

HP Network Node Manager iSPI Performance for Metrics Software

Software Version: 9.20

Deployment by Example



Table of Contents

Introduction	3
Disclaimer	3
Concepts and Definitions	3
Data Flow Pipeline Architecture	4
Installation and Configuration	5
Server Preparation	5
Installation	5
Install License	6
Shared Directories	6
Checking Performance Polling	7
Extend Performance Polling	8
Launching NNM iSPI Performance for Metrics	12
Quick Performance Troubleshooting	13
Threshold Alerts	18
Reporting	21
Quicklaunch ReportViews Folder	22
iSPI Metrics Folder	23
Topology Filtering	31
Running with Prompts	34
Understanding Min, Max, Avg, Pctile, and Aggregations	37
Creating Report Views	41
Scheduling Reports	45
Creating a Location for the Scheduled Reports	47
Creating a Job	50
Scheduling the Job	57
Viewing Versions of the Report Views	59
Showing the Job Run History	60
Maintenance	62
Tracking the Database Size	62
Backup	64

Appendix A: E-mailing Reports on Linux	65
Appendix B: Installation with Missing Libraries	68

Introduction

This white paper shows you a sample deployment of the NNM iSPI Performance for Metrics. This is a very simplified sample, but will give you a general idea of a procedure that you can follow. This paper will not show all of the variations available. Instead, it will give a very specific example. This paper assumes an understanding of the NNMi product.

This paper is targeted at the version 9.20 of the NNMi and iSPI Performance for Metrics though many of the concepts are applicable to previous releases.

Disclaimer

This example is done from within a test lab; some of the numbers are not representative of a real network. Sometimes nodes and interfaces are shown with inappropriate performance values. This is due to intentional misconfigurations of the nodes and interfaces.

Also, there are many ways to accomplish the same procedure with NNMi and NNM iSPI Performance for Metrics. This paper will show various procedures but they are usually not the only way to do a task.

Concepts and Definitions

What is the NNM iSPI Performance for Metrics?

It is a Smart Plug-in (iSPI) for the NNMi product. NNMi is principally a fault management tool. It alerts operators when there are faults on the network. NNMi does not have a lot of capability to monitor performance issues. With the addition of the NNM iSPI Performance for Metrics, the capability of the solution expands to include performance management. This includes threshold alerts, historical analysis, and presentation of data via reports including graphs and charts.

Since it is a “plug-in,” it works in conjunction with NNMi. The products are tightly coupled together. For example, all of the performance polling is done via the NNMi state poller service; the NNM iSPI Performance for Metrics does not poll any devices. However, the storage of the performance data is not done in NNMi. Most of the presentation of the performance data is not done by NNMi either.

What is the Network Performance Server?

The Network Performance Server (NPS) is a collection of pieces of software that are common among many of the iSPIs for NNMi. It consists of the database for storing the performance data and the Business Intelligence software for presenting the performance data to the user in the form of reports. The NNM iSPI Performance for Metrics uses NPS for its data storage and presentation capabilities as do other iSPIs. NPS uses Sybase for the database storage and Business Intelligence for the report generation and presentation. This document will refer to “NPS” and the “NNM iSPI Performance for Metrics” server interchangeably. They are the same server.


NPS is installed as a component of the NNM iSPI Performance for Metrics.

iSPI licensing is all enforced in NNMi rather than in NPS. All licenses are installed within NNMi.

Data Flow Pipeline Architecture

The flow of the performance data is shown in the image below. For the NNM iSPI Performance for Metrics, all performance polling is done via the NNMi poller. This data is written to Comma-Separated Values (CSV) files. The CSV files represent all of the performance information related to a "polling cycle." In addition, topology information is occasionally written in CSV files. If NPS is on a separate server, as it will be for this paper, a shared directory structure is established between NNMi and NPS. For our example, this will be on Linux using the Network File System (NFS) protocol. If NPS is on the same server as NNMi, the same directory structure is used by all files are local to both processes.

After the CSV files are released by NNMi for NPS to consume, NPS consumes the CSV files through a set of ETL processes. ETL stands for Extract, Transform, and Load. These processes read the CSV files and write the data into the database. With data in the database, reports can be run against the data. The BI Server uses data from the database to build reports and presents the data to the user. The reports can be interactive or can be stored for future reference or e-mailed.



There are some safeguards in place at various stages through this pipeline. The first one to note is that the NNMi poller does not write more than 1G (configurable) worth of CSV files if the files are not consumed. When this limit is reached, the NNMi poller stops generating CSV files and posts an error message in the incident browser. This might be due to a stoppage of the ETL processes or a broken NFS link. NNMi stops writing so that it won't use up all the disk storage.

Installation and Configuration

Server Preparation

This example uses a two-server solution—one server for NNMi and one server for the NNM iSPI Performance for Metrics. This is the preferred method over a single-server solution. It helps distribute the load better and facilitates easier separation of tasks and maintenance.

Make sure that you have adequate hardware for the software. Consult the “System and Device Support Matrix” document. Because the NNM iSPI Performance for Metrics can retain a lot of data, you must size your environment adequately for a good experience. The support matrix provides a table with disk space requirements based on the number of polled interfaces and polled components as well as the length of retention. Make sure that storage is sized appropriately. Also an important item, though not specified in the support matrix, is disk I/O speed. NPS is very disk I/O intensive and having a fast I/O is very important.

Recommended disk I/O: RAID 1+0 or 5/6 with write cache for a local disk or high performance SAN storage.

Installation

The first step is to install the software. You should look over all of the prerequisites for the installation. Especially look at the required libraries on Linux. This is a common problem. Because it is so common, this document will show the process for working through this (see [Appendix B: Installation with Missing Libraries](#)).

Begin the installation procedure by installing NNMi first. Nothing specific needs to be done for the NNM iSPI Performance for Metrics during the NNMi installation.

Next, install the NNM iSPI Performance for Metrics on a dedicated, standalone Linux server.

At the end of the installation, the installer will run the Configuration Utility. Once you get this prompt, go to the NNMi management server and run `/opt/OV/bin/nmenableperfspi.ovpl`. This script enables the NNM iSPI Performance for Metrics functionality on the NNMi management server, checks for licensing and configures some important items like file sharing.

The `nmenableperfspi.ovpl` script asks if you want to install an evaluation license. Type Y (yes) for this. It is best to always type Y to this option even if an NNM iSPI Performance for Metrics license has already been installed.

```
# nmenableperfspi.ovpl
nmenableperfspi.ovpl[$$] -
Mon Apr 16 12:06:09 2012
```

```
Initialising...
Using OvInstallDir of </opt/OV>
Using iSPI Performance file location of
</var/opt/OV/shared/perfSpi/datafiles>
```

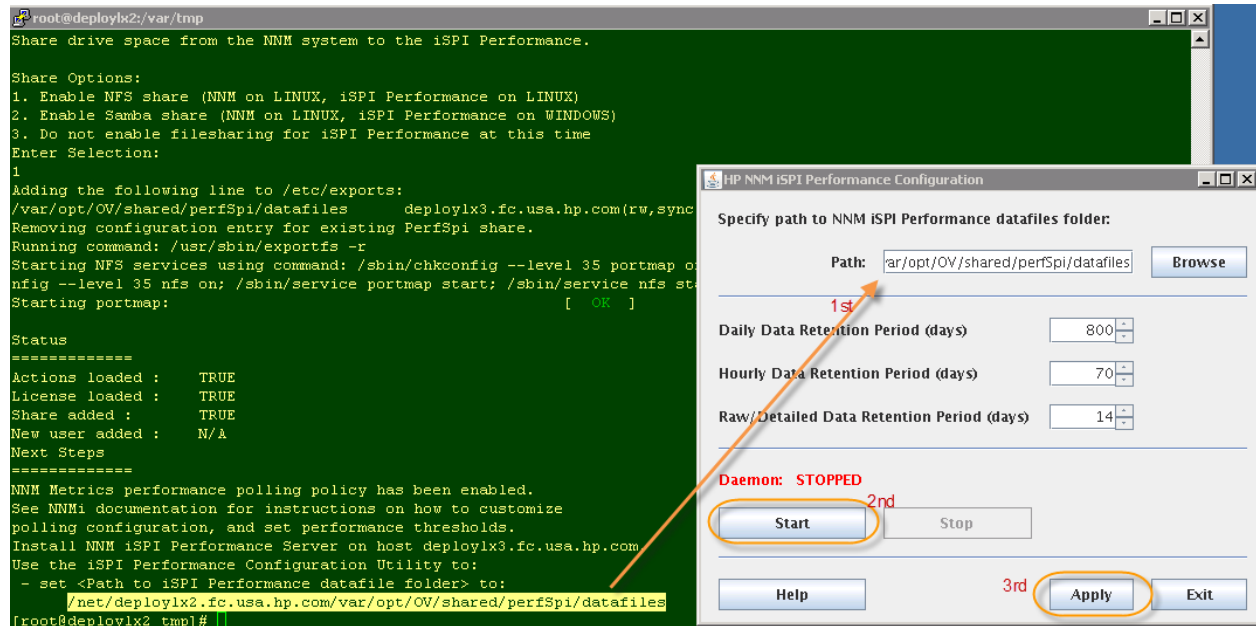
```
-----
Enable NNM iSPI Performance, Network Performance Server
-----
```

The following tasks need to be completed before performance reporting can begin.

1. Add additional URL Actions to the NNMi Actions menu
2. Share drive space between NNMi and the iSPI Performance
 - If necessary, add a user account to the NNMi server operating system
3. Enable shared single-signon security between NNMi and iSPI Performance

```
Do you want to enable the iSPI Performance for Metrics evaluation license (Y/N)? (H = Help)
Y
```

In the command line console, the script shows the path for the shared directory of the CSV files. You must then copy this value from the command line console to the Configuration Utility as shown in this image below. After you copy the path, click Start to start the daemon, and then click Apply. Finally, click Exit.



Install License

You can install a license for NNMi and for NNM iSPI Performance for Metrics with the GUI or command line.

This example assumes that you received two license files. These are not supplied by default. You can create a small test setup using the Instant-On licenses that are provided during installation.

To install licenses specified in the iAdv.5000.perm.key (for NNMi) and PerfSPI.5000.perm.key file (for the NNM iSPI Performance for Metrics), run the following commands:

- `nnmlicense.ovpl NNM -f iAdv.5000.perm.key`
- `nnmlicense.ovpl PerfSPI -f PerfSPI.5000.perm.key`

Shared Directories

If you run the “df” command, you can see that there is an NFS mount between the NNM iSPI Performance for Metrics server and the NNMi management server.

```
# df
Filesystem                1K-blocks      Used Available Use% Mounted on
deploylx2.fc.usa.hp.com:/var/opt/OV/shared/perfSpi/datafiles
210667744    8016256 191777408   5%
/net/deploylx2.fc.usa.hp.com/var/opt/OV/shared/perfSpi/datafiles
```

If you go to this directory (from either server since it is shared), you can see that it has a few subdirectories. Each subdirectory has two subdirectories—final and working.

```
# ls -F
```

```
extension/ metric/ nnm_details.xml* nps_baselinestate/ topoDump/
```

```
# ls -F metric/  
final/ working/
```

While NNMi is polling devices during a polling cycle, it accumulates the results of each poll into the `working` subdirectory. The files are compressed CSV files. These files continue to grow during a polling cycle.

```
# ls metric/working/
```

```
ComponentMetrics_20120424104506352.csv.gz InterfaceMetrics_20120424104503566.csv.gz
```

Once the polling cycle has completed, NNMi releases the files to the NNM iSPI Performance for Metrics by closing the files and moving them into the `final` directory. The NNM iSPI Performance for Metrics will regularly look for files in the `final` directory. When the iSPI sees files there, it consumes them and removes them from the `final` directory.

Check that NNM iSPI Performance for Metrics is consuming data by simply looking for any file accumulation in the “final” directories. Note that a few files actually flow in the other direction (from NNM iSPI Performance for Metrics to NNMi). Also note that the final “group” directories (node groups and interface groups) are not completely emptied by the NNM iSPI Performance for Metrics. One file of each group type stays in the directory.

Checking Performance Polling

If you want to make sure that NNMi is polling the performance data, an easy method is to do a status poll of a router from the NNMi console and look for any performance policies like the one shown below.

This indicates that performance metrics are being polled on the router.

```
Policy: Interface LAN Performance Monitoring (Etherlike SNMP v2 & v3)
Target: 16.78.56.77
Poller: NnmSnmpPoller, Target Responding: true, Poll Successful: true, Poll Duration: 49 mSec
sysUpTime 2869504454
```

Object	outUcastPkts	outNUcastPkts	AlignmentErrors	CollisionsPct	FCSLANErrorCount	FCSLANErrorCountRate	dot3Sta
Fa0/1	697191487	33815746	0	0	0	0	0
Vl133	43098894	34152195	0	100	0	0	0
Fa0/0	343227590	58796113	0	0	0	0	0
Vl141	0	47728	0	100	0	0	0
Fa2/0	90966886	1105396	0	0	3224	0	0
Fa2/8	2657817	0	0	0	0	0	17
Vl136	0	47728	0	100	0	0	0
Vl134	38827245	34485461	0	100	0	0	0
Vl36	0	47728	0	100	0	0	0
Fa2/6	479197	0	2	0	0	0	0
Vl137	393065928	39003865	0	100	0	0	0
Fa1/1	20375641	33995858	0	0	0	0	0
Vl39	0	47728	0	100	0	0	0
Vl149	0	47728	0	100	0	0	0
Fa1/0	997361306	70210351	0	0	0	0	0

Extend Performance Polling

By default, NNMi polls the performance data from all routers with out-of-the-box settings. To see this, in the NNMi console, click Configuration -> Monitoring -> Monitoring Configuration...

The screenshot shows the NNMi console with the 'Monitoring' menu open and 'Monitoring Configuration...' selected. The main window displays a table of nodes with columns: Status, Device, Name, System Location, and Device Profile. The table lists various devices like 'access-server-2', 'accsw1', 'cisco2522', etc. The 'Monitoring Configuration' dialog is open, showing 'Global Control' and 'Interface Settings' tabs. The 'Interface Settings' tab is active, showing a table of interface settings with columns: Name, Order, Enable SNMP Polling, Enable Management Polling, Enable IP Address Polling, Enable Interface Fault Polling, Enable Node Component Fault Polling, Enable Node Component Performance Polling, Poll Unconnected Interface IP Address, Enable Interface Performance Polling, Enable DSX Interface Performance Polling, Enable SONET Interface Performance Polling, Enable ATM Interface Performance Polling, Enable Frame Relay Interface Performance Polling, Enable Humint of Chassis Polling, and Enable Notes. The 'Routers' group is highlighted, and the 'Enable Node Component Performance Polling' checkbox is checked.

Status	Device	Name	System Location	Device Profile
✓		access-server-2	building 6 Annex North	cisco2621XM
✓		accsw1	building 6 Annex North	cisco2621
✓		cisco2522	5U E CPU RM	cisco2522
✓		cisco2k1	building 6 Annex North	cisco2621
✓		cisco4k1	"back of car"	cisco4500
✓		cisco4k2	5 upper east compute	cisco4500
⚠		cisco6509-loop0	building 6 North LIT LA	ciscocat6509
⚠		core6509-1	building 6 Annex North	ciscocat6509
⚠		core6509-2	building 5 parking lot	ciscocat6509
✓		dc6509-1	building 6 Annex North	ciscocat6509

Updated: 4/16/12 12:38:09 PM

Analysis

Summary

No Objects Selected

You can see the columns specific to performance polling.

The screenshot shows the 'Monitoring Configuration' dialog with the 'Interface Settings' tab selected. The dialog displays a table of interface settings with columns: Name, Order, Enable SNMP Polling, Enable Management Polling, Enable IP Address Polling, Enable Interface Fault Polling, Enable Node Component Fault Polling, Enable Node Component Performance Polling, Poll Unconnected Interface IP Address, Enable Interface Performance Polling, Enable DSX Interface Performance Polling, Enable SONET Interface Performance Polling, Enable ATM Interface Performance Polling, Enable Frame Relay Interface Performance Polling, Enable Humint of Chassis Polling, and Enable Notes. The 'Routers' group is highlighted, and the 'Enable Node Component Performance Polling' checkbox is checked.

Name	Order	Enable SNMP Polling	Enable Management Polling	Enable IP Address Polling	Enable Interface Fault Polling	Enable Node Component Fault Polling	Enable Node Component Performance Polling	Poll Unconnected Interface IP Address	Enable Interface Performance Polling	Enable DSX Interface Performance Polling	Enable SONET Interface Performance Polling	Enable ATM Interface Performance Polling	Enable Frame Relay Interface Performance Polling	Enable Humint of Chassis Polling	Enable Notes
Routers	100	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Includes nodes which do not
Networking Infrastructure I	200	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Networking Infrastructure I
Microsoft Windows System	300	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Any system running Microsoft
Non-SNMP Devices	400	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	Nodes which have never been

If you want to extend performance polling, select the checkboxes in the Monitoring Configuration form for existing groups or you can create new groups. Here is an example of a newly created interface group based on custom attributes.

Interface Groups Interface Group *

Save and Close Delete Interface Group

▼ Basics

* Name: Vital Perf Interfaces

Add to View Filter List: ☒

Node Group: []

Notes

You can filter interface groups using ifType Filters and Additional Filters. If you use both ifType Filters and Additional Filters, interfaces must match at least one ifType Filter and the Additional Filters specifications to belong to this Interface Group. If you select a Node Group, the interface must belong to a Node that is a member of that Node Group. See Help → Using the Interface Group form.

To test your Interface Group definition, select File → Save, then Actions → Interface Group Details → Show Members (Include Child Groups).

▼ NNM iSPI Performance

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

Add to Filter List: ☒

ifType Filters Additional Filters

When using the like or not like operators, use an * (asterisk) to match zero or more characters in a string and a ? (question mark) to match exactly one character in a string.

To create an inclusive IP address range, use the between operator. Valid example: ipAddress between 10.10.1.1 AND 10.10.1.255

For more information, click [here](#).

Filter Editor

Attribute	Operator	Value
ifAlias	=	

Append
Insert
Replace

Append ▼
AND
OR
NOT
EXISTS
NOT EXISTS
Delete

Filter String

(EXISTS (customAttrName = Vital) AND EXISTS (customAttrValue = 1))

Analysis - Interface Group Summary : Vital Perf Interfaces - Add to View Filter List: true Add to Filter List: false

This group could be used to govern performance polling for member interfaces. Note that you also must select the “Add to Filter List” checkbox in order for this group to appear in the NNM iSPI Performance for Metrics filtering forms.

You can enhance the polling of these interfaces to include fault and performance polling by selecting the Performance Polling checkboxes in the monitoring configuration as shown below. It might be important to set the Ordering to a lower number (higher priority) than the existing interface groups. In this example, monitoring of unconnected interfaces is chosen. This enables you to make sure that NNMi polls performance data from all interfaces that are members of this interface group.

Interface Settings *

Save and Close Delete Interface Settings

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

Basics

* Ordering 50

* Interface Group Vital Perf Interfaces

Fault Monitoring

ICMP Fault Monitoring

Enable IP Address Fault Polling ☐

SNMP Fault Monitoring

Enable Interface Fault Polling ☒

* Fault Polling Interval 5.00 Minutes

SNMP Performance Monitoring

Configuration for the optional NNM iSPI Performance for Metrics.

LAN Performance Monitoring

Enable Interface Performance Polling ☒

WAN Performance Monitoring

Enable DSx Interface Performance Polling ☐

Enable SONET Interface Performance Polling ☐

Enable ATM Interface Performance Polling ☐

Enable Frame Relay Interface Performance Polling ☐

* Performance Polling Interval 5.00 Minutes

Extend the Scope of Polling Beyond Connected Interfaces

By default, only connected interfaces are polled. These settings extend the set of monitored interfaces. It is recommended to use them with small node or Interface Groups. See Help → Using the Monitoring Configuration form.

Poll Unconnected Interfaces ☒

Poll Interfaces Hosting IP Addresses ☐

Threshold Settings Baseline Settings

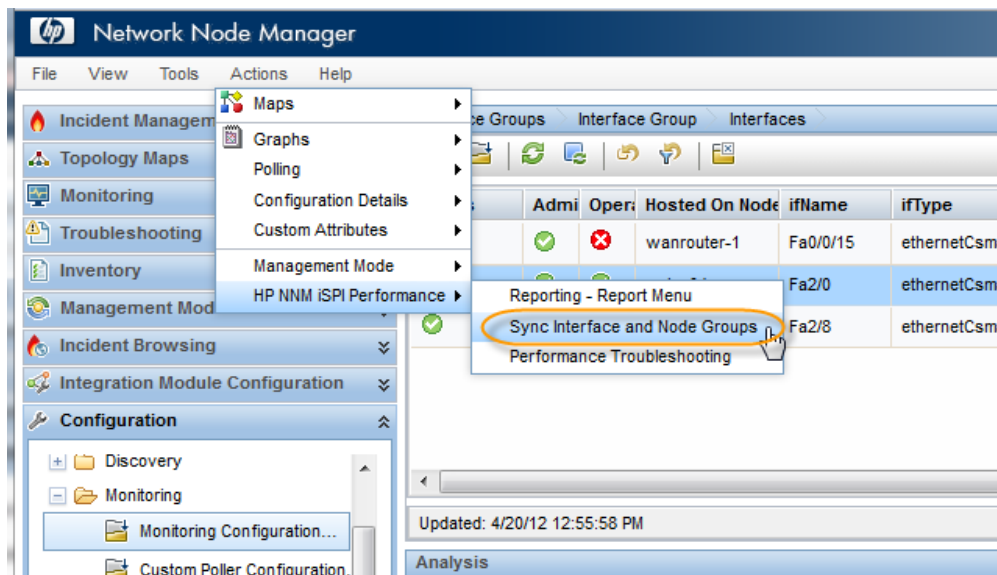
If the optional NNM iSPI Performance for Metrics is enabled, set the low and high values to determine interface performance state.

Monitored Attribute	Threshold Setting Type	High Value	High Value Rearm	Low Value	Low Value Rearm
0 - 0 of 0					

Total: 0 Selected: 0 Filter: OFF Auto refresh: OFF

Analysis

It is important to synchronize the Interface Group to the NNM iSPI Performance for Metrics. NNMi will do this automatically in time. If you want to speed this along, select Actions -> HP NNM iSPI Performances -> Sync Interface and Node Groups.

