
		specified system command. The OS error might give more information (lack of system resources etc.). The program immediately exits.
Failed command:unlink({FILE_LIST}):{OS_ERROR}	ERROR	Failed removing the specified files. The OS error might give more information (directory does not exist, permissions, etc.). The program immediately exits.
Failed command:unlink({FILE_LIST}):{OS_ERROR}	WARN	Failed removing the specified files. The OS error might give more information (directory does not exist, permissions, etc.).
Failed copying {SRC} to {DEST}: {OS_ERROR}	ERROR	Failed copying file from specified source to destination. The OS error might give more information (directory does not exist, permissions, etc.).
Failed creating {FILE}: {OS_ERROR}	ERROR	Failed creating the specified file. The OS error might give more information (directory does not exist, permissions, etc.).
S		**Top**
SEVERE: The web application [/p2pd] created a ThreadLocal with key of type [com.cognos.cm.diagnostics.AbstractCounter$ThreadLocalCounter$1]	ERROR	Ignore this error.
General Warnings		**Top**

Message	Type	Cause/Suggested Actions
⚠ <Package name>	WARN	This warning can appear for any one of the following reasons: • There is no data in the topology tables for the displayed package. • The topology filters chosen will result in empty reports. • The topology filters chosen are not appropriate for the selected package and will be dropped.

Reports take too long to open

The problem could be due to one of the following.

- Did NNMi experience a recent topology reset?

  Check the current NNMi topology:

  a. Log in to NNMi console using your username and password.

  b. Click **Topology Map**.

  c. Click **Network Overview**. This view is updated frequently when the topology goes through a change.

- Were one or more interface groups recently added to performance monitoring?

  Check the interface groups:

  a. Log in to the NNMi console using your username and password.

  b. Click **Inventory**.

  c. Click **Interface Groups**.

  The Interface Group-Interface Groups form displays the interface groups available in the network.

- Are large numbers of users logged in to the system?

  Enable more parallel BI Server processes. Follow the recommendations in the BI Server tuning guides.

- Are the hardware resources over-utilized?

  Check if NNMi generated any resource-related incidents:

  a. Log in to the NNMi console using your username and password.

  b. Click **Incident Management**.

  c. Click **Open Key Incidents** to check for any resource-related incidents.

Check hardware statistics, especially aggregate CPU usage.

Verify that disk speed is adequate. For large networks, a RAID5 8-spindle is recommended.

General Resolution Tips

- Create a report view with the required prompts and schedule this report view to run on a regular basis. View the scheduled output, or bookmark a URL to view the scheduled output. For more information on scheduling reports, see Schedule Delivery of Reports.

- Filters can improve report performance. A report on a specific node group is generally faster than a report on All Nodes/Interfaces. The same logic applies to the time dimension. For more information on report filters, see the following topics:

  - Set Topology Filters

  - Change Time Controls

- Run the same report again. Data caching, which took place the first time you ran the report, will decrease the run time the next time you run the same report.

Reports do not open in a GNM environment

In a GNM environment where each regional manager is configured with an NPS and the global manager is also configured with a dedicated NPS instance, the following error message may appear when you try to launch a report:

```
DPR-ERR-2107
```

```
The user capabilities cookies can not be decoded
```

You are likely to see this error message when you do the following:

1. Log on to the NPS console associated with the global manager.
2. Log on to the NNMi console of one of the regional managers.
3. Launch the NPS console from this regional NNMi console.
4. Launch one of the reports in the NPS console.

This error appears when the Regional Manager and Global Manager are in different domains.

Follow these steps on all NPS systems available in the GNM environment (In a distributed deployment of NPS, do this only on the UiBi Server): Log on to the NPS system as root or administrator.

1. Run the following command:

   - *On Windows:* **%NPSInstallDir%\bin\runBIConfigGui.ovpl**

   - *On Linux:* **/opt/OV/NNMPerformanceSPI/bin/runBIConfigGui.ovpl**

2. In the HPE NNM iSPI Performance BI Configuration window, click **Actions > Edit Global Configuration**.

3. In the Global Configuration dialog box, specify the following values:

   - Domain: specify the FQDN of the NPS system (or the FQDN of the UiBi Server in a distributed deployment of NPS).

- Path: specify /.

- Secure flag enabled: Select yes.

4. Click **OK**.

5. Restart the BI Server by running the following commands:

   a. **stopBI.ovpl**

   b. **startBI.ovpl**

You must log in again when navigating from NNMi to an iSPI Performance product report

This can occur if you launch the NNMi console using the IP address of the NNMi management server and not the fully qualified domain name (FQDN). Always use the FQDN of the NNMi management server to log in to the NNMi console.

