



Intel® Omni-Path Fabric Suite Fabric Manager GUI

User Guide

Rev. 8.0

October 2017



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit <http://www.intel.com/design/literature.htm>.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

No computer system can be absolutely secure.

Intel, the Intel logo, Intel Xeon Phi, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2015–2017, Intel Corporation. All rights reserved.



Revision History


For the latest documentation, go to <http://www.intel.com/omnipath/FabricSoftwarePublications>.

Date	Revision	Description
October 2017	8.0	Updates to this document include: <ul style="list-style-type: none"> The <i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i> has been merged into the <i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>. The only change to this document was in the Preface section. See the Intel® Omni-Path Documentation Library for details.
August 2017	7.0	Updates to this document include: <ul style="list-style-type: none"> Updated Preface section "Intel® Omni-Path Documentation Library". Changed terminology in Home Tab and Performance Tab topics from "failed ports" to "no response ports" Added reference to Appendix SL Key Creation for Fabric Manager GUI in <i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i> in Section 2.0. Minor changes to Troubleshooting Q & A section, Q1 and Q5.
April 2017	6.0	Updates to this document include: <ul style="list-style-type: none"> Updated Preface to include a new section "Intel® Omni-Path Documentation Library". Minor change to Troubleshooting Q & A section, Q1.
December 2016	5.0	Updates to this document include: <ul style="list-style-type: none"> Updated Performance Tab Example figure Updated Port Performance Example figure Added Appendix "Port Counters Overview"
July 2016	4.0	Updates to this document include: <ul style="list-style-type: none"> Added description for zooming options for performance charts Added Switch Properties, Partition Key Enforcements, and Forwarding Information Added Routing Information for IP Address, Routing Mode, and Adaptive Routing
May 2016	3.0	Updates to this document include: <ul style="list-style-type: none"> Added new Subnet Performance Example figure Added content to Trend chart and Top <i>N</i> chart descriptions
February 2016	2.0	No updates for this release.
November 2015	1.0	Initial release of this document.



Contents


Revision History	3
Preface	9
Intended Audience	9
Intel® Omni-Path Documentation Library	9
Cluster Configurator for Intel® Omni-Path Fabric	11
Documentation Conventions	11
License Agreements	12
Technical Support	12
1.0 Introduction	13
2.0 Intel® Omni-Path Fabric Suite Fabric Manager GUI Setup Wizard	15
2.1 Setting up Email Notification	17
2.2 Email Notification of Events	21
3.0 Fabric Manager GUI Window Navigation	24
3.1 Menu Bar	24
3.1.1 Connect to a Subnet	25
3.2 Tool Bar	25
3.3 Main Panel	25
3.4 Event Table	25
3.5 Pin Board	26
4.0 Home Tab	28
4.1 Subnet Summary (Home)	28
4.1.1 Subnet Statistics	28
4.1.2 Subnet Status	29
4.1.3 Health Trend	29
4.1.4 Worst Nodes	31
4.2 Subnet Performance	31
4.2.1 Utilization Group	32
4.2.2 PM Counter Categories Group	33
5.0 Performance Tab	35
5.1 Resource Tree	35
5.2 Tree Search	36
5.3 Summarized Performance	36
5.3.1 Subnet Summary (Performance)	36
5.3.2 General Summary (Performance)	37
5.4 Node Performance	39
5.4.1 Performance (Node)	39
5.4.2 Connectivity (Node)	43
5.4.3 Node Properties	46
5.5 Port Performance	50
5.5.1 Performance (Port)	51
5.5.2 Connectivity (Port Performance)	54
5.5.3 Port Properties	54