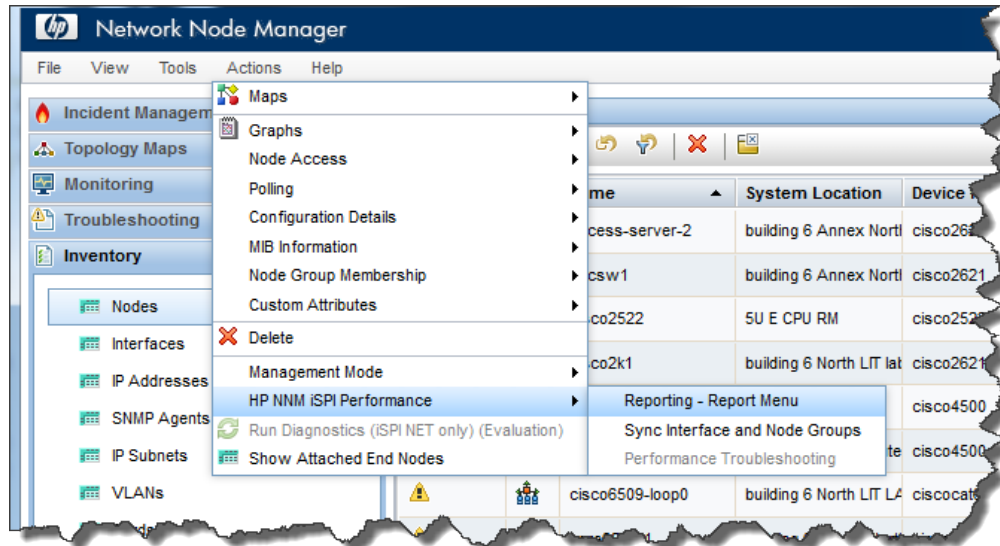


The node and interface groups should be synchronized within five minutes.

Launching NNM iSPI Performance for Metrics

To access the NNM iSPI Performance for Metrics views, select Action -> HP NNM iSPI Performance. There are three choices:

- Reporting – Report Menu: This is the main menu to the BI Server reporting.
- Synch Interfaces and Node Groups: This is a tool to speed up the synchronization when you make a change on NNMi that you want to push to the NNM iSPI Performance for Metrics quickly.
- Performance Troubleshooting: This is a quick graphing utility; it is only available if you have an Interface, Node, or L2 Connection selected.



Quick Performance Troubleshooting

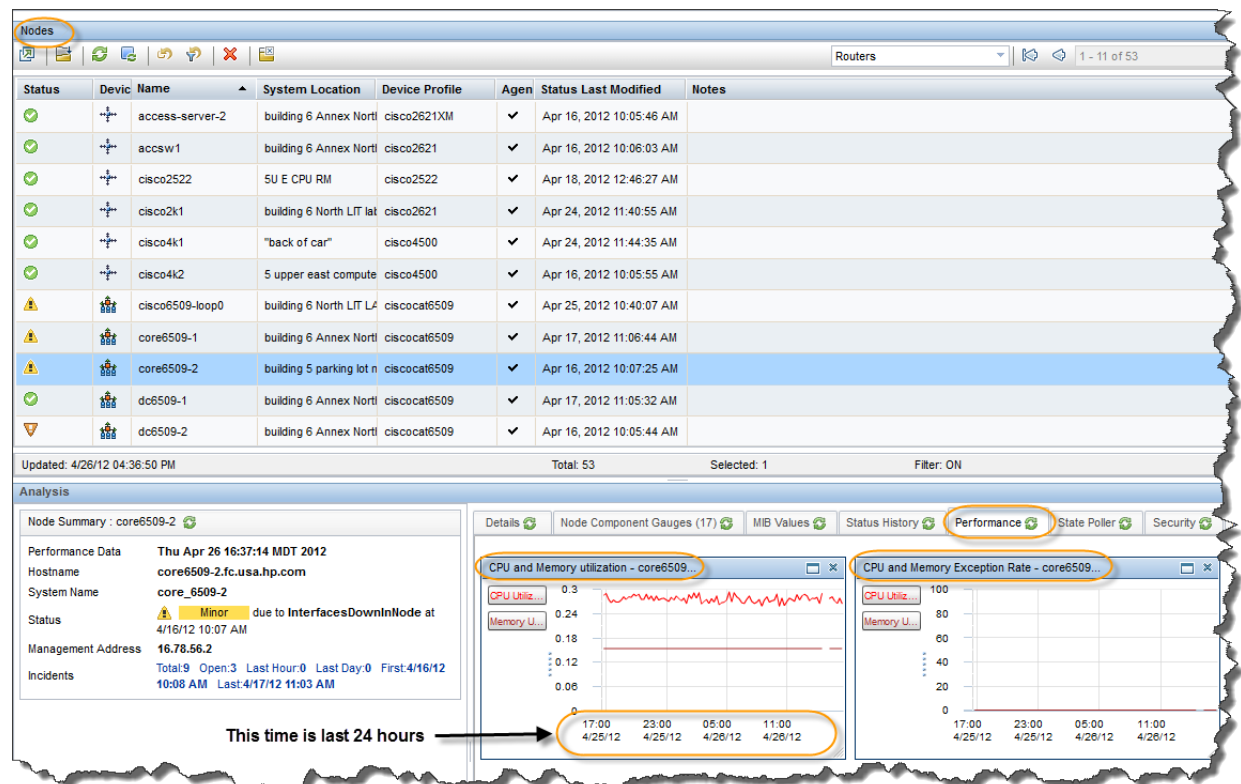
In addition to the BI reporting tool, the NNM iSPI Performance for Metrics provides you with the Performance Troubleshooting tool. It is meant to complement the BI reporting tool. This tool helps you get a quick idea of performance issues on nodes, interfaces, and connections.

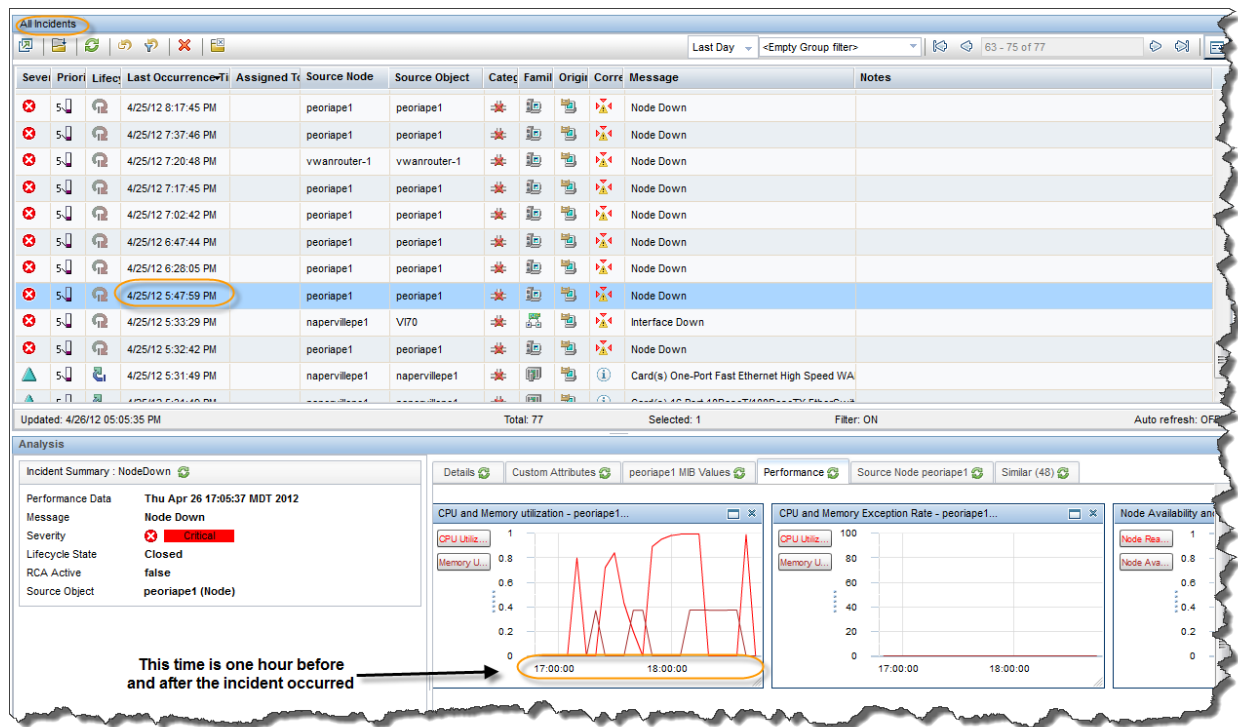
If you click on a node, an interface, or an L2 Connection in the NNMi console, and then select the Performance tab in the Analysis pane, you see a quick graph of performance data pertinent to that object. You can also click on an incident and view the graph in the Analysis pane.

Graphs are built with 24 hours' data except for incidents. For incidents, the graph represents two hours of data—one hour before and one hour after the incident occurred.


Note: Adobe Flash is required to use this feature and must be installed on the client.

Even though this graphing utility runs within the NNMi console, the data is fetched from the NNM iSPI Performance for Metrics database.

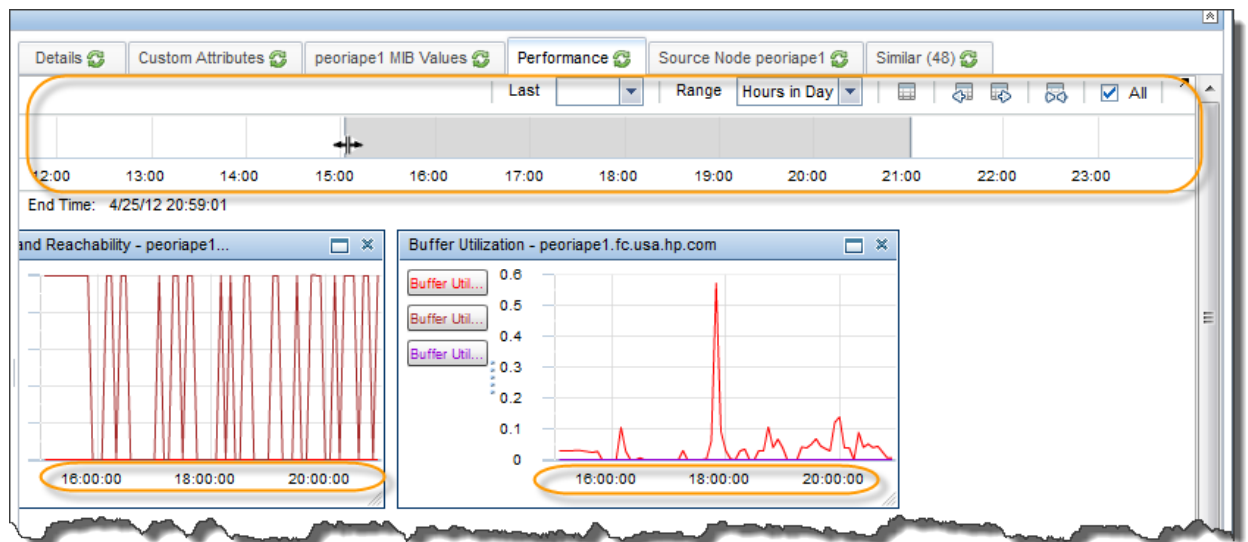
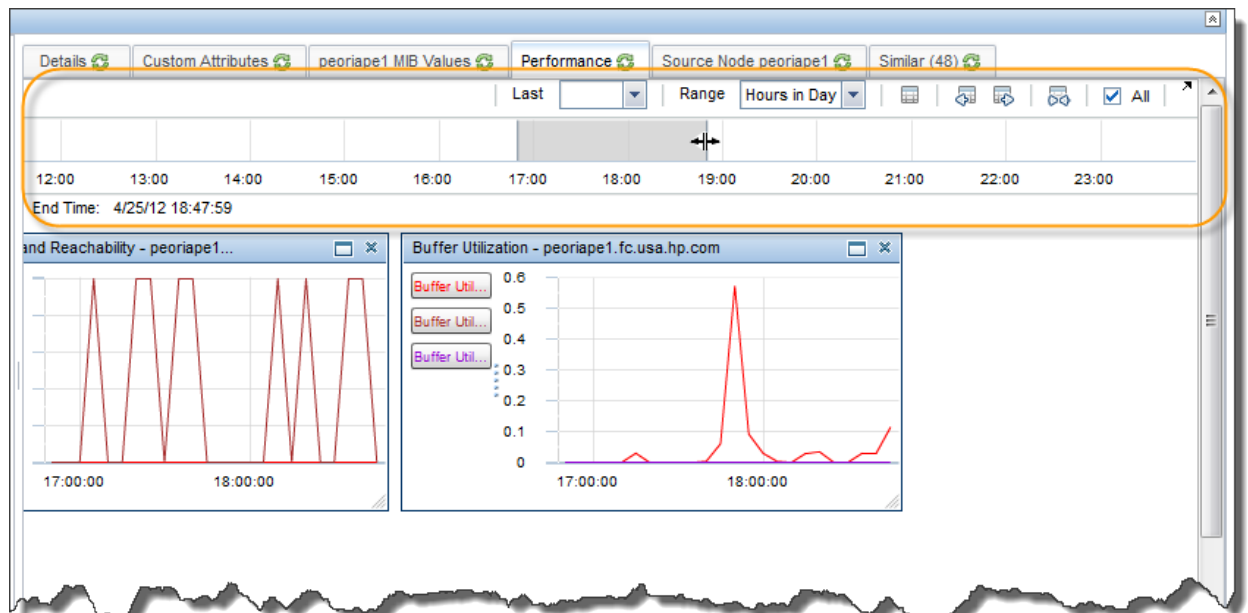




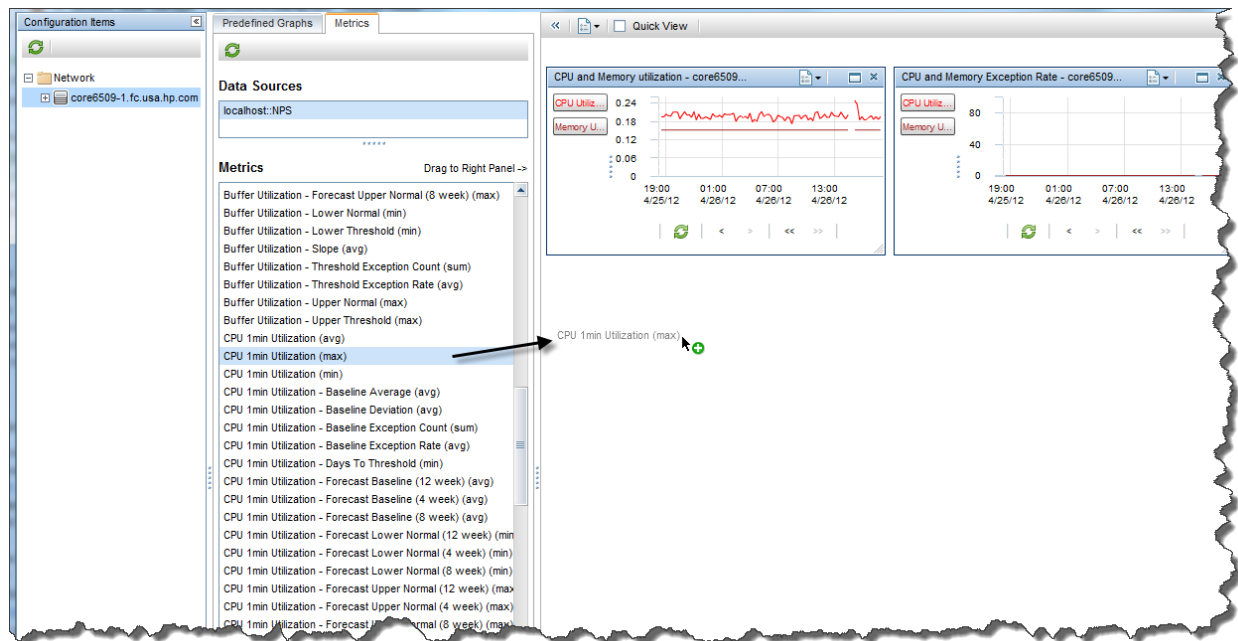
You can alter the time range from within the Performance analysis tab. If you look at the upper right corner of the space, you can see a small arrow. Click the arrow.



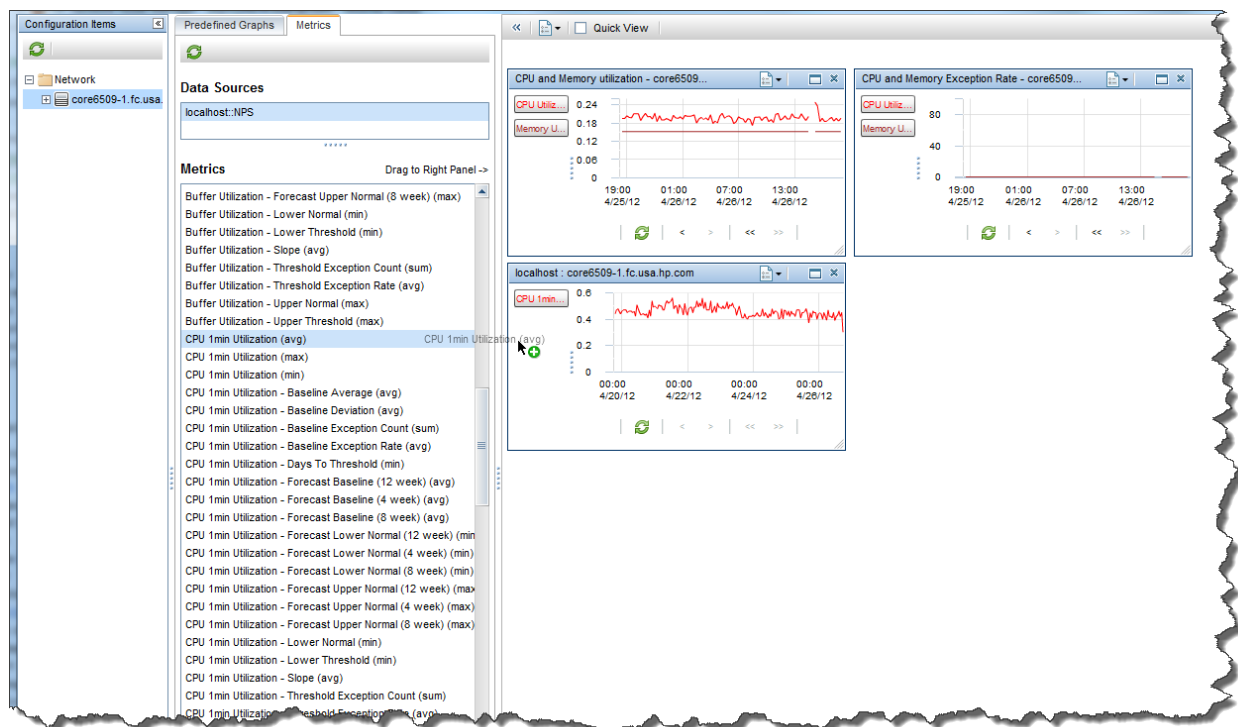
Then you will see time controls at the top of the analysis pane. You can make whatever changes you want and the graphs will reflect the change.



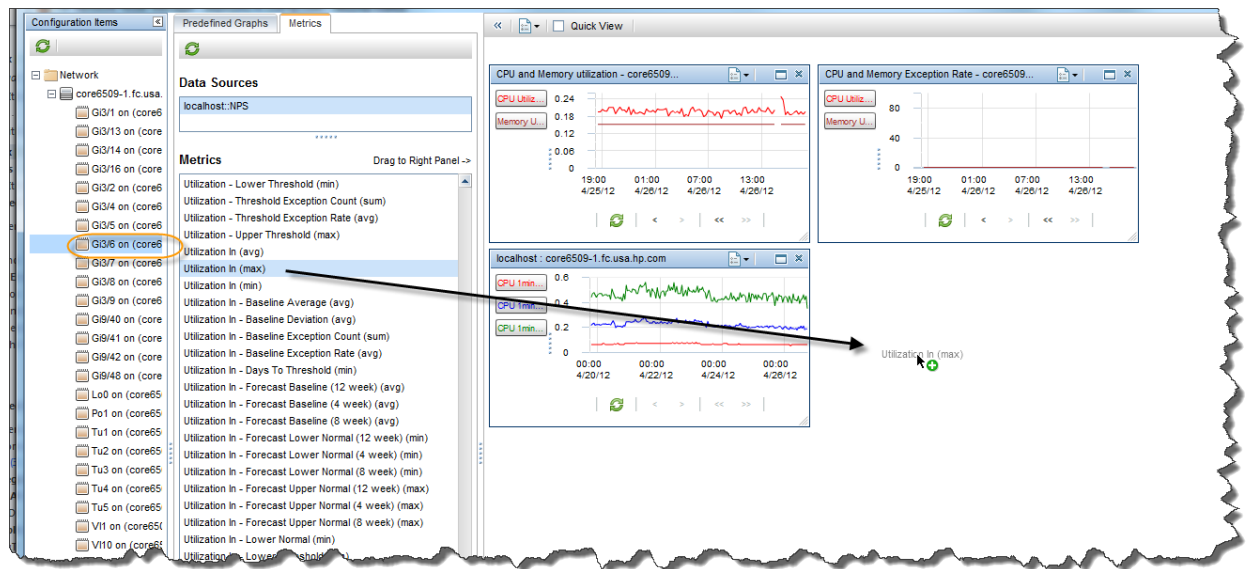
If you want a richer troubleshooting experience, you can launch the Performance Troubleshooting tool in the standalone mode. Select the node, and then click Action -> HP NNM iSPI Performance -> Performance Troubleshooting. This will launch the tool in a separate browser window. From here, you can choose the Metrics tab at the top and choose any metric you like. Simply click on the metric and drag it over to the right pane.



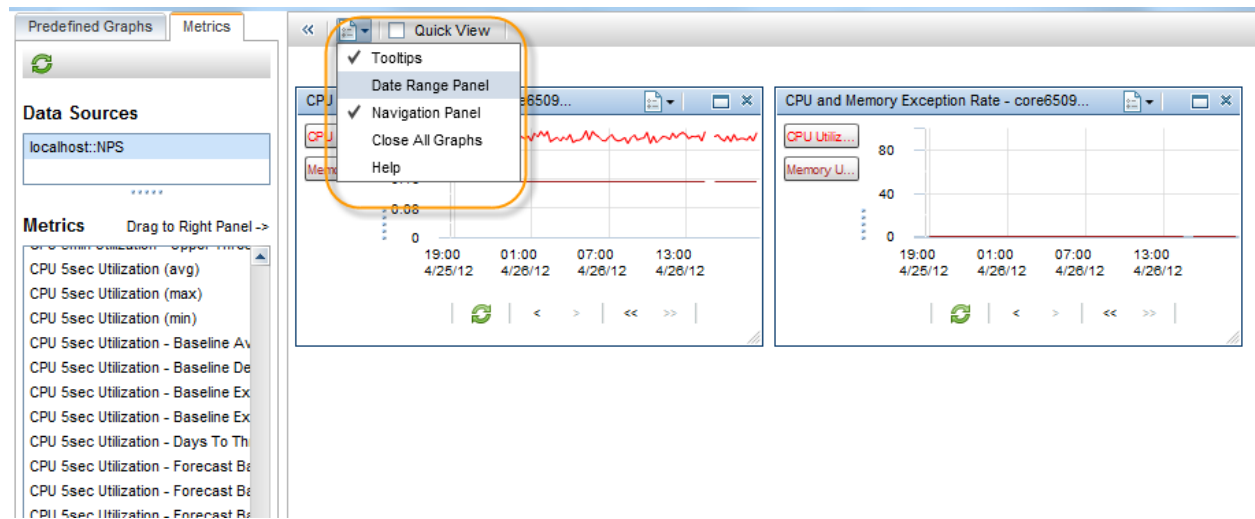
You can drag multiple metrics onto the same graph if you like.