Whenever you navigate from NNMi to a report, account information is passed between your NNMi login session and NPS using cookies in your browser. A function known as Single Sign-On (SSO) makes it possible for you to move from NNMi to a report without having to log in a second time.

For SSO to operate, the following conditions must be met:

- The NNMi system and the NPS system must each be configured with an FQDN.

- You must use the FQDN in your browser (not "localhost" or the machine name minus the domain) when you launch the NNMi console.

An FQDN has the following format:

http://fully_qualified_domain_name:port/nnm

NNMi Console Action menu has no link to the Reporting–Report menu

If you do not completely (and successfully) run the enablement script, the NNMi console does not display the **Reporting-Report menu** item in the **Action** menu.

See the instructions in the *Installation Guide* for running the enablement script (`nnmenableperfspi.ovpl`) on the NNMi management server.

NNMi is not polling one or more objects for performance data

The problem could be due to either of the following.

- Is NNMi polling the specified devices for performance data?

  a. Navigate to the Monitoring Configuration form in NNMi and check the polling attributes.

  b. Examine the polling attributes for the device and any group the device belongs to.

- Can NNMi synchronize topology information with NPS?

  NNMi and NPS share topology information. Most of the time, this information is fully synchronized. However, certain conditions will cause the iSPI Performance's view of the network topology to vary slightly from NNMi's view. A lack of topology synchronization can occur under the following conditions:

  - NNMi is running discovery.

  - NNMi rediscovers an existing network.

  - NNMi is adding or removing nodes, interfaces, or groups from polling policies.

  Should any of these occur, you will see:

  - An error indicating that NNMi is not polling targeted objects for performance data

  - "Unmanaged" next to nodes and interfaces in the Path Health report

Leaving a report open requires re-login

NPS requires you to re-login to view a report under the following conditions:

- You are trying to access NPS without being authenticated by NNMi.
- You pointed your browser at NPS without first logging in to NNMi.
- Your browser session has expired.

You can resolve this by loggin in again to return to NPS.

Leaving the browser open for a while prevents re-login

The inactivity timeout setting for NPS is 1 hour. If your browser is inactive for 1 hour, your session expires. You must log in again. The login window requires the administrator username and password.

1. Close the timeout login window.
2. Log in to the NNMi console using the standard URL:

   http://fully-qualified-hostname:port;/nnm
3. Navigate to the iSPI Performance report you were using.

Exception counts are missing from reports

Did you define thresholds in NNMi?

Reports will not track exception counts unless high and low thresholds are set in NNMi.

If reports are not showing exception counts, the NNMi administrator did not set thresholds.

Setting up thresholds in NNMi requires administrator privileges. To locate applicable help topics:

1. Launch the NNMi console.
2. Click **Help** > **Help for Administrators**.

3. Select **Configure Threshold Monitoring for Nodes**

Immediately after thresholds are set, exceptions will begin registering in reports.

NPS stops functioning and does not show any reports

Did you change the address or hostname of NPS?

If, after installing NPS, you modify the IP address, hostname, or MAC address, you cannot restart NPS. As a result, you cannot view any reports.

To change the IP address, hostname, or MAC address of the NPS server:

1. On the NPS server, run the following command:

   Windows:

   %NPSInstallDir%\NNMPerformanceSPI\bin\stopAll.ovpl

   Linux:

   /opt/OV/NNMPerformanceSPI/bin/stopALL.ovpl

2. Open the BI Server configuration GUI with the following command:

   Windows:

   %NPSInstallDir%\NNMPerformanceSPI\bin\runBIConfigGUI.ovpl

   Linux:

   /opt/OV/NNMPerformanceSPI/bin/runBIConfigGUI.ovpl

3. Click **File** > **Export As**.

4. Specify a location and file name in the dialog box and click **Save**.

   The configuration data is saved in XML format.

5. Manually delete the following directories:

   Windows:

   %NPSInstallDir%\nonOV\cognos\bi\configuration\signkeypair

   %NPSInstallDir%\nonOV\cognos\bi\configuration\encryptkeypair

   Linux:

   /opt/OV/nonOV/cognos/bi/configuration/signkeypair

   /opt/OV/nonOV/cognos/bi/configuration/encryptkeypair

6. Copy the XML file saved in step 4 to the following directory:

   Windows:

   %NPSInstallDir%\nonOV\cognos\bi\configuration

   Linux:

   /opt/OV/nonOV/cognos/bi/configuration

Navigation Panel Does not Show Reports

The navigation panel in the NPS console does not show any reports.