6.0 Topology	57
6.1 Resource Tree (Topology)	58
6.2 Outline	58
6.3 Selected Resources	58
6.4 Local Topology	58
6.4.1 Device Set (Topology)	59
6.4.2 Device Node	60
6.4.3 Links	61
6.4.4 Routes	61
7.0 Admin Tab	63
7.1 Applications Management	63
7.2 Device Group Management	65
7.3 Virtual Fabric Management	65
7.4 Interactive Console	66
7.5 Log View	69
8.0 Troubleshooting Guide	73
8.1 Overview	73
8.2 General Error Prevention	74
8.3 Best Practices	74
8.4 Application Log	74
8.5 Troubleshooting Q & A	75
Appendix A Port Counters Overview	82
A.1 Utilization	82
A.1.1 PortXmitData (TxD) and PortVLXmitData[n]	82
A.1.2 PortRcvData (RxD) and PortVLRcvData[n]	82
A.1.3 PortMulticastXmitPkts (MTxP)	82
A.1.4 PortMulticastRcvPkts (MRxP)	82
A.2 Link Integrity	82
A.2.1 Link Quality Indicator (LQI)	83
A.2.2 LocalLinkIntegrityErrors (LLI) Counter	83
A.2.3 PortRcvErrors (RxE) Counter	83
A.2.4 ExcessiveBufferOverflowErrors (EBO) Counter	83
A.2.5 LinkErrorRecovery (LER) Counter	84
A.2.6 LinkDowned (LD) Counter	84
A.2.7 UncorrectableErrors (Unc) Counter	84
A.2.8 FMConfigErrors Counter (FMC)	84
A.3 Congestion	84
A.3.1 CongDiscards (CD) Counter	84
A.3.2 PortRcvFECN (RxF) Counter	84
A.3.3 PortRcvBECN (RxB) Counter	85
A.3.4 PortMarkFECN (MkF) Counter	85
A.3.5 PortXmitTimeCong (TxTC) Counter	85
A.3.6 PortXmitWait (TxW) Counter	85
A.4 SMA Congestion	85
A.4.1 PortVLXmitWait[15] (VLTxW[15]) Counter	85
A.4.2 VLCongDiscards[15] (VLCD[15]) Counter	85
A.4.3 PortVLRcvFECN[15] (VLRxF[15]) Counter	85
A.4.4 PortVLRcvBECN[15] (VLRxB[15]) Counter	85



A.4.5 PortVLXmitTimeCong[15] (VLTxTC[15]) Counter.....	86
A.4.6 PortVLMarkFECN[15] (VLMkF[15]) Counter.....	86
A.5 Bubble.....	86
A.5.1 PortXmitWastedBW (WBW) Counter.....	86
A.5.2 PortXmitWaitData (TxWD) Counter.....	86
A.5.3 PortRcvBubble (RxBb) Counter.....	86
A.6 Security.....	86
A.6.1 PortRcvConstraintErrors (RxCE).....	86
A.6.2 PortXmitConstraintErrors (TxCE).....	87
A.7 Routing.....	87
A.7.1 PortRcvSwitchRelayErrors (RxSR).....	87
A.8 Other.....	87
A.8.1 PortRcvRemotePhysicalErrors (RxRP).....	87
A.8.2 PortXmitDiscards (TxDc).....	87



Figures

1	Fabric Manager GUI.....	13
2	Fabric Topology.....	14
3	Time Window.....	16
4	Metrics Collected in Time Window.....	17
5	Email Settings Error Dialog.....	18
6	Enable Email Notification Selection.....	18
7	SMTP Host Selection.....	19
8	SMTP Sender Selection.....	19
9	SMTP Test Recipient Selection.....	20
10	SMTP Settings Test.....	20
11	Save SMTP Settings.....	21
12	Email Notification Event Selection.....	21
13	Save Events.....	22
14	Email Notification Recipients List.....	22
15	Email Notification Test.....	23
16	Save Email Recipients.....	23
17	Event Table Example.....	26
18	Cards pinned to the Pin Board.....	26
19	Pin Board Examples.....	27
20	Subnet Statistics Example.....	28
21	Subnet Status Example.....	29
22	Health Trend Example.....	29
23	Health Score Example.....	31
24	Subnet Performance Example.....	31
25	Top N Example.....	33
26	Performance Tab Example.....	35
27	Events Example.....	37
28	General Summary Example.....	37
29	Top N Example.....	39
30	Performance (Node) Example.....	40
31	Connectivity Example.....	43
32	Jump to Performance or Topology.....	44
33	Cable Information Jump.....	44
34	Node Properties Example.....	47
35	Port Performance Example.....	51
36	Topology Example.....	57
37	Local Topology Example.....	58
38	Device Set Example.....	59
39	Device Groups Example.....	60
40	Device Node Example.....	60
41	Links Example.....	61
42	Routes Example.....	62
43	Admin Tab.....	63
44	Applications Management Example.....	64
45	Device Group Management Example.....	65
46	Virtual Fabric Management Example.....	66
47	Interactive Console.....	67
48	Interactive Console Examples.....	68
49	Log Viewer.....	69
50	Configuration Login Screen.....	72
51	Troubleshooting Example.....	75