You can even mix different metrics into the same collection of graphs in the right pane. You can expand the node to see the interfaces and select an interface of your choice. Then you can select the Utilization In (max) metric and drag that to the right pane. Now you have a collection of graphs for both node components as well as interfaces.



There are various other controls available at the top of this tool.



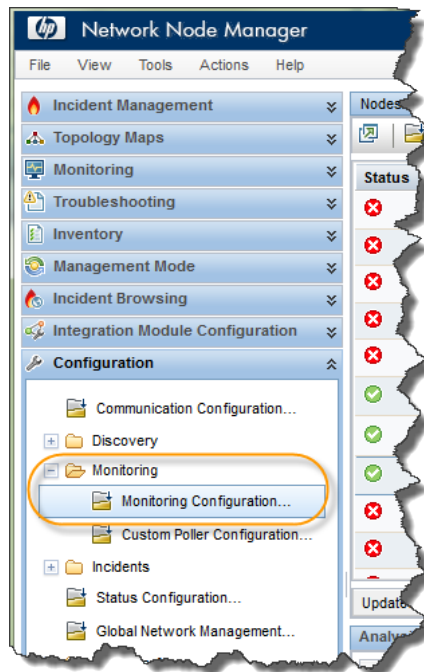
Threshold Alerts

All performance-based alerts are configured on the NNMi management server.

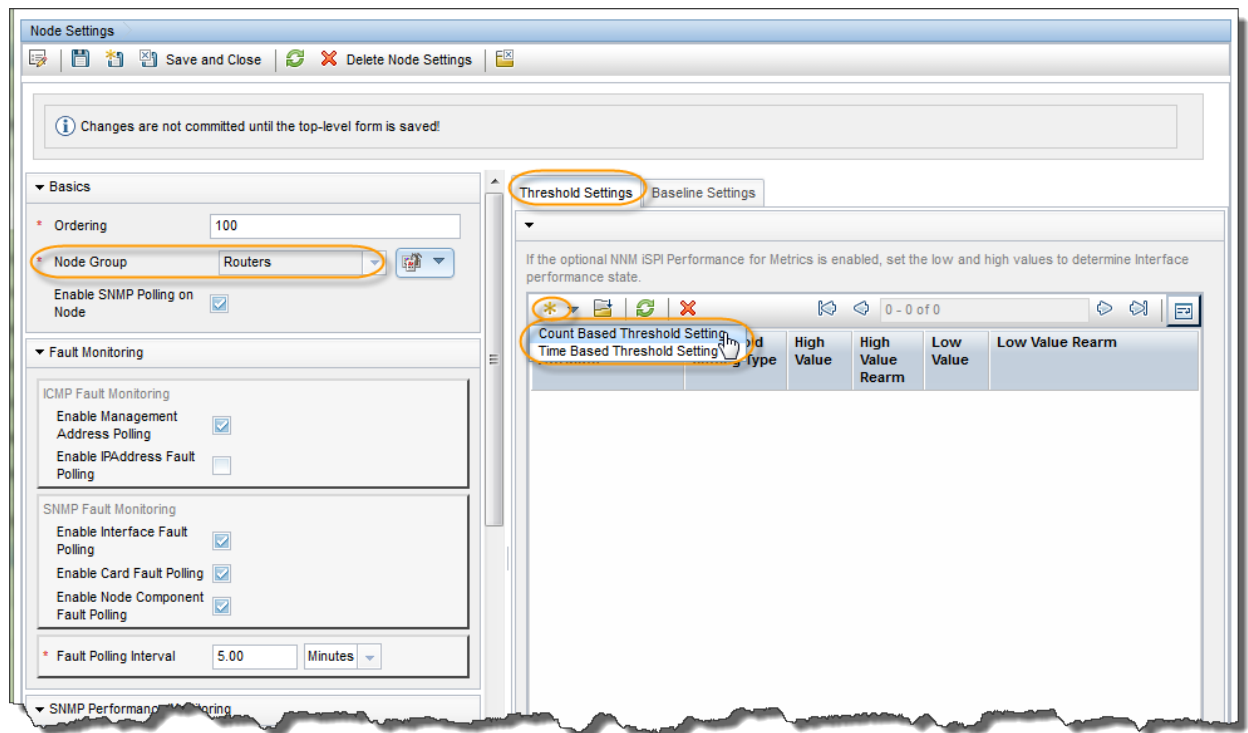
Here is an example:

Suppose you want to generate an alert when CPU on a router is at or above 90% utilization.

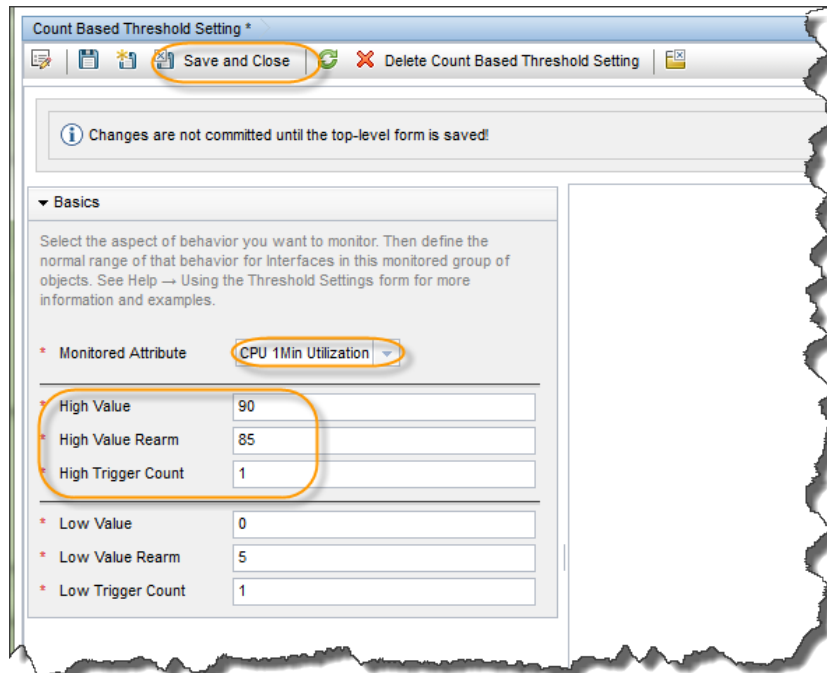
1. First go to the NNMi Configuration workspace, choose Monitoring Configuration and double-click on the Routers node group.



2. Choose the Threshold Settings tab and click on the * icon.
3. Choose Count Based Threshold Settings. (See the NNMi online help for details on different types of threshold settings.)



4. Select the Monitored Object to be CPU 1 Min Utilization. Set the High Value to 90, the rearm value to 85, and the count to 1.
5. Click Save and Close. You need not set the Low values because a low value of zero effectively disables the low threshold.



When a router has a CPU at or above 90%, you will receive an incident in the browser as shown here.

Network Node Manager

File View Tools Actions Help

Incident Management

Open Key Incidents

Unassigned Open Key Incidents

My Open Incidents

Open Key Incidents

Last Week <Empty Group

Sev	Priori	Lifec	Last Occurrence	Assigned To	Source Node	Source Object	Cate	Famil	Origin	Corre	Message	Notes
5	5	5	5/21/12 9:31:01 PM		naperville1	CPU					CPU on naperville1 utilization is too high	
5	5	5	5/20/12 5:34:32 AM		ntc6kgw1	172.20.1.1					Primary device in Router Redundancy Group sw	
5	5	5	5/19/12 12:21:03 PM		vwanrouter-1	10.100.100.1					Primary device in Router Redundancy Group sw	
5	5	5	5/19/12 1:50:48 AM		vwanrouter-1	10.100.100.1					Primary device in Router Redundancy Group sw	
5	5	5	5/18/12 8:17:04 PM		peoriapi1	VI24					Interface Down	
5	5	5	5/18/12 7:33:43 PM		peoriapi1	VI24					Correlation: No New VMI Group	


With this incident selected, you can launch to the NNM iSPI Performance for Metrics GUI and run a report for the node showing the CPU and the threshold exception count. A “threshold incident” in NNMi corresponds to a “threshold exception” in the NNM iSPI Performance for Metrics. In this particular case, the CPU was previously running higher than 90% but the breach obviously did not occur until after the threshold was configured.

NNM iSPI Performance Component Health

[Options](#) [Run Prompts](#) [Show Bookmark](#) [Help](#)

May 21, 2012 9:25:00 PM - May 21, 2012 10:25:00 PM (Last 1 Hour) (America/Denver), Node Name = naperville1.fc.usa.hp.com

Display Grain: 5 Minutes



Legend: Metric in Y1 (solid line)
CPU 1min Utilization (avg)

Metric in Y2 (dashed line)
CPU 1min Utilization - Threshold Exception Count (sum)

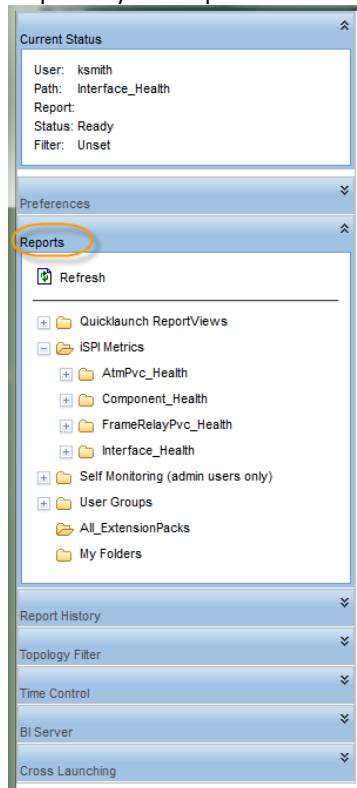
Time	CPU 1min Utilization (avg)	CPU 1min Utilization - Threshold Exception Count (sum)
May 21, 2012 9:25:00 PM	93.00%	0
May 21, 2012 9:30:00 PM	94.00%	1
May 21, 2012 9:35:00 PM	95.00%	1
May 21, 2012 9:40:00 PM	99.00%	1
May 21, 2012 9:45:00 PM	99.00%	1
May 21, 2012 9:50:00 PM	99.00%	1
May 21, 2012 9:55:00 PM	99.00%	1
May 21, 2012 10:00:00 PM	94.00%	1
May 21, 2012 10:05:00 PM	93.00%	1
May 21, 2012 10:10:00 PM	98.00%	1
May 21, 2012 10:15:00 PM	99.00%	1
May 21, 2012 10:20:00 PM	99.00%	1

Generated at : 10:32:46 PM (Server Time)

Reporting

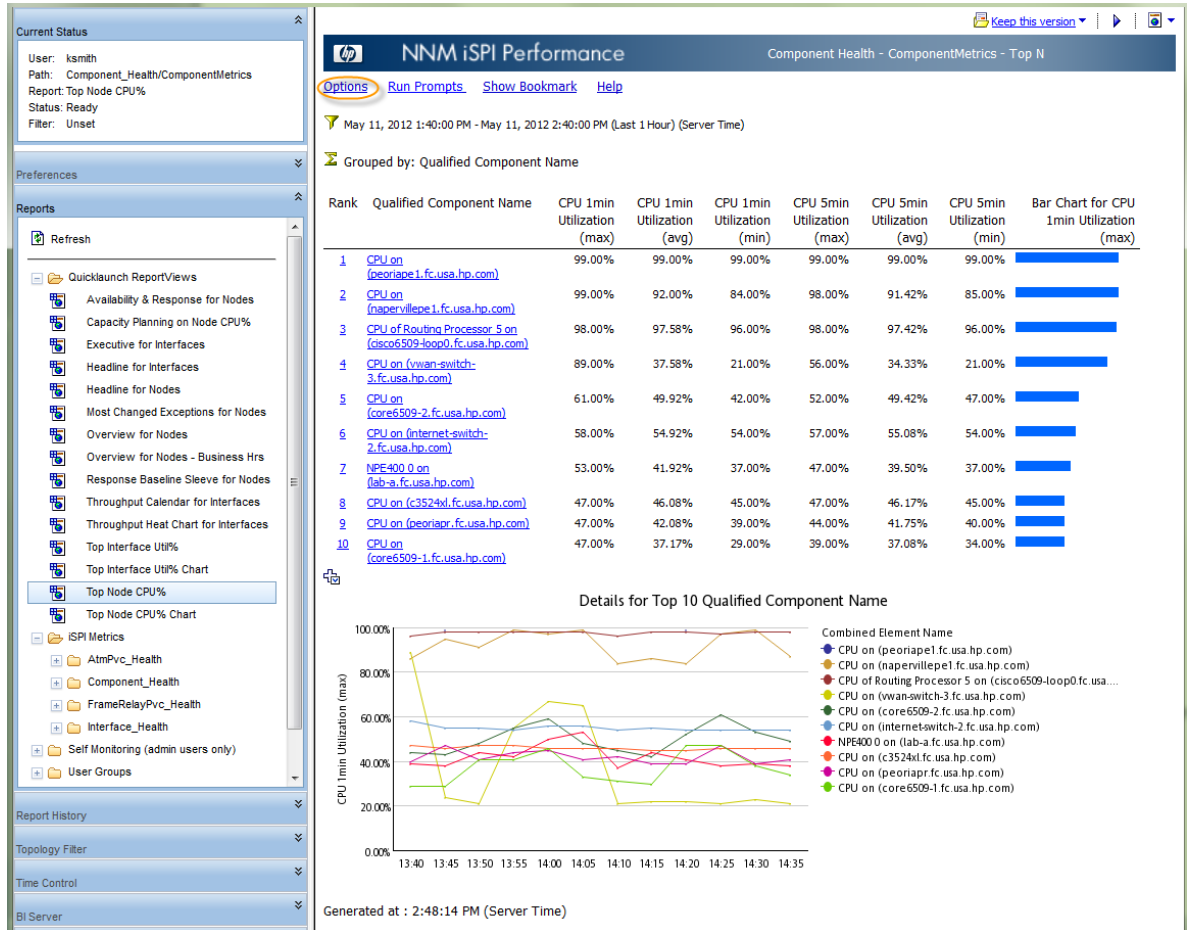
To access the NNM iSPI Performance for Metrics views, select Action -> HP NNM iSPI Performance from the NNMi console. If you have a node or interface selected, the scope of the reports will be isolated to just that object.

The primary workspace for interface and component reports is the Reports workspace.



Quicklaunch ReportViews Folder

The top folder—Quicklaunch ReportViews—is a set of reports that have pre-selected options for easy running of the reports. For example, click Top Node CPU% to view the following report:



If you want to see the pre-selected options, click Options. In this case, you can see that there are six metrics preselected and the report lists the Top 10 CPUs and they are grouped by Qualified Component Name.

The screenshot shows the 'Report Options' dialog for the NNM iSPI Performance report. The 'Top / Bottom 'N'' is set to 'Top 10'. The 'Grouping by:' is set to 'Qualified Component Name'. The 'Select Metric(s):' section shows six metrics preselected: CPU 1min Utilization (max), CPU 1min Utilization (avg), CPU 1min Utilization (min), CPU 5min Utilization (max), CPU 5min Utilization (avg), and CPU 5min Utilization (min). The 'Display Time Series Chart' is set to 'Yes'. A 'Confirm Selection' button is visible.

Top / Bottom 'N': Top 10

Grouping by: Qualified Component Name

Select Metric(s):

- CPU 1min Utilization (max)
- CPU 1min Utilization (avg)
- CPU 1min Utilization (min)
- CPU 5min Utilization (max)
- CPU 5min Utilization (avg)
- CPU 5min Utilization (min)

Display Time Series Chart: Yes

Confirm Selection

You can experiment with changing these. If you make a change to the number of CPUs on the report (for instance, changing it from Top 10 to Top 25), the change will be reflected in the immediate report, but it will not be retained the next time you bring up this report view. Report Views have the options stored as part of the view.

iSPI Metrics Folder

This folder contains a collection of reports for each “report pack.” The NNM iSPI Performance for Metrics provides four Report Packs—ATMPvc_Health, Component_Health, FrameRelayPvc_Health, and Interface_Health. There is a fifth report pack that can appear if you enable custom polling on NNMi and build corresponding report groups in NNMi.

It is important to note that these reports are not Report Views. There are default values for these reports but these defaults usually reflect the last selected value. These changes are actually stored as cookies in the browser. If you clear the cookies, the reports will revert to their original default values. The reports in the iSPI Metrics Folder are meant to be the primary workspace for interactive reporting.

Example

You want to interrogate the network to see what are the top 10 interfaces based on Utilization In. You are interested in Utilization Out as well.

To generate this report, begin by opening Interface_Health and choosing Top N. In this case, the defaults it uses are Volume and Utilization. You can see that reports are always sorted by the left-most column (Volume – Bytes (sum) in this case).

Current Status

User: ksmith
Path: Interface_Health/InterfaceMetrics
Report: Top N
Status: Ready
Filter: Unset

Preferences

Reports

- AtmPvc_Health
- Component_Health
- ComponentMetrics
- FrameRelayPvc_Health
- Interface_Health**
 - InterfaceMetrics
 - Reportlets
 - Baseline Sleeve
 - Calendar
 - Chart Detail
 - Dashboard
 - Executive
 - Headline
 - Headline - Wireless LAN
 - Heat Chart
 - Managed Inventory
 - Most Changed
 - Overview
 - Peak Period
 - Threshold Sleeve
 - Top N**
 - Top N Chart
- Self Monitoring (admin users only)
- User Groups
- All_ExtensionPacks

NNM iSPI Performance Interface Health - InterfaceMetrics

[Options](#) [Run Prompts](#) [Show Bookmark](#) [Help](#)

May 11, 2012 2:25:00 PM - May 11, 2012 3:25:00 PM (Last 1 Hour) (Server Time)

Grouped by: Qualified Interface Name

Rank	Qualified Interface Name	Volume - Bytes (sum)	Percent of ALL for Volume - Bytes (sum)	Utilization (avg)
1	1 on (nsntc-n2840-6.fc.usa.hp.com)	97,085,480,240	17.66%	20.63%
2	1 on (nsntc-n2840-5.fc.usa.hp.com)	78,131,517,360	14.21%	16.46%
3	V154 on (naperville1.fc.usa.hp.com)	45,942,995,456	8.36%	97.35%
4	Fa0/3/0 on (naperville1.fc.usa.hp.com)	45,424,011,776	8.26%	96.21%
5	1 on (nsntc-n2840-7.fc.usa.hp.com)	27,953,644,672	5.08%	5.32%
6	1 on (nsntc-n2840-4.fc.usa.hp.com)	27,089,612,800	4.93%	3.76%
7	Fa2/0/20 on (vwansw-3.fc.usa.hp.com)	20,195,945,984	3.67%	41.21%
8	Gi5/13 on (core6509-2.fc.usa.hp.com)	12,931,152,448	2.35%	2.24%
9	Gi0/1 on (vwansw-2.fc.usa.hp.com)	12,912,300,928	2.35%	2.24%
10	Fa0/24 on (vwansw-2.fc.usa.hp.com)	12,058,999,360	2.19%	23.24%
	<i>Others</i>	170,097,806,442	30.94%	

Generated at : 3:31:19 PM (Server Time)

However, you do not want the report to be based on Volume.

To change this, click Options, and then select the first metric to be Utilization In (max). Then, on the second pull-down list, select Utilization Out (max).

Report Options

Top / Bottom 'N'
Top 10

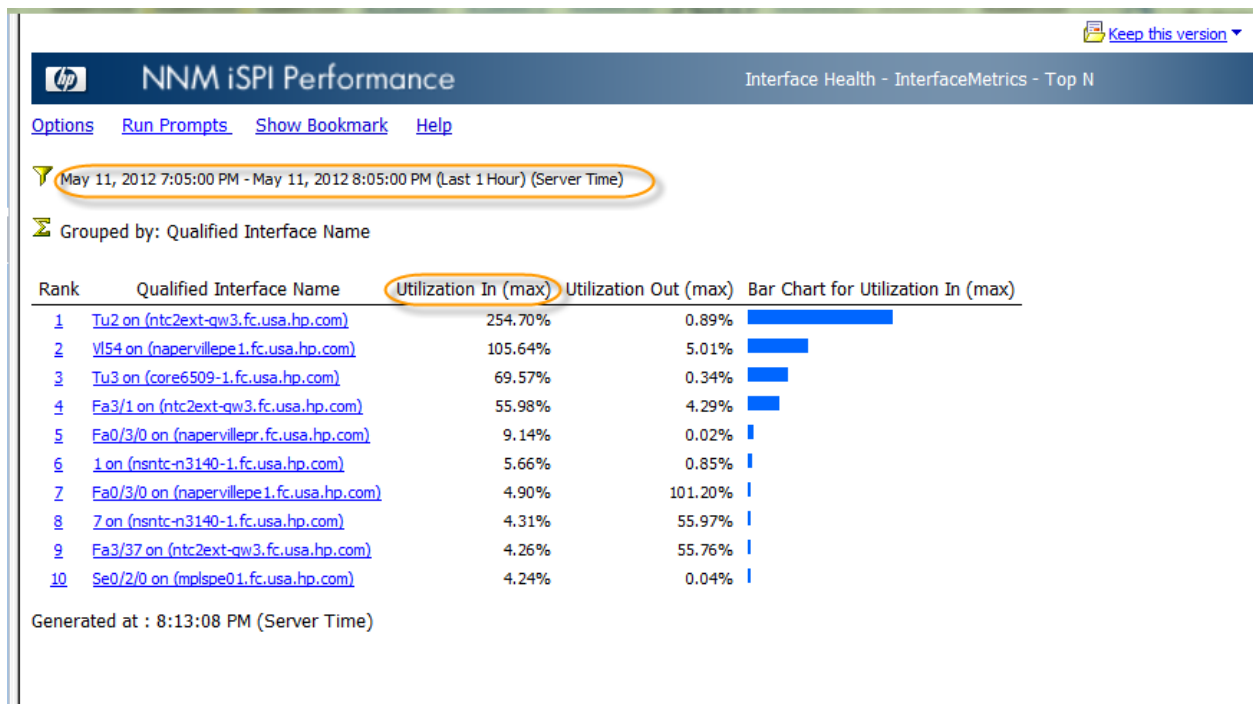
Grouping by:
Qualified Interface Name

Select Metric(s):
Utilization In (max) [selected]
Utilization Out (max) [selected]

Confirm Selection

Rank	Interface Name	Utilization In (max)	Utilization Out (max)
7	1 on (nsntc-n3140-1.fc.usa.hp.com)	5.66%	0.85%
8	Fa0/3/0 on (napervillepe1.fc.usa.hp.com)	4.90%	101.20%
9	7 on (nsntc-n3140-1.fc.usa.hp.com)	4.31%	55.97%
10	Fa3/37 on (ntc2ext-qw3.fc.usa.hp.com)	4.26%	55.76%

There are some important things to note about this report. First is the time duration for the report. It is based on the Last 1 hour. So these maximum values are for the last hour. Also note that it is sorted by the first column Utilization In. If you want it sorted by Utilization Out, you will need to swap the metric selections in the Options link.