To resolve this issue, follow these steps:

1. Go to the Reports tab in the navigation panel and click **Refresh**.

2. If reports still do not appear, log on to the NPS system, and then run the following commands to restart the BI Server (in a distributed deployment of NPS, run these commands on the UiBi Server):

   stopBI.ovpl

   startBI.ovpl

3. If reports still do not appear, log on to the NPS system, and then run the following commands (in a distributed deployment of NPS, run these commands on the UiBi Server):

   stopALL.ovpl

   startALL.ovpl

Performance Tab in the Analysis Pane is Blank

Symptom: The Performance tab in the Analysis pane of an inventory view does not show any data when NPS is configured to use the HTTPS protocol or after upgrading NPS from a previous version.

Solution: To resolve this issue, follow these steps:

1. In the inventory view , right-click the object, and then click **HPE NNM iSPI Performance > Performance Troubleshooting**. A new browser window opens warning you of a security exception.

2. Accept the security exception in the browser. Hereafter, performance graphs appear in the Performance tab in the Analysis pane.

Also, restart the BI Server by running the following commands on the NPS system:

stopBI.ovpl

startBI.ovpl

Report shows gaps in displayed data

This can happen when there is a conflict between the polling interval and the display grain selected for the report.

When you select a finer grain for your report than the polling interval selected in NNMi, the iSPI Performance displays all the intermediate time entries and the data available for them. If these intermediate time entries do not contain any data sample, the reports might display missing entries in tables or gaps in the graphs and charts.

Resolution

Set a greater display grain for the report than the polling interval defined for the device in NNMi. For example, if the polling interval for the interface is configured as 1 minute, make sure that the display grain for the reports related to the interface is at least 1 minute.

No Data in the Node Performance or Interface Performance View

Symptom: The Node Performance and Interface Performance views in the Performance Analysis workspace are empty.

Cause: The NNMi management server and NPS system are configured with different time zone settings.

Solution: Make sure that the NNMi management server and NPS system have the same time zone setting.

Report shows no data

Cause 1: Problem with Extension Packs

If you are not able to see any data on all the reports in an extension pack and if you see the `ETL.<extension_pack> Starting ETL processor` message repeatedly in the `prspi.log` file, verify that the extension pack is installed on the system correctly.

If the extension pack is not installed on the NPS system successfully, reinstall the extension pack.

> **Note:** While installing an extension pack, the iSPI installer places a copy of the extension pack into a local directory on the NPS system. Therefore, you do not need the iSPI installation media to reinstall an extension pack.

To reinstall an extension pack, follow these steps:

1. Log on to the NPS system as administrator or root .
2. Go to the following directory:

   On Windows:

   `%npsinstalldir%\NNMPerformanceSPI\bin`

   On Linux:

   `/opt/OV/NNMPerformanceSPI/bin`

3. Run the following command:

   On Windows:

   installExtensionPack.ovpl -e "%npsinstalldir%\NNMPerformanceSPI/extensionpacks/*<extensionpack>***.tar.gz"**

   On Linux:

   ./installExtensionPack.ovpl -e "/opt/OV/NNMPerformanceSPI/extensionpacks/*<extensionpack>***.tar.gz"**

   The command installs the extension pack again on the NPS system.

In this instance, *<extensionpack>* is the name of the extension pack.

Cause 2: Problem with the file share (only on a dedicated NPS system on Linux configured to use the NFS protocol)

This problem could arise if NPS is not able to access data files using the NFS protocol.

If you run the `nnmenableperfspi.ovpl` script on the NNMi management server for a second time and if you choose the NFS share type, the shared space on the NNMi management server is re-created. The mounted space that was configured on the NPS system originally (at the time of installation) becomes inaccessible.

To resolve this issue, follow these steps:

1. Log on to the NPS system as root.

2. Run the following commands:
   a. **umount /net/**<*NNMi_FQDN*>**/var/opt/OV/shared/perfSpi/datafiles**

   b. **/opt/OV/NNMPerformanceSPI/bin/startALL.ovpl**

In this instance, <*NNMi_FQDN*> is the FQDN of the NNMi management server.

Need more information about polling

Performance polling impacts NNMi and your network. If you enable additional performance polling, NNMi must allocate a portion of its resources to handling the additional polling, and your network will have to accommodate an increase in SNMP traffic. To ensure that the performance polling does not cause wasted resources, it is important to limit performance polling to important devices.

You must have NNMi administrator privileges to enable performance polling.

To find additonal information about polling:

1. Launch the NNM console.

2. Click **Help** > **Help for Administrators**.

3. Select **Monitoring Configuration**.

Need more information about creating new node groups

Creating a new node group requires setting device-level filters. After you create the new node group, you could use the group to filter NNMi views. You could also add the node group as a filter to the iSPI element selection tree.

Creating node groups and using the filter options are NNMi administrator tasks.

To find more information:

1. Launch the NNMi console.

2. Click **Help** > **Help for Administrators**.

3. Select **Monitoring**.

Content store can hang during upgrade when the database is in use for scheduled jobs or reports

If the content store database is in use for writing by schedules or jobs, the export can fail.