Tables

1	Pin Board Options.....	27
2	Ports Table Data.....	40
3	Connectivity Data.....	44
4	Device Information.....	47
5	Port Properties.....	48
6	Switch Properties.....	48
7	Partition Key Enforcement.....	48
8	Forwarding Information.....	48
9	IP Address.....	49
10	Routing Mode.....	49
11	Adaptive Routing.....	49
12	Receive Counters.....	52
13	Transmit Counters.....	53
14	Other Counters.....	54
15	Typographical Conventions.....	73
16	Link Quality Values and Description.....	83



Preface

This manual is part of the documentation set for the Intel® Omni-Path Fabric (Intel® OP Fabric), which is an end-to-end solution consisting of Intel® Omni-Path Host Fabric Interfaces (HFIs), Intel® Omni-Path switches, and fabric management and development tools.

The Intel® OP Fabric delivers a platform for the next generation of High-Performance Computing (HPC) systems that is designed to cost-effectively meet the scale, density, and reliability requirements of large-scale HPC clusters.

Both the Intel® OP Fabric and standard InfiniBand* are able to send Internet Protocol (IP) traffic over the fabric, or *IPoFabric*. In this document, however, it is referred to as *IP over IB* or *IPoIB*. From a software point of view, IPoFabric and IPoIB behave the same way and, in fact, use the same `ib_ipoib` driver to send IP traffic over the `ib0` and/or `ib1` ports.

Intended Audience

The intended audience for the Intel® Omni-Path (Intel® OP) document set is network administrators and other qualified personnel.

Intel® Omni-Path Documentation Library

Intel® Omni-Path publications are available at the following URLs:

- Intel® Omni-Path Switches Installation, User, and Reference Guides
<http://www.intel.com/omnipath/SwitchPublications>
- Intel® Omni-Path Software Installation, User, and Reference Guides (includes HFI documents)
<http://www.intel.com/omnipath/FabricSoftwarePublications>
- Drivers and Software (including Release Notes)
<http://www.intel.com/omnipath/Downloads>

Use the tasks listed in this table to find the corresponding Intel® Omni-Path document.

Task	Document Title	Description
Key: Shading indicates the URL to use for accessing the particular document.		
• Intel® Omni-Path Switches Installation, User, and Reference Guides:	http://www.intel.com/omnipath/SwitchPublications	
• Intel® Omni-Path Software Installation, User, and Reference Guides (includes HFI documents):	http://www.intel.com/omnipath/FabricSoftwarePublications (no shading)	
• Drivers and Software (including Release Notes):	http://www.intel.com/omnipath/Downloads	
<i>continued...</i>		