Since these rates are maximum rates across an hour, you might be interested in knowing if it was just a quick spike or a prolonged spike. A simple way to get this data for the time period of the report (one hour in our case) is to click Options > Display Time Series Charts > Yes.

Keep this version

HP NNM iSPI Performance

Interface Health - InterfaceMetrics - Top N

[Hide Options](#) [Run Prompts](#) [Show Bookmark](#) [Help](#)

Report Options

Top / Bottom 'N'

Grouping by:

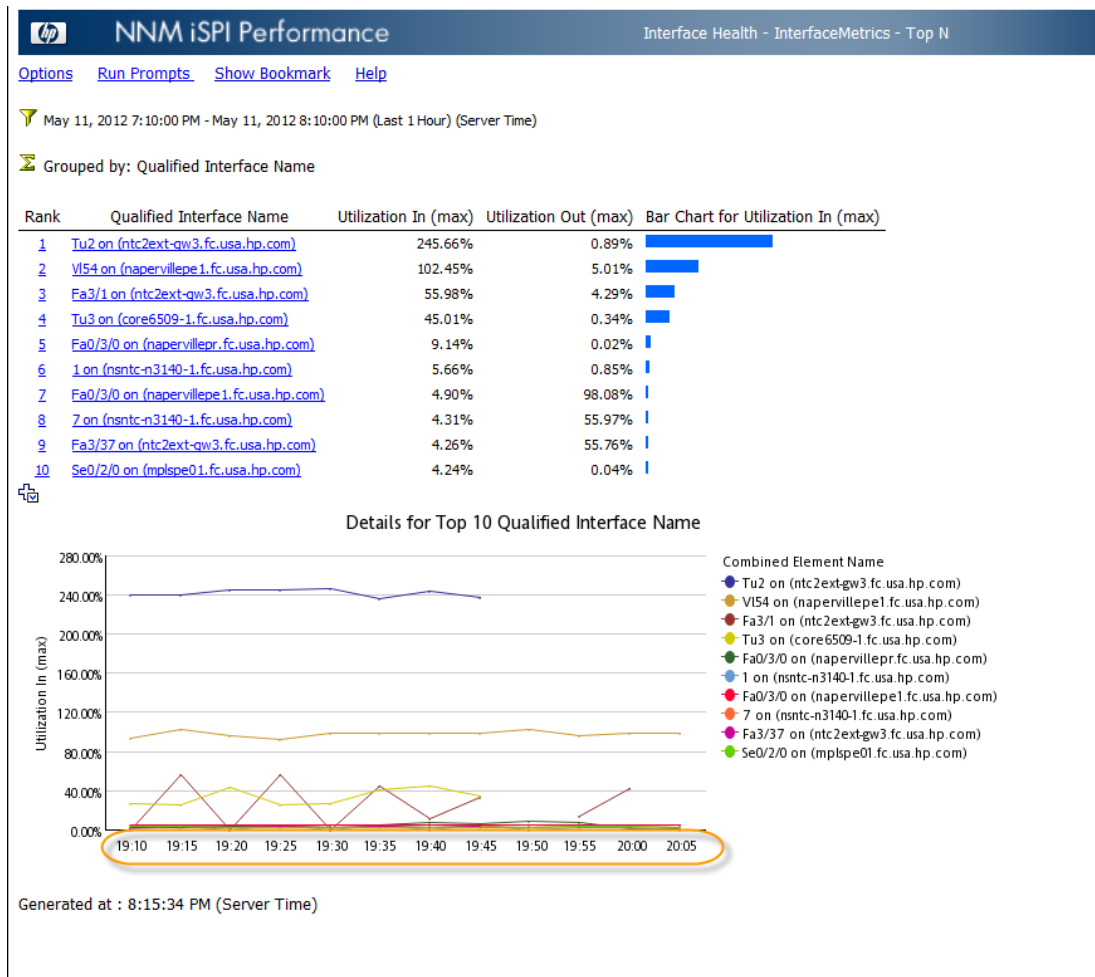
Select Metric(s):

Display Time Series Chart

May 11, 2012 7:05:00 PM - May 11, 2012 8:05:00 PM (Last 1 Hour) (Server Time)

Grouped by: Qualified Interface Name

This will produce a graph based on the first metric across one hour.








This particular graph has some anomalies due to limitations in a lab environment (and polling was stopped on some interfaces for a period of time). But ignoring that, you can see the nature of the data. In this case, it is not a sudden spike.


Now if you want to look further back than one hour, open the Time Control workspace and change the time controls. Change it to one day with the Grain set to 1 hour.

Preferences
 Reports
 Report History
 Topology Filter
Time Control

Data	Available From	Available To
Detail	05/09/12 10:23	05/11/12 20:14
Hourly	05/09/12 10:00	05/11/12 20:14
Daily	05/09/12 00:00	05/11/12 20:14

Yes No  
 Relative Start ☒ ☐  

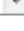
Interval: Last Rolling 


As Polled 1m 5m 15m 30m 

Grain ☐ ☐ ☒ ☐

TimeZone: Default TimeZone

Auto Refresh: Off

Hour of Day: 0:00 1:00 2:00 3:00 4:00 5:00 6:00 

Day of Week: Sunday Monday Tuesday Wednesday Thursday Friday Saturday 

Reset

Submit

The NNM iSPI Performance for Metrics does data aggregation for one hour and one day increments. Based on the report granularity, the iSPI will query the appropriate data for the most efficient reporting. Since you chose a granularity of 1 hour, the report will fetch the data points from the 1 hour aggregate table.

You can now see that the graph covers one day's worth of data.

Options

Run Prompts


Show Bookmark

Help

May 10, 2012 8:00:00 PM - May 11, 2012 8:00:00 PM (Last 1 Day) (Server Time)

Grouped by: Qualified Interface Name



Rank	Qualified Interface Name	Utilization In (max)	Utilization Out (max)	Bar Chart for Utilization In (max)
1	Po1 on (ntc2ext-gw3.fc.usa.hp.com)	516.03%	0.02%	
2	Tu2 on (ntc2ext-gw3.fc.usa.hp.com)	309.04%	0.92%	
3	VI54 on (napervillepe1.fc.usa.hp.com)	108.12%	5.38%	
4	1 on (nsntc-n2840-6.fc.usa.hp.com)	98.34%	84.93%	
5	1 on (nsntc-n2840-5.fc.usa.hp.com)	97.96%	79.65%	
6	Tu3 on (core6509-1.fc.usa.hp.com)	69.57%	0.34%	
7	Fa3/1 on (ntc2ext-gw3.fc.usa.hp.com)	67.50%	4.56%	
8	1 on (nsntc-n2840-7.fc.usa.hp.com)	47.92%	47.62%	
9	1 on (nsntc-n2840-4.fc.usa.hp.com)	40.15%	24.89%	
10	Fa0/3/0 on (napervillepr.fc.usa.hp.com)	9.40%	0.03%	



Generated at : 8:22:37 PM (Server Time)

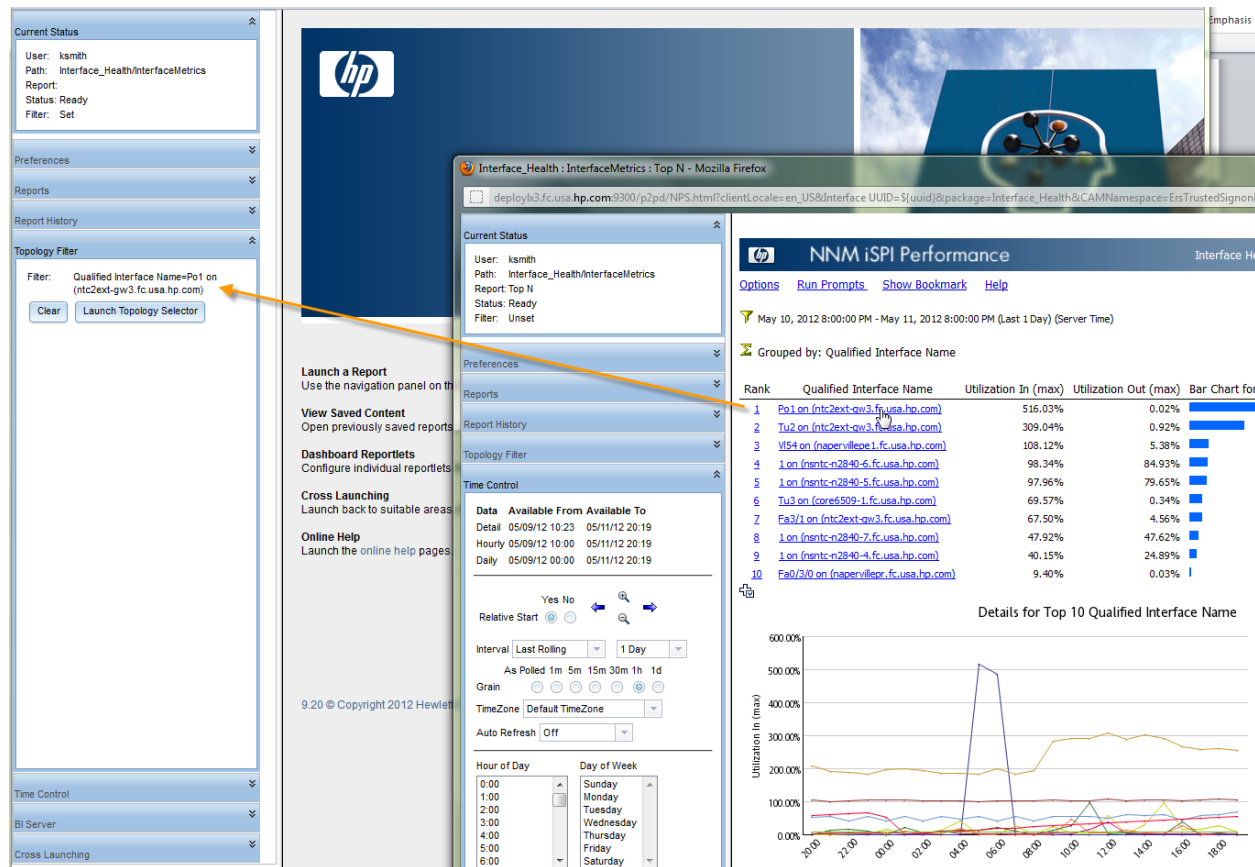
If you now want to study this top interface in more detail, click the hyperlink containing the qualified interface name to narrow down your selection.

Note that some reports do not have such hyperlinks by default. On these graphs, there is usually a plus sign with an arrow at the top of the graph that can be clicked to enable hyperlinks. Examples of these are shown below.

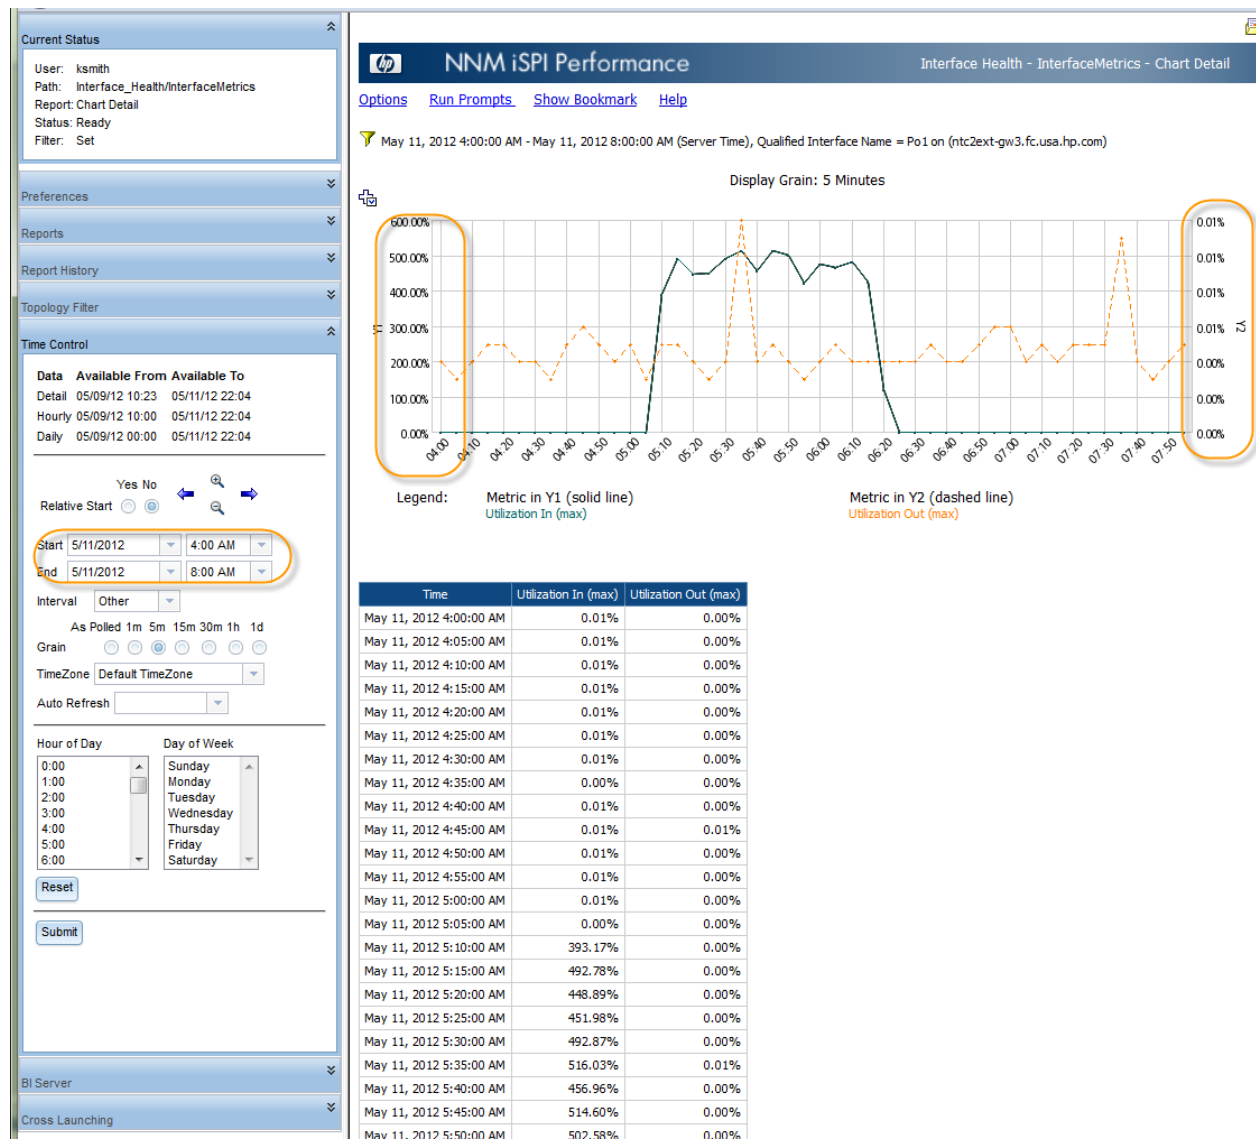


Getting back to the Top N report, click on the qualified interface name link. It will then launch another instance of the BI User Interface into a separate browser window. The original BI User Interface window is still present, but you now have a new window that has the specific interface selected in the topology filter. This means that reports that are run in this new window will be specific to that interface. The original User Interface is left for convenience in case you want to work in the

original context with no interface selected. You can see in this image that clicking on the link launches a separate window with the topology filter selected.



With this specific interface selected, you can run a Chart Detail report and zoom in on the time. There are two different scales on the Y axis on the graph—one on the left and one on the right.



Topology Filtering

Consider narrowing down some of the reports to include a specific node group or interface group. You have a group of nodes that you want to monitor. This group is called Naperville Nodes. It is important to make sure that you have added the Node Group to the NNM iSPI Performance filter list in NNMi as shown below.

The screenshot shows the 'Node Groups' configuration window in NNMi. The 'Node Group' tab is active, showing the configuration for 'Naperville Nodes'. The 'Basics' section includes fields for Name, Calculate Status, Status, and Add to View Filter List. The 'NNM iSPI Performance' section has an 'Add to Filter List' button highlighted with an orange circle. The 'Filter Editor' section on the right shows a filter rule: 'hostname like naperv*'. The status bar at the bottom indicates 'Analysis - Node Group Summary : Naperville Nodes - Status: No Status Status Last Modified: May 12, 2012 9:19:33 AM MDT'.

Node Groups Node Group

Save and Close Delete Node Group

Basics

Name: Naperville Nodes

Calculate Status: ☒

Status: No Status

Add to View Filter List: ☒

Notes

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

To test your Node Group definition, select File → Save, then Actions → Node Group Details → Preview Members (Current Group Only).

NNM iSPI Performance

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

Add to Filter List: ☒

Filter Editor

Attribute Operator Value

tenantUuid !=

hostname like naperv*

Append Insert Replace

Append AND OR NOT EXISTS NOT EXISTS Delete

Analysis - Node Group Summary : Naperville Nodes - Status: No Status Status Last Modified: May 12, 2012 9:19:33 AM MDT

To launch a report view with the context of this node group, you can follow several different ways. This document will show how to select the node group from the NNM iSPI Performance for Metrics user interface.

After you run a report, click the Topology Filter workspace to narrow down the scope to a specific node group and click Launch Topology Selector. There are many selections that can be made.

Click on the Single value select icon for NodeGroup Name. Note that two types of Topology Groups (SCD Type 1 and SCE Type 2) exist. Select Type 1 as that reflects the present definition of the node group. For further details about these types, see the online help.

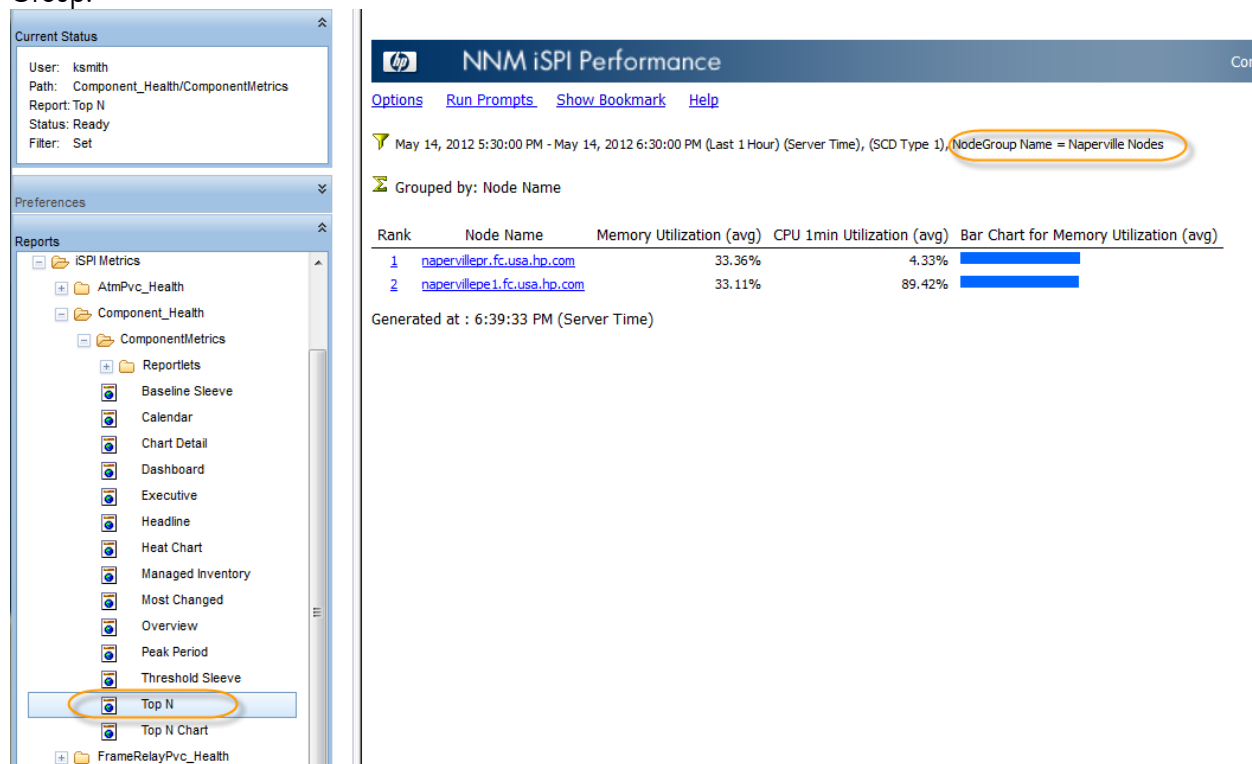
Select the node group you are interested in. Not all node groups are shown in the pull down menu. Instead, only the ones that have been added to the Filter List for NNM iSPI Performance are shown. If an expected node group does not appear in the pull down list, do the following:

- Check that the Add to Filter List box has been checked.
- Run the Sync Interface and Node Groups menu item and wait for 5 minutes.

Also, a node group needs to have nodes populated in it and the nodes need to be polled for performance data in order for the node group to appear in this list.