The solution is to go to the BI Server portal prior to upgrade and disable the schedules and stop some of the services the dispatcher runs.

Follow these steps:

1. Log in to the BI Portal and launch BI Administration.

2. Open the **Status** tab.

3. Click **Schedules** in the list on the left.

4. Click the check box at the top of the list.

5. Click the **Disable** icon in the top right.

6. Select the **Configuration** tab.

7. Select **Dispatchers and Services** from the list on the left.

8. Click the dispatcher in the displayed list. You will see a list of services. Be aware that, by default, there is more than one page to this list.

9. Click **More** for each of the following services, and then immediately click **Stop**.

   - BatchReportService

   - JobService

   - ReportService

Error messages appear after starting the BI Server

The following error messages appear in the command line console when you run the `startBI.ovpl` or `startALL.ovpl` command:

```
WARN: Failed command system(./cogconfig.sh -java:local -s >
"/var/opt/OV/NNMPerformanceSPI/logs/startBIServices.log" 2>&1): 512

BI Server - initializing

ERROR: Problems starting the BI Services.

ERROR: BI Service startup log: Using /opt/OV/nonOV/jdk/hpsw/bin/java

ERROR: BI Service startup log: 22/04/2015,15:30:54,Err,CFG-ERR-0106 HP NNM iSPI
Performance BI Configuration did not receive a response from the HP NNM iSPI Performance
BI service in the time allotted.

ERROR: BI Service startup log: Check that HP NNM iSPI Performance BI service is available
and properly configured.

FATAL: Stacktrace: 0. ./startBI.ovpl:31 [NPS::npslog::FATAL()]

FATAL: Failed starting the HP NNM iSPI Performance BI Service.

Died at /opt/OV/nonOV/perl/a/lib/5.8.8/NPS/npslog.pm line 194,  line 3.
```

This error may appear if a large number of Extension Packs are installed on the NPS system.

To resolve this problem, follow these steps:

1. Log on to the NPS system as root or administrator.

2. Go to the following directory:

   - Windows: `%npsinstalldir%\..\nonOV\cognos\bi\configuration`

   - Linux: `/opt/OV/../nonOV/cognos/bi/configuration`

3. Open the `cogconfig.prefs` file with a text editor.

4. Add the following lines:

ServiceMaxTries=1200

ServiceWaitInterval=2400

5. Save the file.
6. Run the following commands:
   a. **stopBI.ovpl**
   b. **startBI.ovpl**

Error messages in the command line console

The following error messages appear in the command line console when you run the `startBI.ovpl` or `startALL.ovpl` command:

WARN: Failed command system(./cogconfig.sh -java:local -s > "/var/opt/OV/NNMPerformanceSPI/logs/startBIServices.log" 2>&1): 512

BI Server - initializing

ERROR: Problems starting the BI Services.

ERROR: BI Service startup log: Using /opt/OV/nonOV/jdk/hpsw/bin/java

ERROR: BI Service startup log: 22/04/2015,15:30:54,Err,CFG-ERR-0106 HP NNM iSPI Performance BI Configuration did not receive a response from the HP NNM iSPI Performance BI service in the time allotted.

ERROR: BI Service startup log: Check that HP NNM iSPI Performance BI service is available and properly configured.

FATAL: Stacktrace: 0. ./startBI.ovpl:31 [NPS::npslog::FATAL()]

FATAL: Failed starting the HP NNM iSPI Performance BI Service.

Died at /opt/OV/nonOV/perl/a/lib/5.8.8/NPS/npslog.pm line 194, line 3.

This error may appear if a large number of Extension Packs are installed on the NPS system.

To resolve this problem, follow these steps:

1. Log on to the NPS system as root or administrator.
2. Go to the following directory:

   - *Windows:* `%npsinstalldir%\..\nonOV\cognos\bi\configuration`

   - *Linux:* `/opt/OV/../nonOV/cognos/bi/configuration`

3. Open the `cogconfig.prefs` file with a text editor.
4. Add the following lines:

   ServiceMaxTries=1200

   ServiceWaitInterval=2400

5. Save the file.
6. Run the following commands:

- **stopBI.ovpl**

- **startBI.ovpl**

NPS Console Fails to Open

The NPS console fails to open and the following error message appears in the browser:

WARNING: Using less secure method to encode cookie - invalid key length, or no key file specified

<time_and_date> com.hp.ov.cr.sso.SSOTokenCodec getNNMDataDirPath

To resolve this problem, clear the web browser cache, restart the browser, and then log on to the NPS console again.

Send Documentation Feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Online Help (Network Node Manager iSPI Performance for Metrics Software 10.20)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to network-management-doc-feedback@hpe.com.

We appreciate your feedback!