Task	Document Title	Description
Using the Intel® OPA documentation set	<i>Intel® Omni-Path Fabric Quick Start Guide</i>	A roadmap to Intel's comprehensive library of publications describing all aspects of the product family. It outlines the most basic steps for getting your Intel® Omni-Path Architecture (Intel® OPA) cluster installed and operational.
Setting up an Intel® OPA cluster	<i>Intel® Omni-Path Fabric Setup Guide</i> (Old title: <i>Intel® Omni-Path Fabric Staging Guide</i>)	Provides a high level overview of the steps required to stage a customer-based installation of the Intel® Omni-Path Fabric. Procedures and key reference documents, such as Intel® Omni-Path user guides and installation guides are provided to clarify the process. Additional commands and BKM's are defined to facilitate the installation process and troubleshooting.
Installing hardware	<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>	Describes the hardware installation and initial configuration tasks for the Intel® Omni-Path Switches 100 Series. This includes: Intel® Omni-Path Edge Switches 100 Series, 24 and 48-port configurable Edge switches, and Intel® Omni-Path Director Class Switches 100 Series.
	<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>	Contains instructions for installing the HFI in an Intel® OPA cluster. A cluster is defined as a collection of nodes, each attached to a fabric through the Intel interconnect. The Intel® HFI utilizes Intel® Omni-Path switches and cabling.
Installing host software Installing HFI firmware Installing switch firmware (externally-managed switches)	<i>Intel® Omni-Path Fabric Software Installation Guide</i>	Describes using a Text-based User Interface (TUI) to guide you through the installation process. You have the option of using command line interface (CLI) commands to perform the installation or install using the Linux* distribution software.
Managing a switch using Chassis Viewer GUI Installing switch firmware (managed switches)	<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	Describes the Intel® Omni-Path Fabric Chassis Viewer graphical user interface (GUI). It provides task-oriented procedures for configuring and managing the Intel® Omni-Path Switch family. Help: GUI online help.
Managing a switch using the CLI Installing switch firmware (managed switches)	<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	Describes the command line interface (CLI) task information for the Intel® Omni-Path Switch family. Help: -help for each CLI.
Managing a fabric using FastFabric	<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i> (Merged with: <i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i>)	Provides instructions for using the set of fabric management tools designed to simplify and optimize common fabric management tasks. The management tools consist of TUI menus and command line interface (CLI) commands. Help: -help and man pages for each CLI. Also, all host CLI commands can be accessed as console help in the Fabric Manager GUI.
Managing a fabric using Fabric Manager	<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>	The Fabric Manager uses a well defined management protocol to communicate with management agents in every Intel® Omni-Path Host Fabric Interface (HFI) and switch. Through these interfaces the Fabric Manager is able to discover, configure, and monitor the fabric.
	<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>	Provides an intuitive, scalable dashboard and set of analysis tools for graphically monitoring fabric status and configuration. It is a user-friendly alternative to traditional command-line tools for day-to-day monitoring of fabric health. Help: Fabric Manager GUI Online Help.
continued...		



Task	Document Title	Description
Configuring and administering Intel® HFI and IPoIB driver Running MPI applications on Intel® OPA	<i>Intel® Omni-Path Fabric Host Software User Guide</i>	Describes how to set up and administer the Host Fabric Interface (HFI) after the software has been installed. The audience for this document includes both cluster administrators and Message-Passing Interface (MPI) application programmers, who have different but overlapping interests in the details of the technology.
Writing and running middleware that uses Intel® OPA	<i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>	Provides a reference for programmers working with the Intel® PSM2 Application Programming Interface (API). The Performance Scaled Messaging 2 API (PSM2 API) is a low-level user-level communications interface.
Optimizing system performance	<i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>	Describes BIOS settings and parameters that have been shown to ensure best performance, or make performance more consistent, on Intel® Omni-Path Architecture. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.
Designing an IP or storage router on Intel® OPA	<i>Intel® Omni-Path IP and Storage Router Design Guide</i>	Describes how to install, configure, and administer an IPoIB router solution (Linux* IP or LNet) for inter-operating between Intel® Omni-Path and a legacy InfiniBand* fabric.
Building a Lustre* Server using Intel® OPA	<i>Building Lustre* Servers with Intel® Omni-Path Architecture Application Note</i>	Describes the steps to build and test a Lustre* system (MGS, MDT, MDS, OSS, OST, client) from the HPDD master branch on a x86_64, RHEL*/CentOS* 7.1 machine.
Building Containers for Intel® OPA fabrics	<i>Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note</i>	Provides basic information for building and running Docker* and Singularity* containers on Linux*-based computer platforms that incorporate Intel® Omni-Path networking technology.
Writing management applications that interface with Intel® OPA	<i>Intel® Omni-Path Management API Programmer's Guide</i>	Contains a reference for programmers working with the Intel® Omni-Path Architecture Management (Intel OPAMGT) Application Programming Interface (API). The Intel OPAMGT API is a C-API permitting in-band and out-of-band queries of the FM's Subnet Administrator and Performance Administrator.
Learning about new release features, open issues, and resolved issues for a particular release	<i>Intel® Omni-Path Fabric Software Release Notes</i>	
	<i>Intel® Omni-Path Fabric Manager GUI Release Notes</i>	
	<i>Intel® Omni-Path Fabric Switches Release Notes</i> (includes managed and externally-managed switches)	