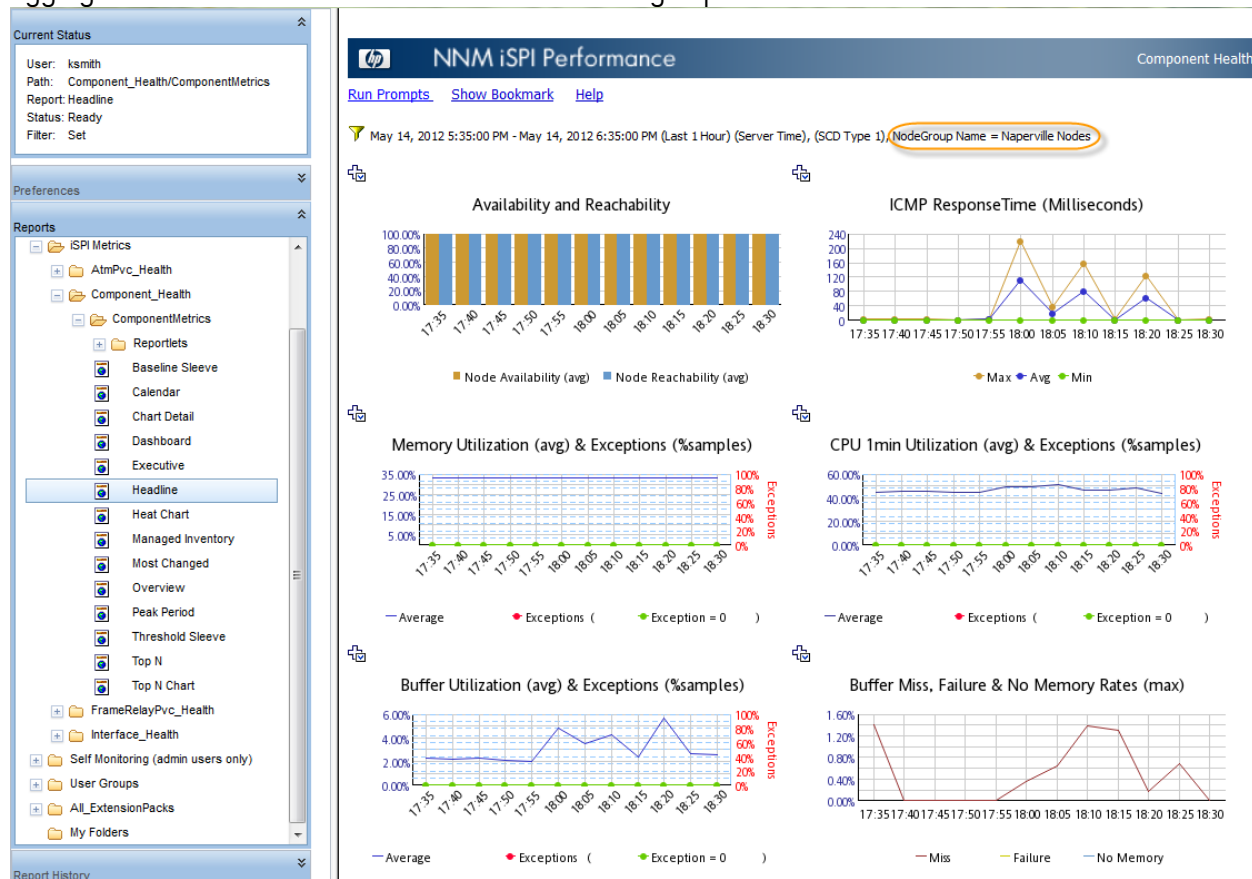
After you select the node group, click Apply, and then click Finish. The report will be re-built.

You can see the filtering criteria at the top indicating the Node Group.



This filter now applies to all subsequent reports until the selection is cleared.

If you run the Component Headline report, you will see that the numbers on the report are an aggregation of the nodes that are members of the filter group.



Running with Prompts

A handy feature available in 9.20 is the ability to run with prompts. This feature enables you to quickly select a set of attributes for a report.

There are two ways to use this feature:

- Right-click on the report and select Run with prompts. The advantage of making this selection with a right click is that you can set the attributes for the report prior to running the report.
- After you run the report, click Run Prompts at the top of reports.

The screenshot displays the NNM iSPI Performance web interface. On the left, a sidebar shows a tree of reports under 'ComponentMetrics', with 'Top N' selected. A right-click context menu is open over 'Top N', showing options: 'Properties', 'Run', 'Run with prompts' (highlighted with a red circle), and 'Create ReportView'. The main panel shows the 'Run Prompts' button in the top navigation bar, also highlighted with a red circle. Below it, a table displays performance metrics for various nodes.

Rank	Node Name	Memory Utilization (avg)	CPU 1min Utilization (avg)	Bar Chart for Memory Utilization (avg)
1	internet-switch-3.fc.usa.hp.com	97.57%	16.33%	
2	mpls04.fc.usa.hp.com	90.37%	4.58%	
3	mpls2950-1.fc.usa.hp.com	74.28%	15.58%	
4	internet-switch-4.fc.usa.hp.com	73.80%	15.25%	
5	c2900xl-1.fc.usa.hp.com	69.49%	32.42%	
6	c3524xl.fc.usa.hp.com	64.88%	45.17%	
7	internet-switch-2.fc.usa.hp.com	64.48%	54.17%	
8	vwan-switch-3.fc.usa.hp.com	60.76%	25.50%	
9	wanrouter-1.fc.usa.hp.com	58.39%	7.90%	
10	wan-bo1-sw.fc.usa.hp.com	58.38%	10.83%	

Generated at : 6:47:12 PM (Server Time)

The first prompt page is the topology selector. You can make changes here. If you choose to change other attributes, click Next. Otherwise, click Finish.

Topology Filter

Attribute Selection

Topology group tracking method SCD Type 1

NodeGroup Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	=	topo_NodeGroupName
Qualified Component Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Component Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Component Type	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Component ID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Component UUID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Short Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Contact	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Location	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Family	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Vendor	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node ID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node UUID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node ODBID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Tenant Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Tenant UUID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
SecGroup Name	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
SecGroup UUID	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		
Node Annotation	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>		

Apply Reset

Cancel < Back **Next >** Finish

The next prompt page enables you to change the Time Control for this report. In this example, it is changed to the last 12 hours with the granularity at 15 minutes.

Time Controls

Data From:	Relative Start	Yes	Hour of Day	Day of Week
Detailed: 05/09/12 10:21			0	Sunday
Hourly: 05/09/12 10:00		Last Rolling Interval	1	Monday
Daily: 05/09/12 00:00			2	Tuesday
			3	Wednesday
			4	Thursday
			5	Friday
			6	Saturday

Data To: 05/14/12 18:49
 Interval: 12 Hours
 Display Grain: Qtr Hour
 Auto Refresh: No
 TimeZone: TimeZone

[Select all](#) [Deselect all](#)
[Select all](#) [Deselect all](#)

Cancel < Back **Next >** Finish

The final prompt page is the report options. In this example, choose to view graph at the bottom of the report. Note that, depending on the particular report selected, sometimes this third prompt page is not given.

Report Options

Top / Bottom 'N'

Top 10

Grouping by:

Node Name

Select Metric(s):

Memory Utilization (avg)

CPU 1min Utilization (avg)

Display Time Series Chart

Yes

Show Chart

Yes

No

Confirm Selection

Cancel

< Back

Next >

Finish

Understanding Min, Max, Avg, Pctile, and Aggregations

The majority of the metrics available for reporting include a Minimum, Maximum, Average, and four percentiles (5, 90, 95, and 99). It is important to know how these metrics can change based on the time frame of a report and the metric selected for the report.

Data points on a graph can be individual samples from a router or they may be aggregates of data. Aggregation can be done in a few different dimensions. For the “Buffer Utilization” metric, this may not be an individual data point if the node supports multiple buffers. Many devices have multiple buffers (big, huge, large, medium, and small). A component on a graph may represent an aggregation of multiple components.

Another form of aggregation can be due to the granularity of the graph. If the data is sampled at 5 minutes, but the granularity of the graph is set to 15 minutes, then each data point on the graph will be an aggregation of three data points.

Yet another form of aggregation is at one hour and one day intervals. These aggregations are actually stored in the database rather than being computed dynamically to improve performance and reduce data retention size.


The Minimum and Maximum aggregations are straightforward. They represent the minimum and maximum samples across the granularity of the graph. Average represents the average of the samples across the granularity. If the graph uses a 15-minute granularity, each point on the graph is an average of three points. Therefore, when you create a graph of an individual component like Medium Buffer Utilization with a granularity equal to the polling cycle, the minimum, maximum, and average graphs look identical. This is because each point represents an individual sample from the router. But if the graph has wider granularity or is an aggregate of components, then the points on the graph will likely not track the individual samples directly.

Example

Here is a graph of Buffer Utilization of an individual node. The granularity is at 5 minutes (same as the poll rate), but the min, max, and average are not identical. This means that there is aggregation occurring at the component level. This node actually has five buffer sizes and they are aggregated to show the points on the graph.

[Options](#) [Run Prompts](#) [Show Bookmark](#) [Help](#)May 22, 2012 7:40:00 PM - May 22, 2012 8:35:00 PM (Server Time), Node Name **mplspe04.fc.usa.hp.com**


Display Grain: 5 Minutes





Time	Buffer Utilization (max)	Buffer Utilization (min)	Buffer Utilization (avg)
May 22, 2012 7:40:00 PM	93.24%	0.00%	24.20%
May 22, 2012 7:45:00 PM	93.24%	0.00%	24.40%
May 22, 2012 7:50:00 PM	93.24%	0.00%	24.00%
May 22, 2012 7:55:00 PM	93.24%	0.00%	23.45%
May 22, 2012 8:00:00 PM	92.00%	0.00%	23.03%
May 22, 2012 8:05:00 PM	85.19%	0.00%	21.33%
May 22, 2012 8:10:00 PM	85.19%	0.00%	21.49%
May 22, 2012 8:15:00 PM	85.19%	0.00%	21.33%
May 22, 2012 8:20:00 PM	85.19%	0.00%	21.65%
May 22, 2012 8:25:00 PM	85.19%	0.00%	21.65%
May 22, 2012 8:30:00 PM	85.19%	0.00%	21.65%

If you run a slightly different graph, shown below, you can break down the graph for individual components. With a five minute granularity, now min, max and average are all identical since they are a single sample from the router.

[Options](#)
[Run Prompts](#)
[Show Bookmark](#)
[Help](#)

 May 22, 2012 7:40:00 PM - May 22, 2012 8:35:00 PM (Server Time), Node Name = mplspe04.fc.usa.hp.com

 Grouped by: Qualified Component Name



Finally, if you take an individual component but change the time granularity to be 15 minutes, then the min, max, and average are no longer always identical because the data points are actually aggregations of three sampled data points. That can be seen in the graph below.

Current Status

User: ksmith
 Path: Component_Health/ComponentMetrics
 Report: Top N Chart
 Status: Ready
 Filter: Set

Preferences

Reports

Report History

Topology Filter


Time Control

Relative Start: Yes No

Start: 5/22/2012 7:30 PM
 End: 5/22/2012 8:35 PM
 Interval: Other
 As Polled: 1m 5m 15m 30m 1h 1d
 Grain: ☒ 15m
 TimeZone: Default TimeZone
 Auto Refresh: ☐

Hour of Day: 0:00 1:00 2:00 3:00 4:00 5:00 6:00
 Day of Week: Sunday Monday Tuesday Wednesday Thursday Friday Saturday

Reset Submit



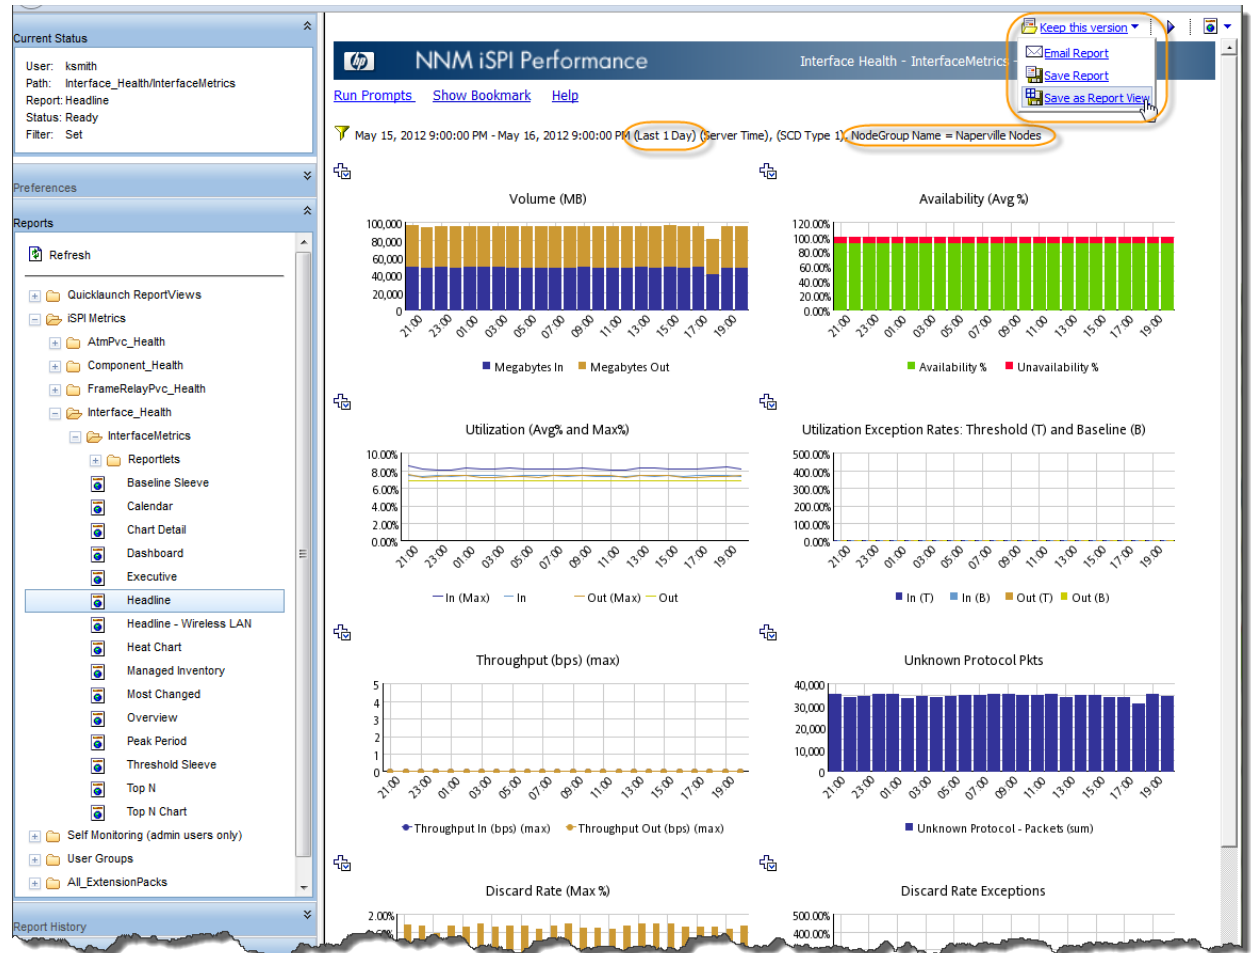
It is important to remember that graphs can be aggregated across multiple dimensions and the data needs to be interpreted with this in mind.

In addition, there is also "percentile." The NNM iSPI Performance for Metrics tracks 90, 95, 99, and 05 percentiles. These metrics are used to help even out spikes in the data. Sometimes, the maximum value is not the best value to be concerned with, especially if it involves a temporary spike in the data. For planning purposes, usually the percentile values are better metrics to work with since they tend to eliminate short anomalies in the data.

These metrics cannot be run against hourly or daily aggregates because all samples are needed for accurate computation. So reports can only be run as far back as the raw storage is available. So in our sample deployment, this goes back 14 days. Here, you cannot report percentiles further back than 14 days.

Creating Report Views

A Report View is like a copy of a report template with all of the properties for the report saved. This is a great way to save strategic reports for easy access. The easiest way to create a report view is to first build the report. In this example, you will build an Interface Headline report for the Naperville Nodes over the last one day. Then go to the top, click Keep this version, and then choose "Save as Report View."



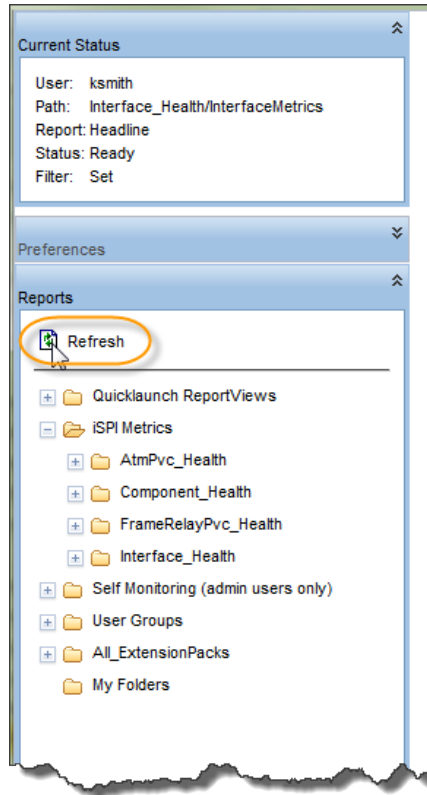
Always create a Report View with a descriptive name. You must also decide how much visibility you want the Report View to have. If you leave it in the Public Folders, all users will be able to view this Report View. If you save the Report View in My Folders, it will only be visible only to you.

The 'Save as report view' dialog box is shown. It prompts the user to specify a name and location for the report view. The 'Name' field contains 'Naperville IF Headline Last Day ReportView'. The 'Location' field shows 'Public Folders > iSPI Metrics'. Below the location field, there are two buttons: 'Select another location...' and 'Select My Folders'. The 'OK' and 'Cancel' buttons are at the bottom.

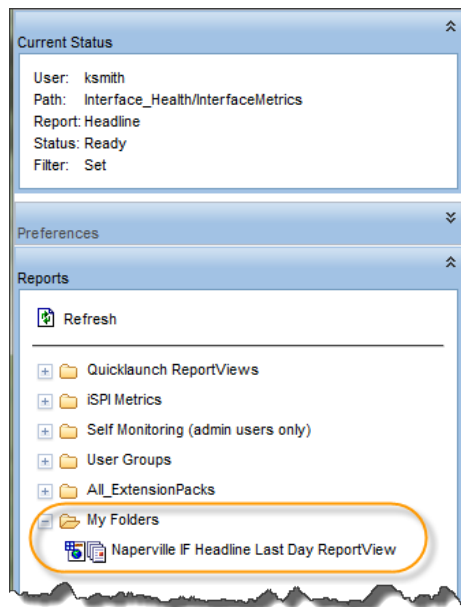
Report Views can also be saved into folders that allow all Level 1 or Level 2 operators to view them. Also, the multi-tenancy feature is in effect for Report Views and they can be made visible only to certain tenants.

After changing the location, click OK.

Next, go to the Reports workspace and click Refresh. Sometimes it is necessary to refresh a few times.

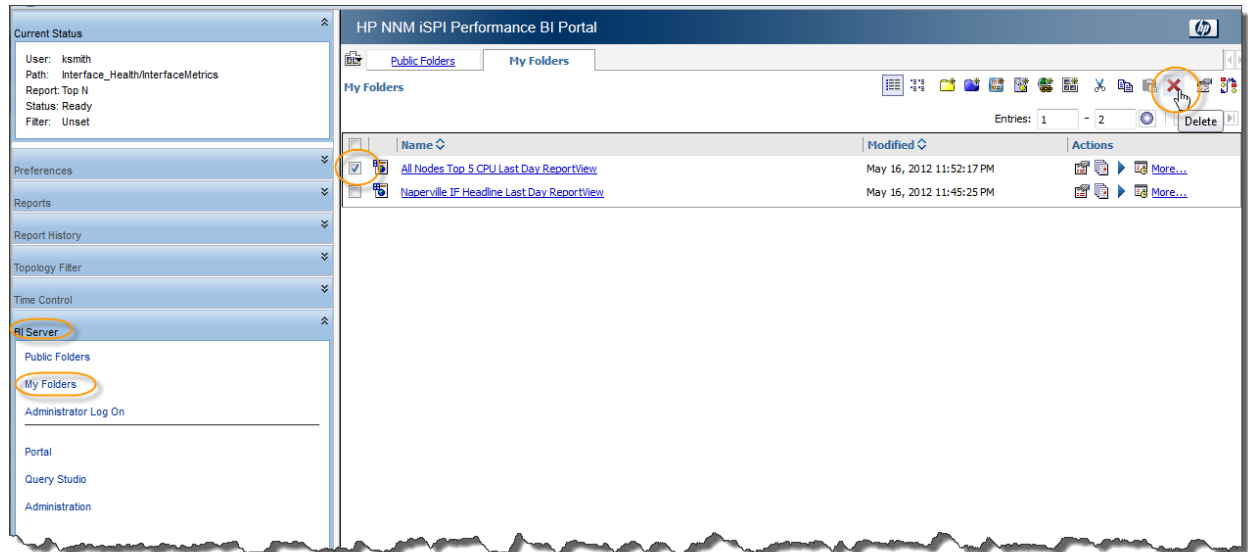


Eventually, the report view will appear under My Folders.

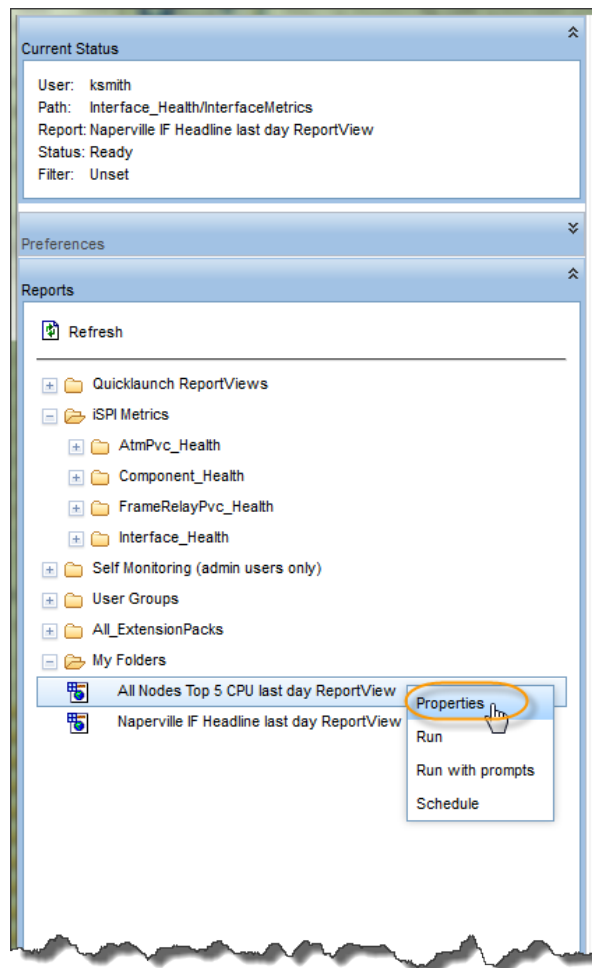


Now create one more Report View called All Nodes Top 5 CPU Last Day ReportView and place it under My Folders.

If you want to delete one of these Report Views, go to BI Server, select My Folders, and then select the report view you want to delete and click Delete.



Also, if you want to change the properties of the report view, go to the Report View, right-click, and then choose Properties.



You can go to the Report View tab and make various changes.

Set properties - All Nodes Top 5 CPU Last Day ReportView

General **Report view** Permissions

Select the default options to use for this entry.