Cluster Configurator for Intel® Omni-Path Fabric

The Cluster Configurator for Intel® Omni-Path Fabric is available at: <http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-configurator.html>.

This tool generates sample cluster configurations based on key cluster attributes, including a side-by-side comparison of up to four cluster configurations. The tool also generates parts lists and cluster diagrams.

Documentation Conventions

The following conventions are standard for Intel® Omni-Path documentation:

- **Note:** provides additional information.
- **Caution:** indicates the presence of a hazard that has the potential of causing damage to data or equipment.



- **Warning:** indicates the presence of a hazard that has the potential of causing personal injury.
- Text in [blue](#) font indicates a hyperlink (jump) to a figure, table, or section in this guide. Links to websites are also shown in blue. For example:
See [License Agreements](#) on page 12 for more information.
For more information, visit www.intel.com.
- Text in **bold** font indicates user interface elements such as menu items, buttons, check boxes, key names, key strokes, or column headings. For example:
Click the **Start** button, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.
Press **CTRL+P** and then press the **UP ARROW** key.
- Text in `Courier` font indicates a file name, directory path, or command line text. For example:
Enter the following command: `sh ./install.bin`
- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:
Refer to *Intel® Omni-Path Fabric Software Installation Guide* for details.
In this document, the term *chassis* refers to a managed switch.

Procedures and information may be marked with one of the following qualifications:

- **(Linux)** – Tasks are only applicable when Linux* is being used.
- **(Host)** – Tasks are only applicable when Intel® Omni-Path Fabric Host Software or Intel® Omni-Path Fabric Suite is being used on the hosts.
- **(Switch)** – Tasks are applicable only when Intel® Omni-Path Switches or Chassis are being used.
- Tasks that are generally applicable to all environments are not marked.

License Agreements

This software is provided under one or more license agreements. Please refer to the license agreement(s) provided with the software for specific detail. Do not install or use the software until you have carefully read and agree to the terms and conditions of the license agreement(s). By loading or using the software, you agree to the terms of the license agreement(s). If you do not wish to so agree, do not install or use the software.

Technical Support

Technical support for Intel® Omni-Path products is available 24 hours a day, 365 days a year. Please contact Intel Customer Support or visit <http://www.intel.com/omnipath/support> for additional detail.