Default action:
View most recent report ▼

Report options:
☐ Override the default values
Formats:
 Default
Accessibility:
 Default
Languages:
 Default

Prompt values:
 auto_refresh: 'No', dow: element_level: '...' [View all](#)
Edit... [Clear](#)
☐ Prompt for values

Run as the owner:
☐ ksmith

[Advanced options](#) ▼

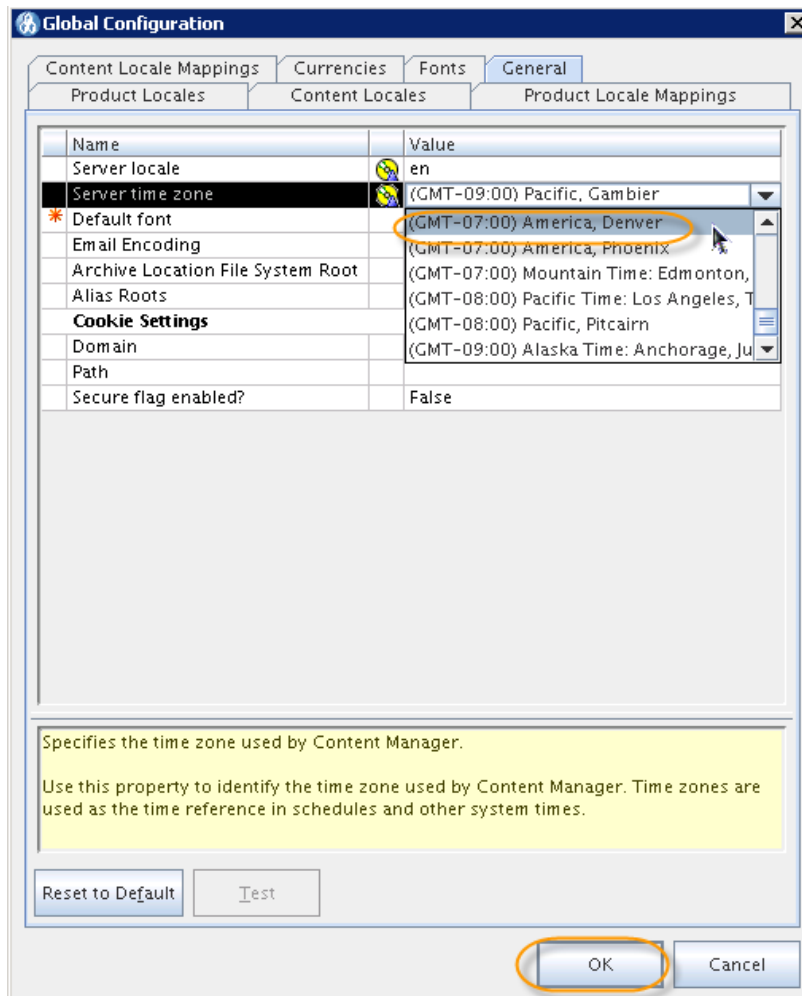
OK Cancel

Now you have easy access to these favorite reports without having to rebuild all the parameters each time.

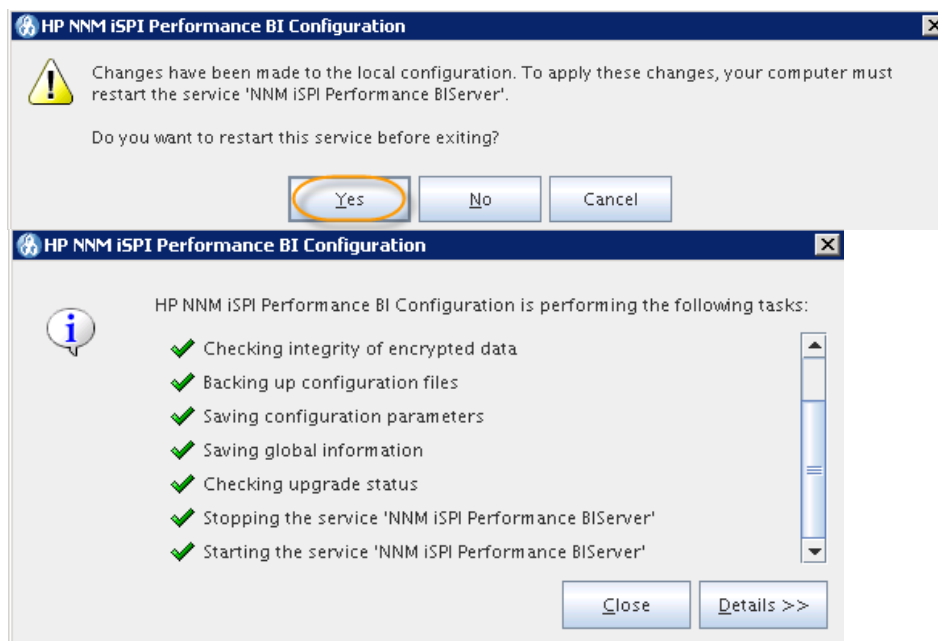
Scheduling Reports

It is a best practice to schedule Report Views rather than Reports because all the parameters must be pre-specified since no human interaction is applicable for a scheduled report. In this deployment example, you will first create a folder to hold scheduled reports. You will then make these Report Views visible to all users.

The default time zone for the BI Server is Eastern Time (US & Canada). This does NOT affect reports, but does affect timestamps used by the BI Server for catalog objects like output version creation time. As a workaround, run the `runBIConfigGUI.ovpl` tool from the command line on the NPS system, and then choose Actions -> Edit Global Configuration, click General Tab, and set the Server time zone.



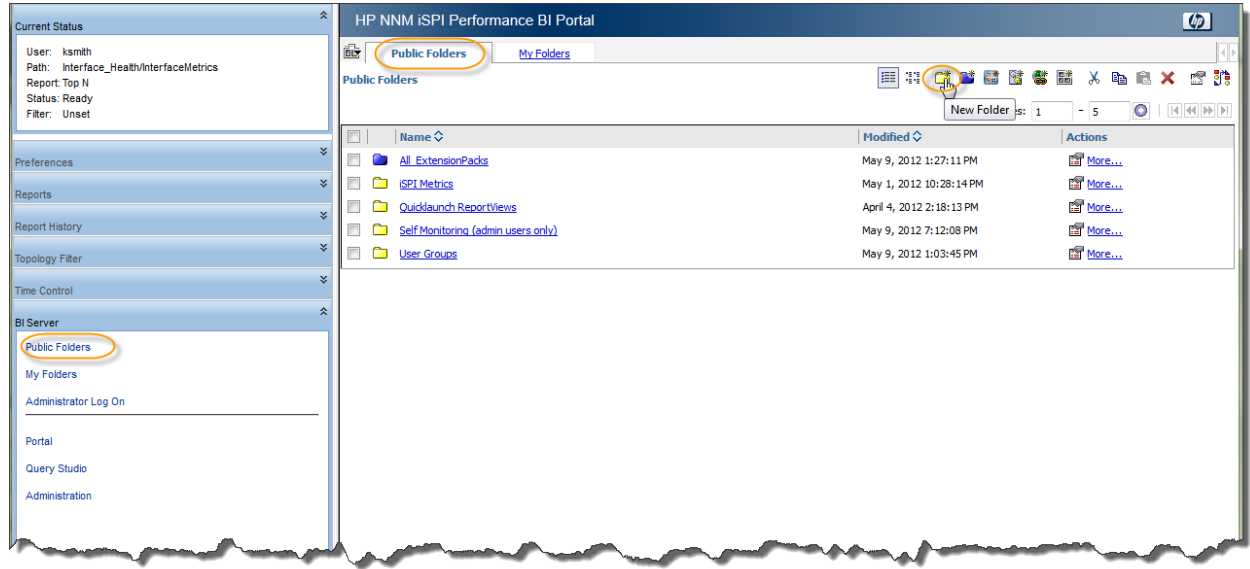
You are then prompted to restart the service. Click Yes.



Now you have the time zone correctly configured.

Creating a Location for the Scheduled Reports

You need to create the location for the scheduled reports.
Go to the BI Server workspace and select Public Folders.



Name the folder, validate the location, and then click Finish.

Specify a name and description - New Folder wizard

Specify a name and location for this entry. You can also specify a description and screen tip.

Name:
Standard Scheduled Reports

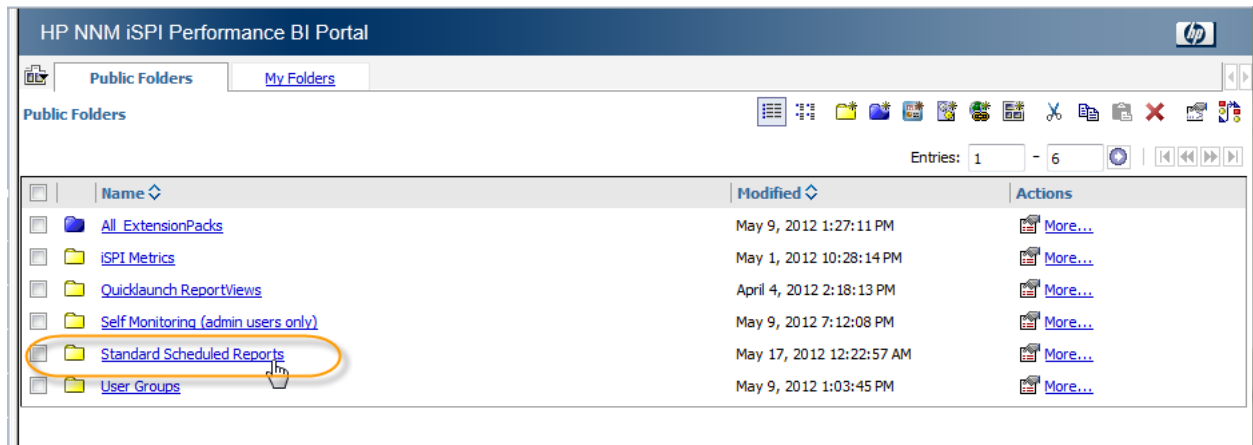
Description:

Screen tip:

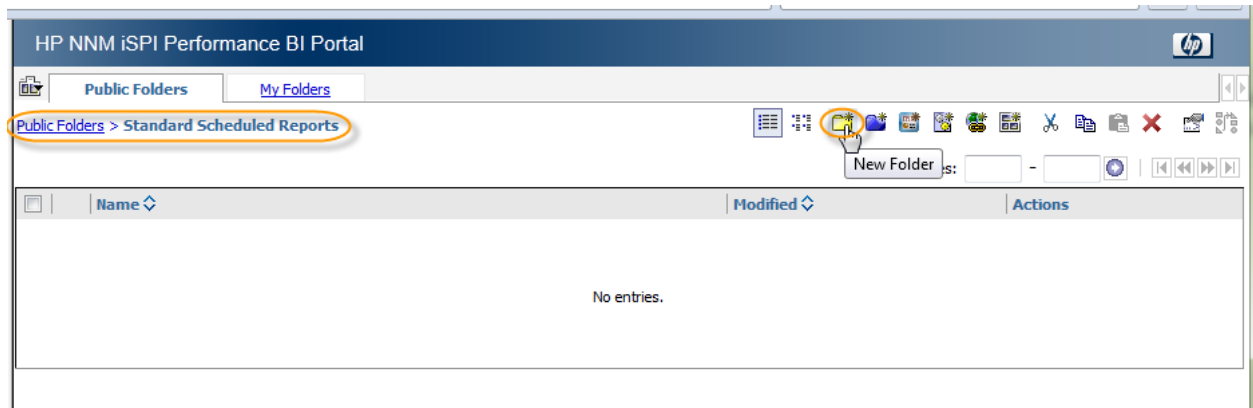
Location:
Public Folders
[Select another location...](#) [Select My Folders](#)

Cancel < Back Next > Finish

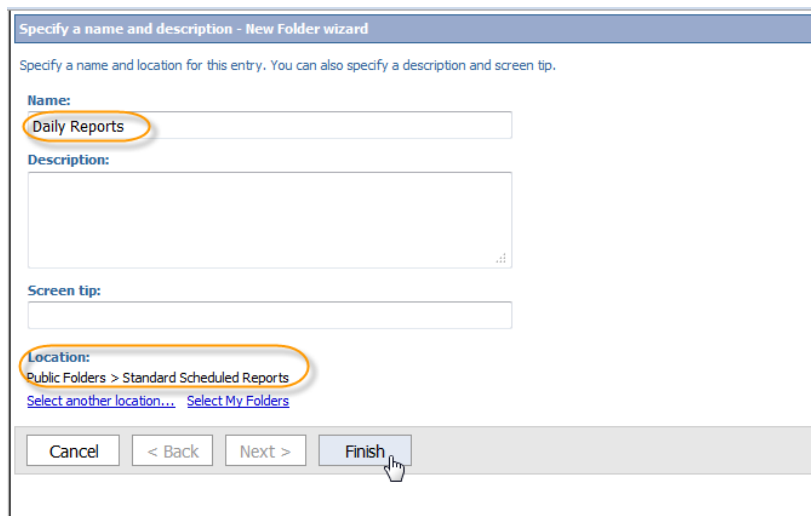
Now create a subfolder inside of this folder for Daily Reports. Click on the Standard Schedule Reports folder.



Click on the New Folder icon.



Name the folder Daily Reports, validate the location, and then click Finish.



Now you have a folder for daily reports. Next, create a couple of report views and place them in this location. The procedure is the same as described earlier. For example, when saving the report view, choose "Select another location."

Save as report view [Help](#) ✕

Specify a name and location for this entry. A report view shares the same report specification as the source report.

Name:

Location:
 Public Folders > ISPI Metrics
[Select another location...](#) [Select My Folders](#)

Then navigate to the Public Folders > Standard Scheduled Reports > Daily Reports folder. Then click OK.

Select a location (Navigate) [Help](#) ✕

Navigate the folders or search to find where you want to place the entry. [Search](#)

[Cognos](#) > [Public Folders](#) > [Standard Scheduled Reports](#) > [Daily Reports](#)

Entries: -

Name
No entries.

[New folder...](#)

Confirm the location, and then click OK.

Save as report view [Help](#) ✕

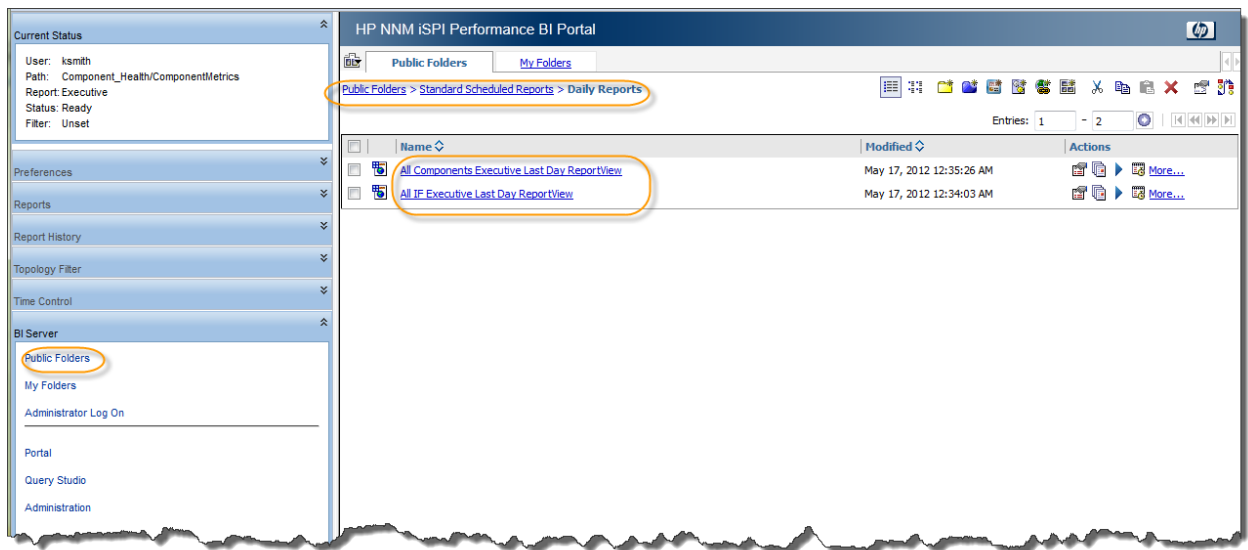
Specify a name and location for this entry. A report view shares the same report specification as the source report.

Name:

Location:
[Public Folders > Standard Scheduled Reports > Daily Reports](#)
[Select another location...](#) [Select My Folders](#)

Repeat this process for a second report view.

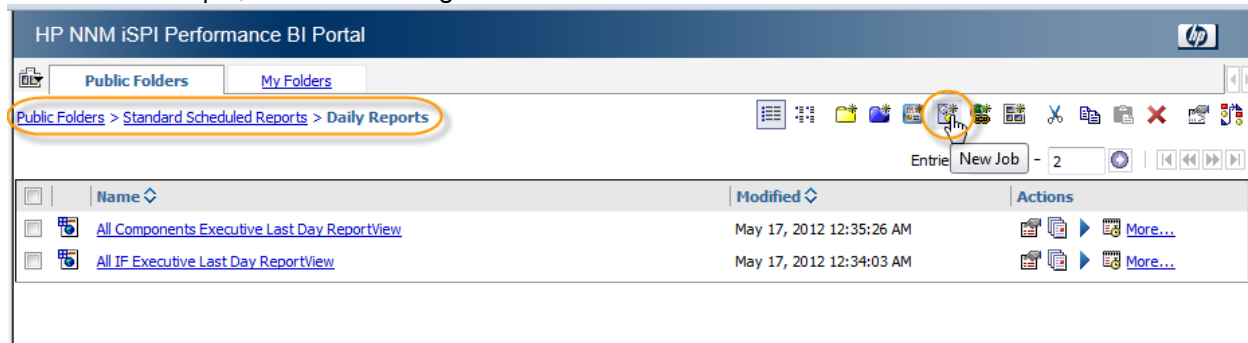
You can confirm that the report views are in the public folders. Go to the BI Server workspace and select Public Folders. Then navigate to the Daily Reports folder and you can see the two report views.



Creating a Job

You will create a “job” for scheduling reports. A “job” is a method of running a collection of reports. When you have a collection of reports that you want to run daily, it is easier to create a daily report job, and then schedule that job rather than scheduling individual reports. Then if you want to move the starting time from 11pm to 1am, you only need to make the change at one place (the job) rather than on each report. Note that Report Views can be individually scheduled without first creating a job.

To create a new job, continue working in the BI Server and click on the New Job icon.



Give the job a name and validate the location, and then click Next.

Specify a name and description - New Job wizard

Specify a name and location for this entry. You can also specify a description and screen tip.

Name:

Description:

Screen tip:

Location:

[Select another location...](#) [Select My Folders](#)

You can change the “submission of steps” to run the reports sequentially rather than all at once. When you do this, it is always a good idea to check the “Continue on error” box or there is a possibility that some of the reports will not run due to a previous error. Then click the Add link to add report views to this job.

Select the steps - New Job wizard Help

Select the entries to include as steps of this job and the options to use when this entry runs.

Steps:

<input type="checkbox"/>	Name	Options and prompt values
No entries.		

[Add...](#)
[Remove](#)
[Modify the sequence...](#)
[Reset to default value](#)

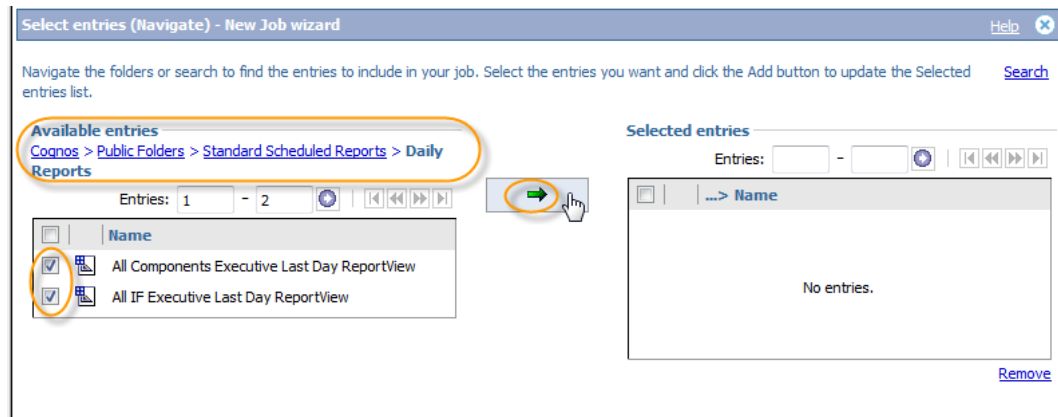
Submission of steps:
 Submitting steps in sequence implies that a step is submitted only upon completion of the step before it.

☐ All at once
☒ In sequence
☒ Continue on error

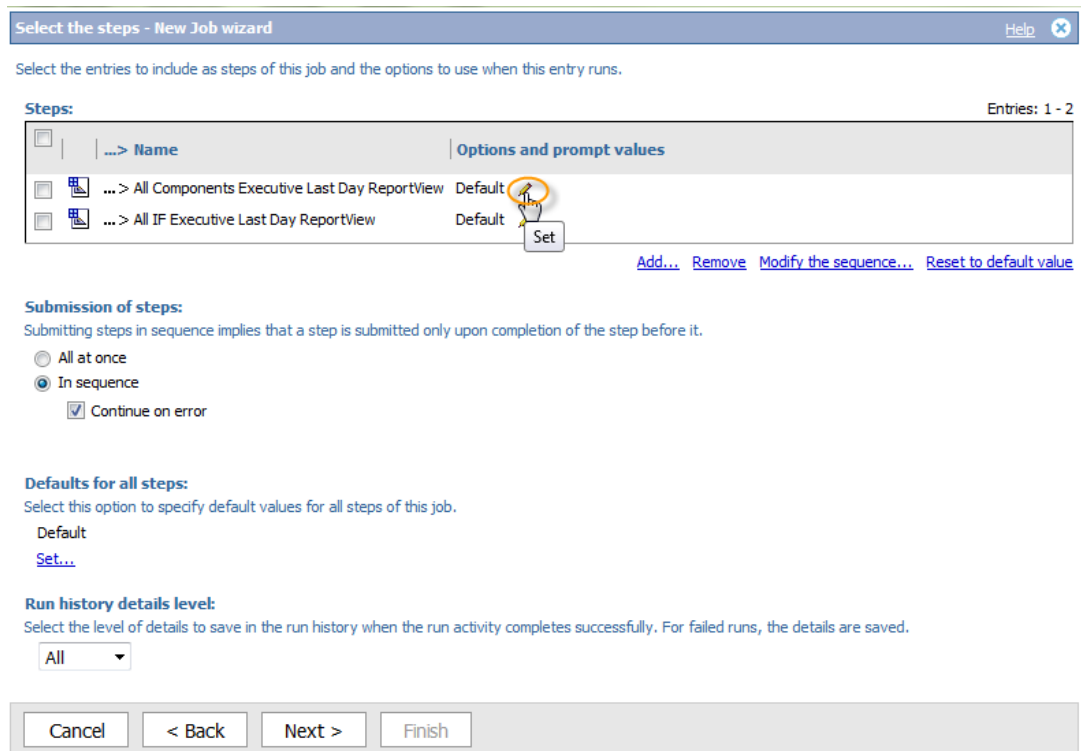
Defaults for all steps:
 Select this option to specify default values for all steps of this job.
 Default
[Set...](#)

Run history details level:
 Select the level of details to save in the run history when the run activity completes successfully. For failed runs, the details are saved.

Select the two Report Views that you want to include in this job and click the arrow.



You can make some changes to the Report View if you choose to. To do this, click on the pencil icon.



Select the “override the default values” checkbox.

Select the report options - New job - All Components Executive Last Day ReportView

Select the run options to use for this step. You can run this step to produce report outputs or to refresh the report pages.

Run the report to:
Produce report outputs

Report options:
☒ Override the default values
Formats:
 Default
Accessibility:
 Default
Languages:
 Default
Delivery:
 Save the report

Prompt values:
☐ Override the default values
 auto_refresh: '0'. Baseline Exception Rate... [View all](#)

OK Cancel

Add PDF to the report output. If you have configured an e-mail server setup, you can choose to send the report via e-mail. See the appendix in this document for instructions on how to do this on Linux.

Select the report options - New job - All Components Executive Last Day ReportView Help

Select the run options to use for this step. You can run this step to produce report outputs or to refresh the report cache to optimize the performance of prompt pages.

Run the report to:
Produce report outputs

Report options:
☒ Override the default values
Formats:
☒ HTML
 Number of rows per Web page: 20
☒ Enable selection-based interactivity
☒ PDF
 No options saved
[Set...](#)
☐ Excel 2007
☐ Excel 2002
☐ Delimited text (CSV)
☐ XML

Accessibility:
☐ Enable accessibility support

Prompt values:
☐ Override the default values
 auto_refresh: '0'. Baseline Exception Rate... [View all](#)

Languages:
 English (United States) [Select the languages...](#)

Delivery:
 Select at least one delivery method. For burst reports, the email recipients are determined by the burst specification.
☒ Save the report
☐ Print the report
Printer location:
 [Select a printer...](#)
☒ Send a link to the report by email [Edit the options...](#)
 0 recipients

OK Cancel

There are some options for e-mail. Chose to attach the report and an e-mail destination.

Set the email options - New job - All Components Executive Last Day ReportView Help

Specify the recipients and contents of the email. To add recipients, click [Select the recipients...](#) or type the email addresses separated by semi-colons. To include an HTML report as the message body, leave the Body box empty and select the report as the only attachment.

To:

Cc:

[Select the recipients...](#) [Show Bcc](#)

Subject:

Body: [Change to plain text >>](#)

☐ Include a link to the report
☒ Attach the report

After repeating this for both the reports, check the options and prompt values, and then click Next.

Select the steps - New Job wizard Help

Select the entries to include as steps of this job and the options to use when this entry runs.

Steps: Entries: 1 - 2

...	Name	Options and prompt values
<input checked="" type="checkbox"/>	... > All Components Executive Last Day ReportView	Produce report outputs, HTML (20 rows per page, Enable selection-based interactivity), PDF, English (United States), No prompt values saved, Save the reports, Send the report by email: 1 recipient ✖
<input checked="" type="checkbox"/>	... > All IF Executive Last Day ReportView	Produce report outputs, HTML (20 rows per page, Enable selection-based interactivity), PDF, English (United States), No prompt values saved, Save the reports, Send the report by email: 1 recipient ✖

[Add...](#) [Remove](#) [Modify the sequence...](#) [Reset to default value](#)

Submission of steps:
 Submitting steps in sequence implies that a step is submitted only upon completion of the step before it.

☐ All at once
☒ In sequence
☒ Continue on error

Defaults for all steps:
 Select this option to specify default values for all steps of this job.
 Default
[Set...](#)

Run history details level:
 Select the level of details to save in the run history when the run activity completes successfully. For failed runs, the details are saved.

Finally, choose the "Save and run once" action and click Finish.

Select an action - job

Select whether you want to run, schedule, or save only, when the wizard closes.

Action:

- ☒ Save and run once
- ☐ Save and schedule
- ☐ Save only

Cancel < Back Next > Finish

You are then prompted to run the job. Choose to run it now.

Run with options - Daily Reports Job

Specify when you want to run this job.

Time:

☒ Now

☐ Later:

May 17, 2012

8 : 26 PM


Steps:

...	Name
...	> All Components Executive Last Day ReportView
...	> All IF Executive Last Day ReportView

Run Cancel

You are then presented with the opportunity to watch the details of this job while it is running. Select the checkbox and click OK.

HP NNM iSPI Performance BI software

 You selected to run 'Daily Reports Job' as follows:
Time: now

☒ View the details of this job after closing this dialog

Click OK to submit the job or click Cancel to return to your selection.

☒ Show this dialog in the future

OK Cancel

You can see the status of the job. Clicking on the Refresh link at the top left of the page helps you follow along better.

View run history details - Daily Reports Job
Help

View the details of this particular run.
Refresh

Start time:
May 17, 2012 8:28:33 PM

Completion time:
Unavailable

Status:
Waiting

Messages

Severity: (All)
Entries: 0 - 0

Time	Message
No entries.	

Job

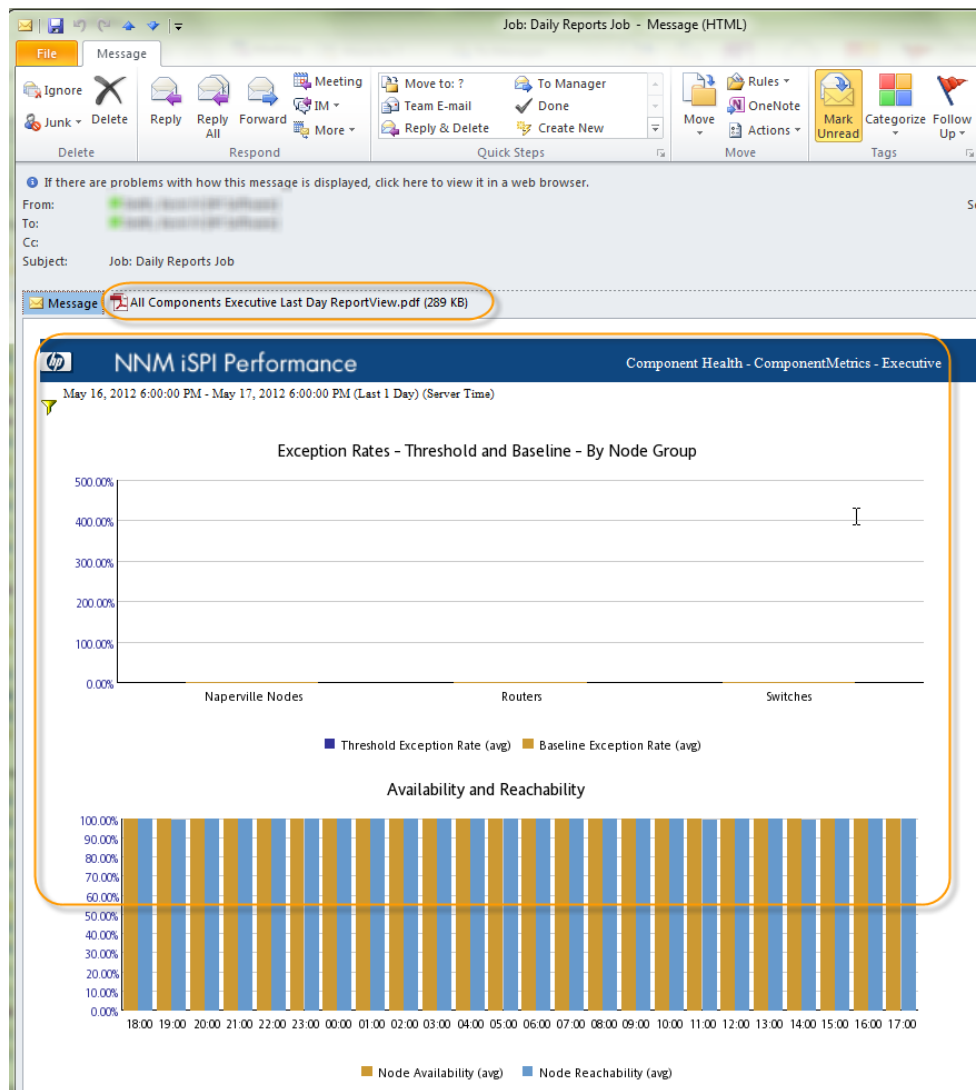
Steps:

Entries: 1 - 2

Name	Request time	Start time	Completion time	Status	Actions
...> All Components Executive Last Day ReportView	May 17, 2012 8:28:35 PM	May 17, 2012 8:28:41 PM	May 17, 2012 8:29:08 PM	Succeeded	
...> All IF Executive Last Day ReportView	May 17, 2012 8:28:35 PM			Pending	

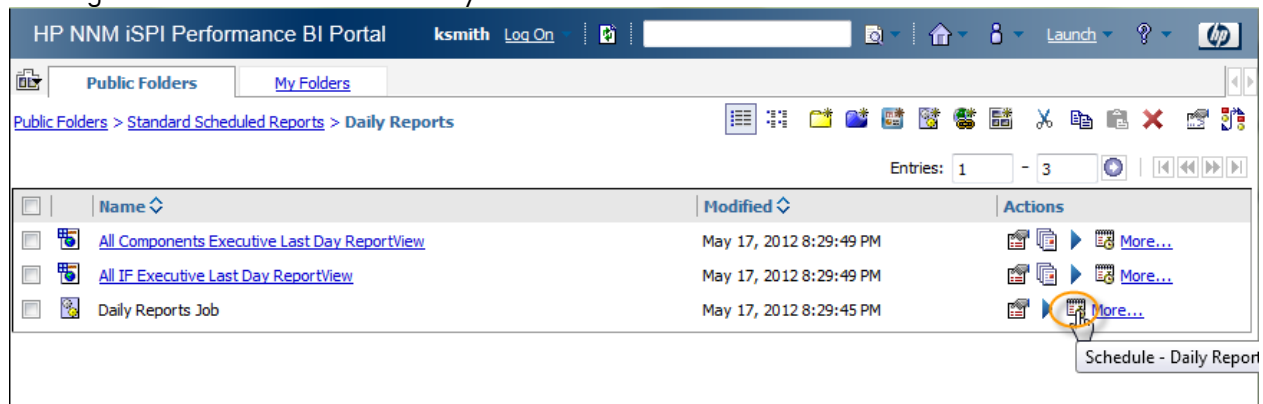
Close

Eventually, you will see both Report Views succeed. You will also receive an e-mail with the attached PDF report as well as the HTML report in the message.

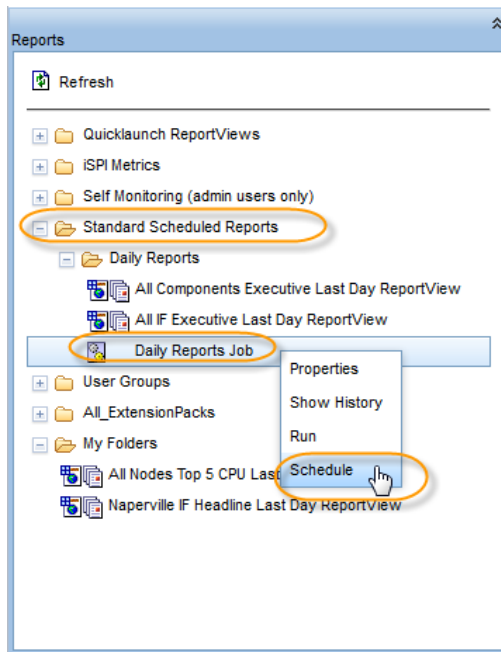


Scheduling the Job

Now that you have created the job and have run it once, schedule it to run daily. There are multiple places where you can schedule it. One place is via the BI Portal page where you have been working. There is a schedule icon that you can click.



Another handy location is the Reports workspace. After making previous changes, if you click Refresh, you can see a new folder based on our previously built structure. You can go to the Daily Reports Job, right-click, and then choose Schedule.



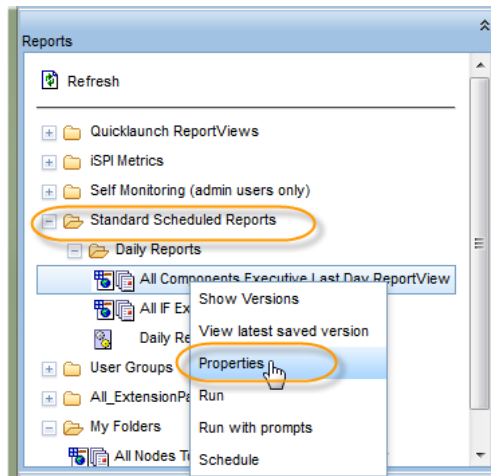
This launches the scheduling page. Validate the frequency and the start time and click OK. For testing purposes, temporarily set the frequency to be every 3 minutes just so you can test the scheduling easier.

Usually this is set to "Every 1 day" as shown below.

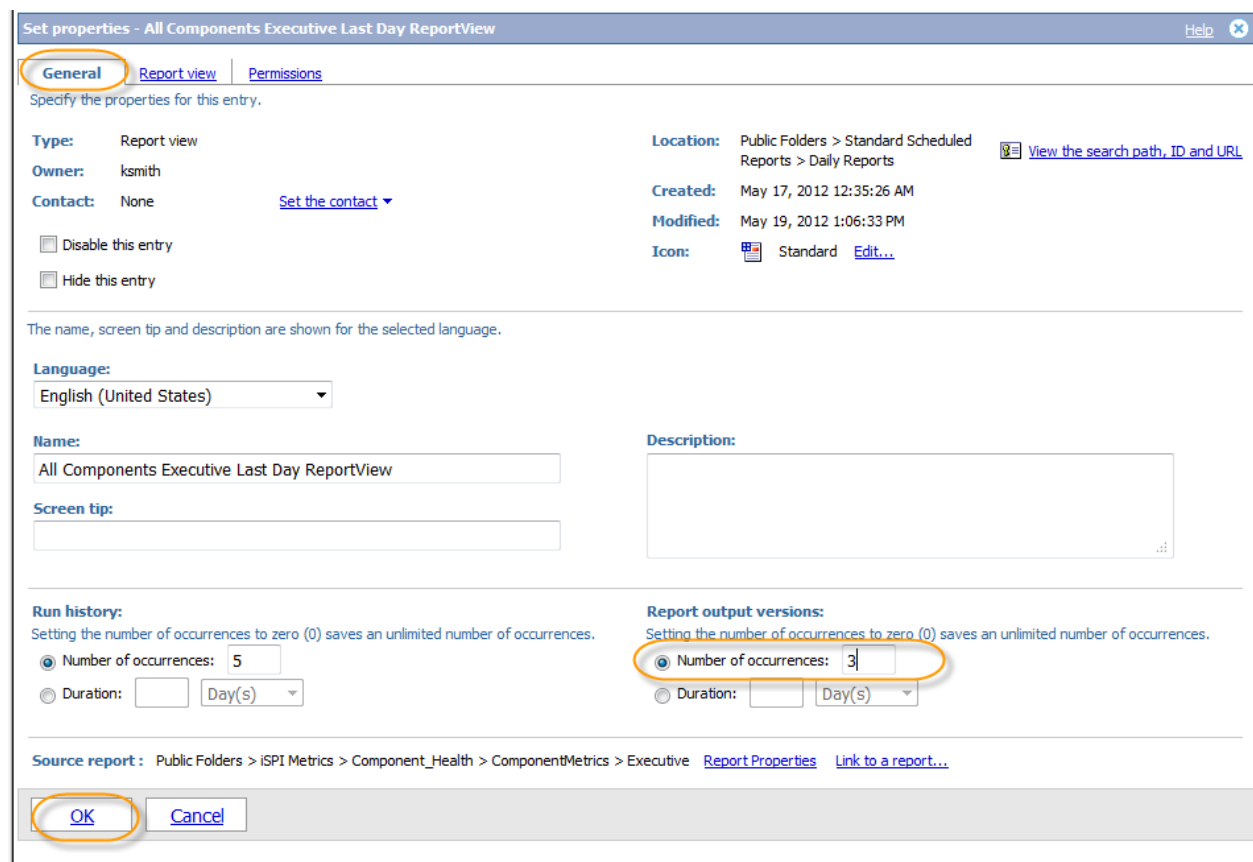
 A screenshot of the 'Schedule - Daily Reports Job' dialog box. The dialog has a title bar with 'Schedule - Daily Reports Job' and a 'Help' button. Below the title bar is a description: 'Schedule the entry to run at a recurring date and time. You can disable the schedule without losing any of its details.' There are several sections:

- Disable the schedule:** A checkbox that is currently unchecked.
- Priority:** A dropdown menu set to '3'.
- Frequency:** A section with a link 'Select the frequency by clicking on a link.' Below it are five buttons: 'By Day', 'By Week', 'By Month', 'By Year', and 'By Trigger'. The 'By Day' button is highlighted with an orange circle.
- Frequency options:** Three radio buttons: 'Every 1 minute(s)', 'Every 1 hour(s)', and 'Every 1 day(s)'. The 'Every 1 day(s)' option is selected and highlighted with an orange circle.
- Daily Frequency:** A section with a checkbox 'Every 1' and a dropdown 'Hour(s)' set to 'Hour(s)'. It also has 'between' and 'and' fields with time values '9 : 00 AM' and '5 : 00 PM'.
- Credentials:** A text field containing 'ksmith (ksmith)'.
- Steps:** A section with a list of report views: '... > All Components Executive Last Day ReportView' and '... > All IF Executive Last Day ReportView'.
- Start/End times:** Two sections. The 'Start' section has a date 'May 19, 2012' and a time '11 : 39 AM' (highlighted with an orange circle). The 'End' section has a radio button for 'No end date' (selected) and a radio button for 'End by:' with a date 'May 19, 2012' and a time '11 : 39 AM'.
- Buttons:** At the bottom are 'OK' and 'Cancel' buttons. The 'OK' button is highlighted with an orange circle.

Change the settings of the Report Views to keep three copies of older reports on the server. Return to the Reports workspace, navigate to Standard Scheduled Report > Daily Reports, and then right-click on the Report View and select Properties.

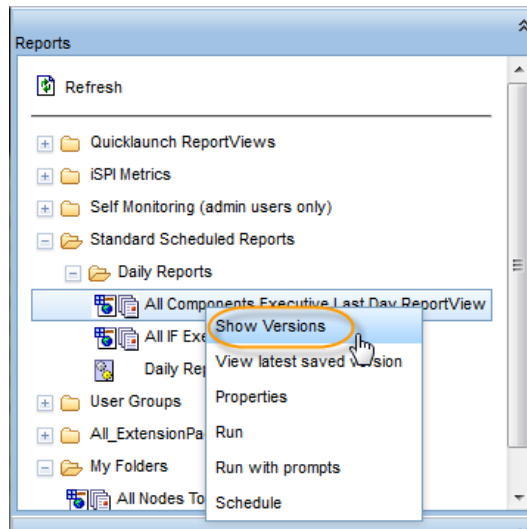


Then select the General tab and increase the “number of occurrences” value under Report output versions and click OK.

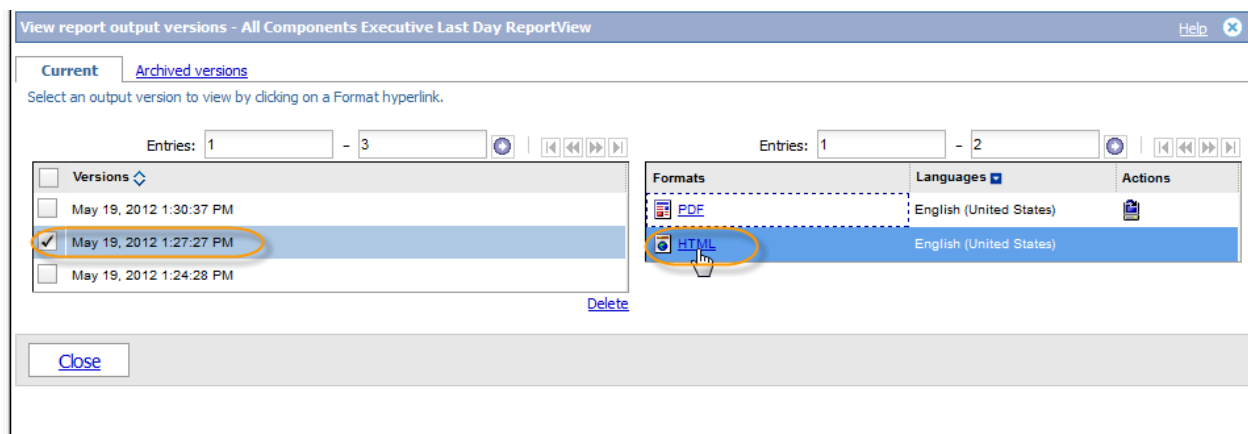


Viewing Versions of the Report Views

Now let it run for a while. Remember that you are running this report view every 3. Go back to the Report View, right-click and choose Show Versions.

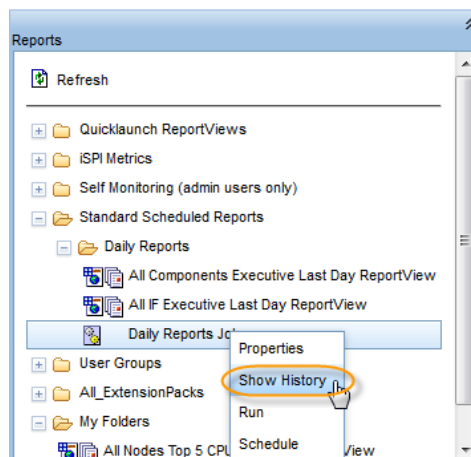


You can now see three current versions (last three). They are each about three minutes apart. You can then click on the preferred output and view the report.



Showing the Job Run History

To see a historical report of the job run history, right-click on the Daily Reports Job and select Show History.



You can see the last five runs of the job. Note that you can control how many status reports are stored in the history via the same job properties where you changed the number of stored output versions. You can also click on the action icon to view the run history details if you like.

View run history - Daily Reports Job





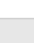
Help

View the run history for this entry.

Refresh

Status:
All statuses

Entries: 1 - 5

Request time	Start time	Completion time	Status	Actions
May 19, 2012 1:39:02 PM	May 19, 2012 1:39:11 PM	May 19, 2012 1:39:54 PM	Succeeded	
May 19, 2012 1:36:02 PM	May 19, 2012 1:36:09 PM	May 19, 2012 1:36:51 PM	Succeeded	
May 19, 2012 1:33:04 PM	May 19, 2012 1:33:11 PM	May 19, 2012 1:33:53 PM	Succeeded	
May 19, 2012 1:30:03 PM	May 19, 2012 1:30:14 PM	May 19, 2012 1:31:02 PM	Succeeded	
May 19, 2012 1:27:02 PM	May 19, 2012 1:27:09 PM	May 19, 2012 1:27:52 PM	Succeeded	

View run history details - Daily Reports Job

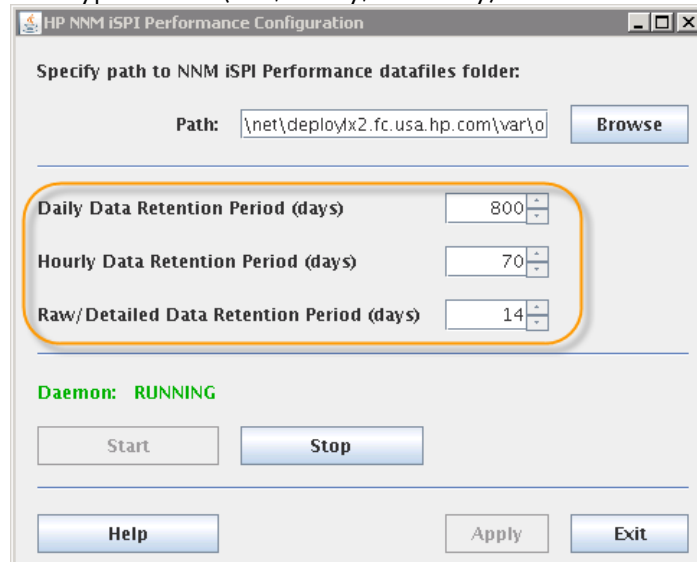
Close

Maintenance

Tracking the Database Size

It is important to track the size of the database to understand disk storage usage. The database size is influenced by the number of polled objects and the retention configuration (in other words, how long you choose to store the various metrics).