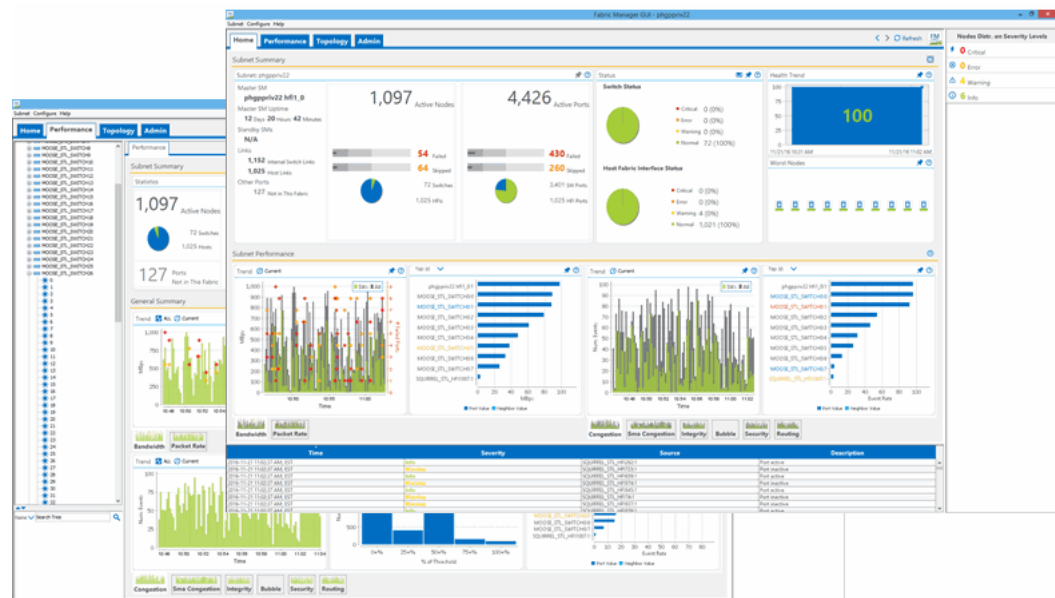


1.0 Introduction

Intel® Omni-Path Fabric Suite Fabric Manager GUI provides an intuitive, scalable dashboard and set of analysis tools for graphically monitoring fabric status and configuration. It is a user-friendly alternative to traditional command-line tools for day-to-day monitoring of fabric health.

For details about the other documents for the Intel® Omni-Path product line, refer to [Intel® Omni-Path Documentation Library](#) on page 9 of this document.

Figure 1. Fabric Manager GUI



Fabric Manager GUI can be run on a Linux* or Windows* desktop/laptop system with TCP/IP connectivity to the Intel® Omni-Path Fabric Suite Fabric Manager. Network connectivity to the Fabric Executive (FE) component of the Fabric Manager allows continuous, remote, “out of band” monitoring. If configured with alternate FE connection information, the Fabric Manager GUI automatically fails over to an alternative FE if the connection to the original is lost.

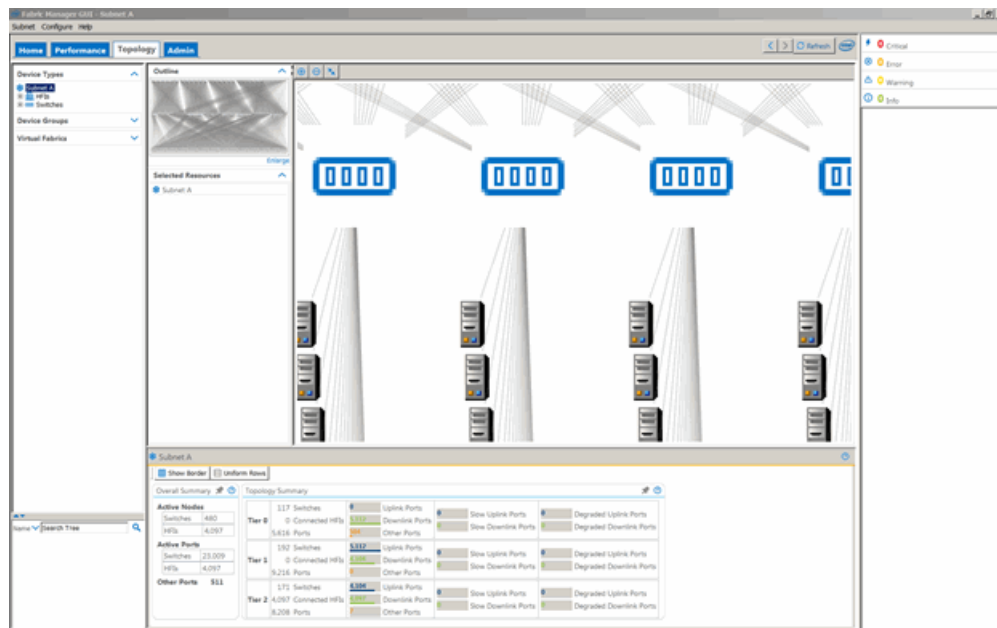
Fabric Manager GUI feature highlights include:

- The ability to monitor multiple fabric “subnets” simultaneously
- A dashboard summary of the state of the connected fabric
- An interactive way to visualize the fabric topology (see the following figure)
- A Pin Board that allows a user-selected set of GUI indicators to remain visible at all times



- Reliable “out-of-band” connectivity to the fabric over both secured and non-secured links
- Email and other alerts if user-selected events occur within the fabric
- Connectivity to the Fabric Manager’s console for access to the rich set of Intel® Omni-Path Fabric Suite FastFabric command-line tools
- Connectivity to the Fabric Manager’s log file for troubleshooting purposes
- The ability to modify the configuration of the Fabric Manager

Figure 2. Fabric Topology






2.0 Intel® Omni-Path Fabric Suite Fabric Manager GUI Setup Wizard

Setup Wizard guides a user to set up the Fabric Manager GUI , including the areas described below.

- Hosts
- Events
- Preferences

Hosts

To add a new subnet, click the  button on the Hosts tab and provide the following attributes:

- Name – unique name of the subnet
- IP address or Host name – address used to connect with the FM, specifically the address used to connect to the FE
- Port number - the port number used to connect

To add a subnet, user name is not required. Even for the secure connection (when the **Secure** check box is checked), user name is not necessary.

For a secure connection, check the **Secure** check box and enter **Key Store** and **Trust Store** information. See the "SSL Key Creation for Fabric Manager GUI" appendix in the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*.

The **Auto Connect** checkbox must be selected to make the connection automatic.

Events

In the course of exploring a fabric, the Fabric Manager may send an event message to the Fabric Manager GUI to notify the user about an incident occurring on the network. The Events tab lists the classification and type of each event, and allows the user to select the severity level and action to take when an event is received.

Events are defined with the following attributes:

- Event Class – Subnet Event, Miscellaneous Event, etc.
- Event Type – Lost Connection with FE, Inactive Port, etc.
- Severity – Informational, Warning, Error, Critical
- Action – Set notification with an email, display, or both

All events have predefined default values, so you need to change only those events of interest.



User Preferences

On the Preferences tab, you define the following attributes:

Refresh Rate – the rate at which the UI is refreshed with the latest data from the Fabric Manager. This rate must not be less than the Fabric Manager sweep rate.

Time Window – the duration of the sliding time window over which the number of active nodes and ports:

- Reported a Critical, Error, or Warning event type
- Was classified as a Worst Node based on its performance

In the following figure, a time window of 1 hour is defined. As time transpires the data collected prior to the time window is removed as new information is reported.