To view and change the storage retention, run the script `runConfigurationGUI.ovpl` on the NNM iSPI Performance for Metrics server. You can see the number of days of retention you have configured for each type of data (raw, hourly, and daily).



For guidance on how much disk space to anticipate using for various retention periods, see the *NNM iSPI Performance for Metrics Support Matrix*.

An example of the System Requirements table in the *NNM iSPI Performance for Metrics Support Matrix* is shown below.

Note: Always use the latest version of the Support Matrix available at:
<http://h20230.www2.hp.com/selfsolve/manuals>

iSPI Minimum System Requirements

Management Environment Size		iSPI Minimum System Requirements						
Approximate managed environment tier	Number of performance polled interfaces / components at 5 minute polling interval	CPU (64-bit) x86-64 AMD64	RAM	Disk space for Application installation (NNMInstallDir)	Disk space for database and data during execution (NNMDataDir)	Disk space for database and data during execution (NNMDataDir) 14 day raw/detailed retention 70 day hourly retention 800 day daily retention	Disk space for database and data during execution (NNMDataDir) 70 day raw/detailed retention 800 day daily retention	Disk space for database and data during execution (NNMDataDir) 70 day raw/detailed retention 800 day daily retention
Small	Up to 5K/5K	8 CPU (2.5GHz for x64)	16 GB	10 GB	1 SCSI or SATA disk drive	300 GB	400 GB	1 TB
Medium	Up to 60K/60K	8 CPU (2.5GHz for x64)	24 GB	10 GB	RAID 1+0 or 5/6 with write cache recommended	800 GB	1.5 TB	4 TB
Large	Up to 130K/130K	16 CPU (2.5GHz for x64)	48 GB	10 GB	RAID 1+0 or 5/6 with write cache recommended	2 TB	3 TB	10 TB
Very Large	Up to 400K/200K	16 CPU (2.5GHz for x64)	72 GB	10 GB	High performance SAN storage	4 TB	8 TB	20 TB

In the table, you can cross-reference the number of objects against the period for retention and come up with anticipated disk space for the data directory (which is mostly composed of the database).

To check our current database size, run the command `dbsize.ovpl` on the NPS system.

```
# dbsize.ovpl
[1211] /opt/OV/NNMPerformanceSPI/bin/dbsize.ovpl(17) INFO: Starting dbsize.ovpl.. checking
for concurrent executions...
[1211] /opt/OV/NNMPerformanceSPI/bin/dbsize.ovpl(21) INFO: Continuing...
DBSPACE Usage Summary for Database: "DSN=PerfSPIDSN"
=====
DbSpace      Size      Res      Used(%)
IQ_SYSTEM_MAIN 1.59G    32G      21
IQ_SYSTEM_TEMP 10G      64G       1
USER_MAIN      5G      198G     37
[1211] /opt/OV/NNMPerformanceSPI/bin/dbsize.ovpl(187) INFO: Done.
```

There are three DB spaces described in this output.

The USER_MAIN space is the location of all the stored metrics in NNM iSPI Performance for Metrics; it is also important to watch the other spaces, especially the IQ_SYSTEM_TEMP space. This temporary space is used when executing a report and can become large under certain circumstances.

The Size column is the amount of disk space that is presently pre-allocated for the database space. As the database contents increase, this space increases. See the man page for `dbsize.ovpl` for details on this algorithm. In this case, the size is presently at 5G. This can be seen on the file system on Linux by going to `/var/opt/OV/NNMPerformanceSPI/database` and looking at these files:

```
-rw-r--r-- 1 root root 1.0G May 20 07:50 perfspi_MAIN_01.iq
-rw-r--r-- 1 root root 1.0G May 20 07:50 perfspi_MAIN_02.iq
-rw-r--r-- 1 root root 1.0G May 20 07:50 perfspi_MAIN_03.iq
-rw-r--r-- 1 root root 1.0G May 20 07:50 perfspi_MAIN_04.iq
-rw-r--r-- 1 root root 1.0G May 20 07:50 perfspi_MAIN_05.iq
```

You can see that these files account for 5G of disk space.

The Used(%) column shows how much space is currently consumed by USER_MAIN. In this case, it is 37% full. About 1.8G of the allocated 5G space is being used for metrics data.

If you want to see a more detailed breakdown of the usage, run `dbsize.ovpl -s`. This output includes a breakdown of the data including daily, hourly, and raw retention. This can help you better anticipate database growth because some values like the raw data will typically hit “steady-state” quickly (in this case, 14 days) and then will not experience much growth beyond that as old raw data is discarded after 14 days.

```
f_Day_AtMvncMetrics      0 MB      0 Rows
f_Day_ComponentMetrics   80 MB    7071 Rows 2012-05-09 -> 2012-05-19
f_Day_DiagnosticMetrics  28 MB    385 Rows 2012-05-09 -> 2012-05-16
f_Day_FrameRelayPvcMetrics 0 MB      0 Rows
```

f_Day_InterfaceMetrics	78 MB	6534 Rows	2012-05-09 -> 2012-05-19
f_Hour_AtmPvcMetrics	0 MB	0 Rows	
f_Hour_ComponentMetrics	128 MB	165068 Rows	2012-05-09 -> 2012-05-20
f_Hour_DiagnosticMetrics	29 MB	5829 Rows	2012-05-09 -> 2012-05-16
f_Hour_FrameRelayPvcMetrics	0 MB	0 Rows	
f_Hour_InterfaceMetrics	142 MB	143951 Rows	2012-05-09 -> 2012-05-20
f_Raw_AtmPvcMetrics	0 MB	0 Rows	
f_Raw_ComponentMetrics	716 MB	2279556 Rows	2012-05-09 -> 2012-05-20
f_Raw_DiagnosticMetrics	25 MB	38795 Rows	2012-05-09 -> 2012-05-16
f_Raw_FrameRelayPvcMetrics	0 MB	0 Rows	
f_Raw_InterfaceMetrics	574 MB	1858662 Rows	2012-05-09 -> 2012-05-20

Backup

The NNM iSPI Performance for Metrics provides you with the `backup.ovpl` tool. This is run on the NPS system. There are multiple options available. See the man page for more details.

In this example, the `-b`, `-d`, `-t` and `-x` options are used.

```
-b dir specifies backup directory
-d backup database
-t suppress creation of single backup tar file
-x exclude demo database from backup [default: false]
```

When you run this, it creates a directory `/var/tmp/backup/backup.20120520081844`. The space used for this directory is approximately the same as the `dbsize.ovpl` usage percentage. As the database gets larger, the `dbsize` percentage more closely tracks the backup size.

```
# du -h backup.20120520081844
2.3G    backup.20120520081844
```

You can assume you need about the same size disk space for backup as for the database usage.

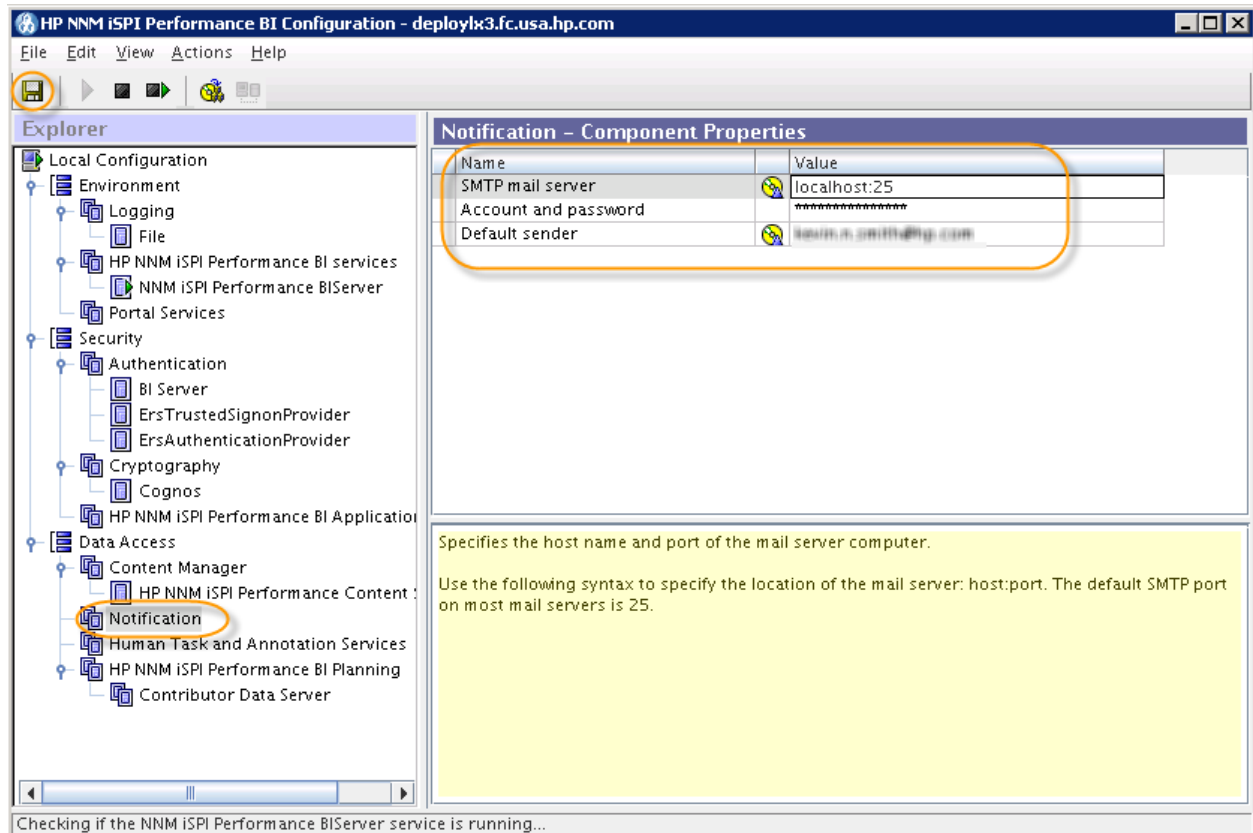
Note: If you compress the output of `backup.ovpl` by omitting the `-t` option, the compression process can take a very long time to run. It can take hours to compress a multi-terabyte file. Also note that running backup can add a fairly high load to the NPS system.

Note the `-d` option can take a directory as well as a file.

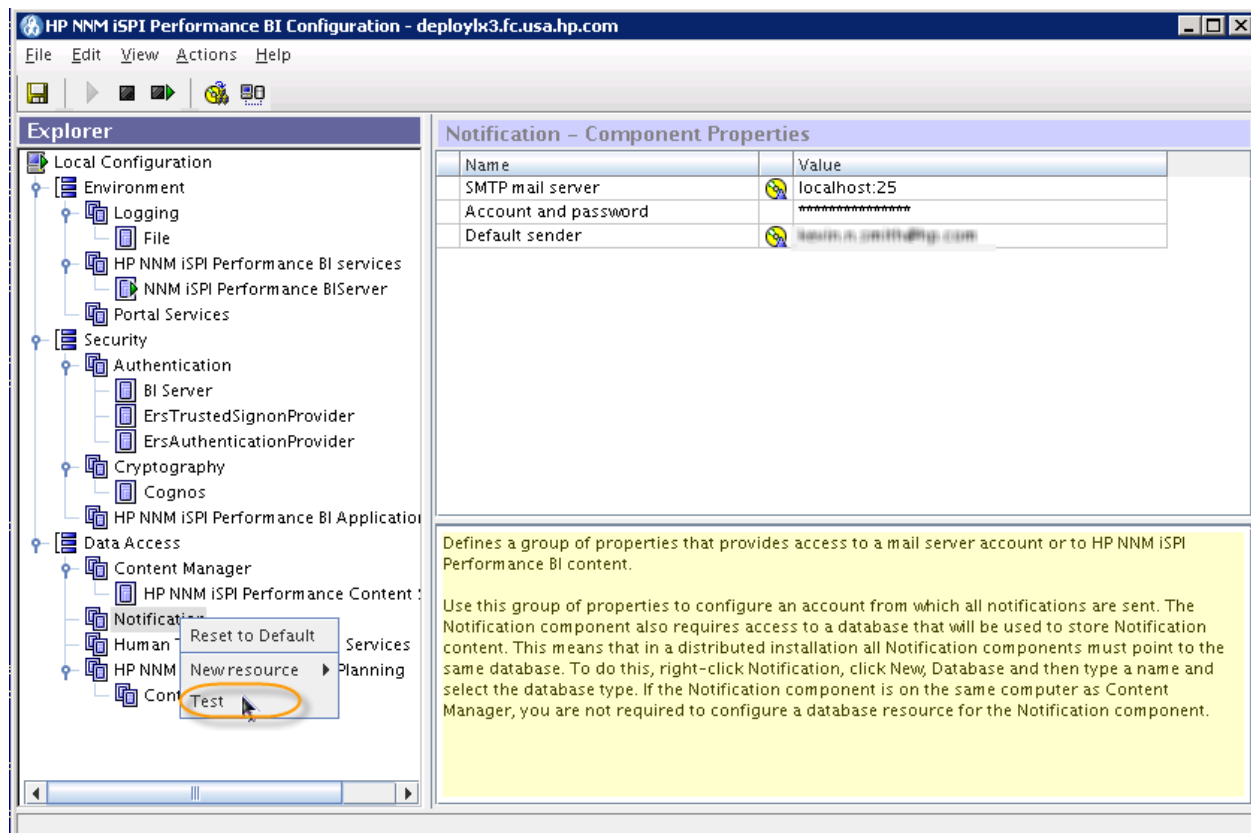
Appendix A: E-mailing Reports on Linux

If you want to use NPS to email reports, you must configure the BI Server to use a Simple Mail Transfer Protocol (SMTP) email server.

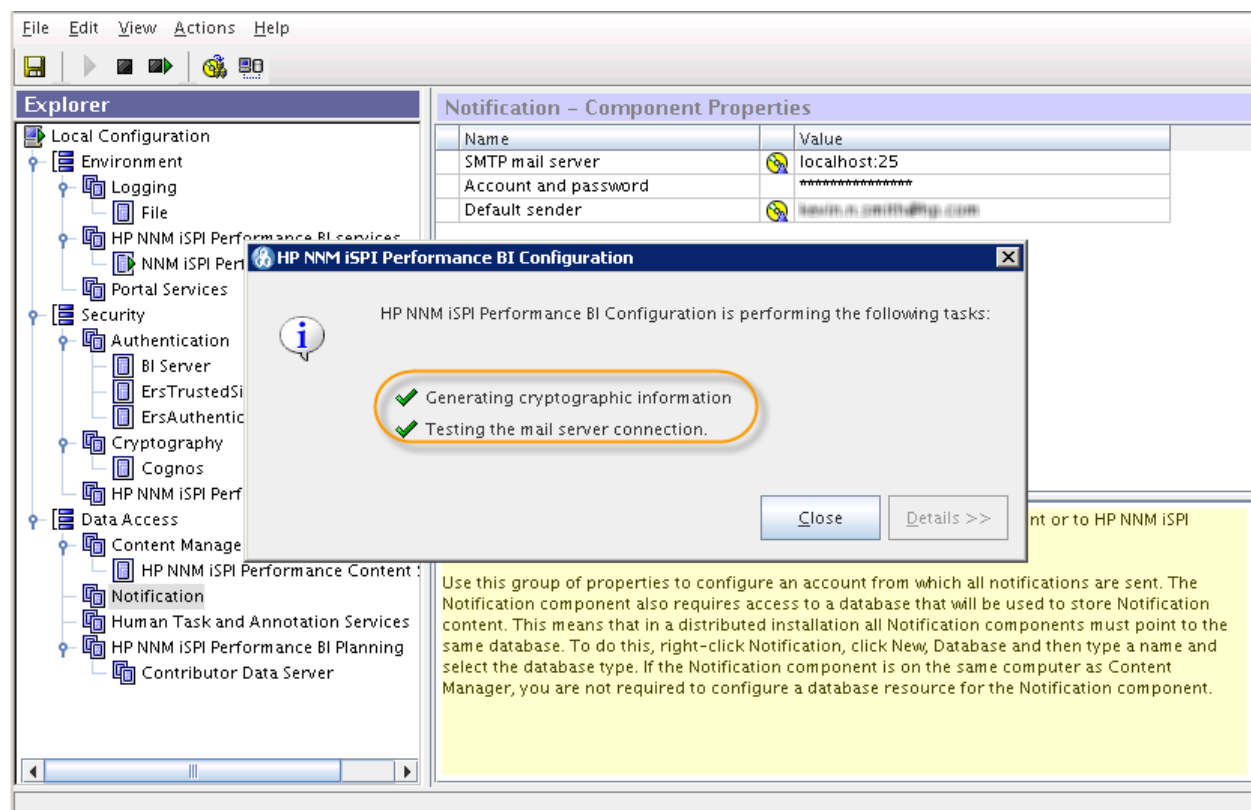
1. Launch the BI Server Configuration graphical user interface:
`/opt/OV/NNMPerformanceSPI/bin/runBIConfigGUI.ovpl`
2. In the Explorer pane, select Notification.



1. Specify appropriate values for the following fields:
 - SMTP Mail server (usually localhost:25 is a good choice on Linux)
 - Account and password
 - Default Sender
2. Click Save. To test the connection, right-click Notification in the left pane, and then click Test.

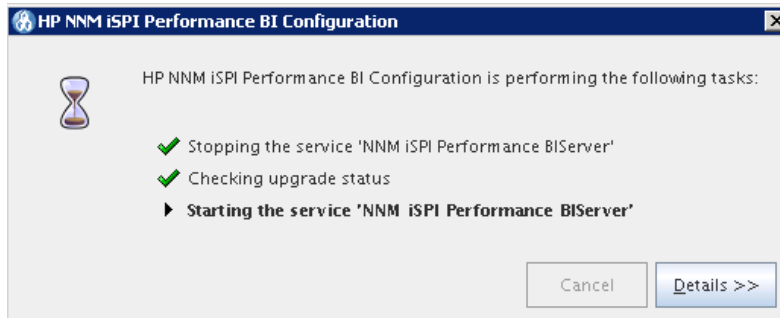


The following message appears on success:



3. Click Close and File > Save.

4. Select Yes when prompted to restart the service NNM iSPI Performance BI Server.

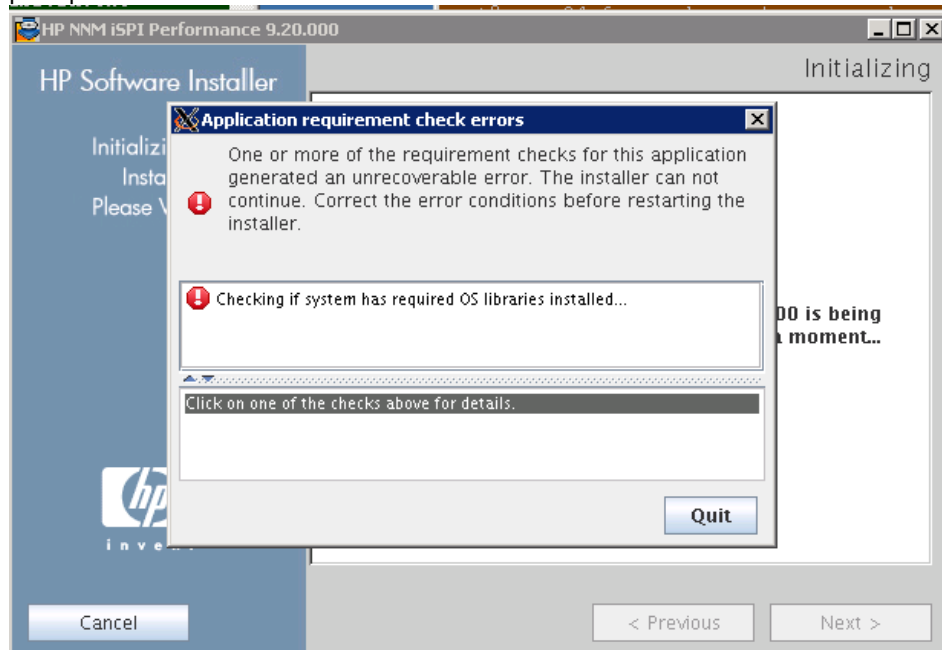


Now reports can be e-mailed via the NPS system.

Tip: You can also set up the email server with the help of the `configureBIEmailServer.ovpl` script.

Appendix B: Installation with Missing Libraries

If you are missing some prerequisite libraries on the system during the installation, you will see a prompt similar to the one shown below.



Click Quit to exit the installer.

The details of this error suggest you to go to the `/var/tmp` directory and look at the `PerfSPI_AppCheckReqdLibs.sh.txt` file. This file shows the libraries that are missing along with a “yum” command for installing the missing packages. If you have a yum server setup, you can run the command from the log file to install the packages with the libraries.

```
# cat PerfSPI_AppCheckReqdLibs.sh.txt
=====
Wed Apr 18 21:38:10 MDT 2012
OS: Linux
-----
compat-libstdc++-296-2.96-138
compat-libstdc++-33-3.2.3-61
compat-libstdc++-33-3.2.3-61
libjpeg-6b-37
libjpeg-6b-37
libpng-1.2.10-7.1.el5_3.2
libpng-1.2.10-7.1.el5_3.2
libXp-1.0.0-8.1.el5
libXp-1.0.0-8.1.el5
ncurses-5.5-24.20060715
ncurses-5.5-24.20060715
openmotif22-2.2.3-18
openmotif22-2.2.3-18
tcsh-6.14-14.el5
package unixODBC.i386 is not installed
INFO: Required library not installed: unixODBC.i386
package unixODBC.x86_64 is not installed
INFO: Required library not installed: unixODBC.x86_64
package unixODBC-devel.i386 is not installed
INFO: Required library not installed: unixODBC-devel.i386
package unixODBC-devel.x86_64 is not installed
INFO: Required library not installed: unixODBC-devel.x86_64
INFO: Run the following command, or otherwise install the identified missing libraries
```

```
INFO: #yum install unixODBC.i386 unixODBC.x86_64 unixODBC-devel.i386 unixODBC-devel.x86_64
ERROR: Please install the required libraries, and then re-run the installer (Logfile
/var/tmp/PerfSPI_AppCheckReqdLibs.sh.txt)
```

After successfully running the yum command to install the libraries, run the NNM iSPI Performance for Metrics installer again and it will succeed.

© 2012 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