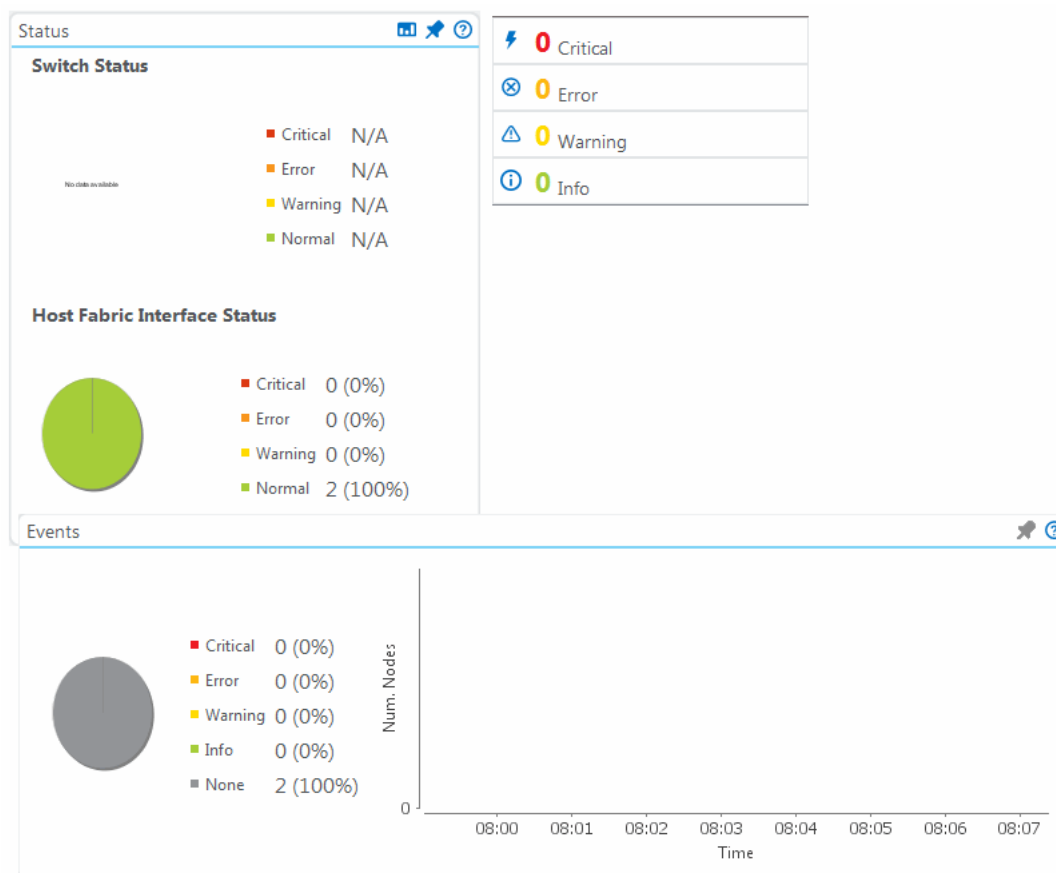
Figure 3. Time Window



Indicators for these values can be found on the Home and Performance pages as shown in the highlighted panels below.



Figure 4. Metrics Collected in Time Window



Note: This value must be larger than $10 * \text{FM sweep rate}$.

Worst Nodes – number of the worst nodes to display.

Email recipients – used to list recipients of email notifications for the subnet. See [Setting up Email Notification](#) on page 17 for more details.

2.1 Setting up Email Notification

The email notification feature allows a specified list of addresses to receive an email notification if a user-selected event occurs in the Fabric Manager GUI or in the fabric itself. Both the list of recipients and the list of events are specified on a per-subnet basis using the Configuration Setup Wizard: Configure menu > **Setup Wizard** (refer to [Intel® Omni-Path Fabric Suite Fabric Manager GUI Setup Wizard](#) on page 15). A single SMTP server is used to send messages for all defined subnets and is specified using the Email Settings dialog (Configure menu > **Email setup**).

Server Invalid/No Server Configured

The notification in the following figure is displayed any time the Fabric Manager GUI cannot connect to the specified SMTP server. This can be due to the Fabric Manager GUI host being offline (for example, no network connection), invalid SMTP settings, or

SMTP settings left blank (with email notification enabled). If the dialog appears in response to an SMTP server settings test, you can dismiss the dialog and the SMTP settings can still be saved.


Figure 5. Email Settings Error Dialog



SMTP Configuration

A single SMTP server is used to send messages for all defined subnets and is specified using the Email Settings dialog (Configure menu > **Email setup**). To begin SMTP server configuration, click the **Enable Email Notifications** check box in the Email Settings dialog shown in the following figure.

Figure 6. Enable Email Notification Selection



You may enable/disable email notifications at any time but the SMTP settings may be edited only when email notifications have been enabled.

Enter the SMTP server address (for example, "mail.mydomain.com") in the **SMTP Host** field as shown in the following figure.



Figure 7. SMTP Host Selection

The 'Email Settings' dialog box is shown. It has a title bar with a close button. Inside, there is a checkbox labeled 'Enable Email Notifications' which is checked. Below it, the 'SMTP Host' field is highlighted with a red oval. The 'Port #' field contains the value '25'. The 'Sender' field contains the placeholder text 'Please enter valid email address'. Below these fields, there is a section titled 'Please provide recipient's email address to test SMTP settings.' which contains an empty text box and a 'Test' button. At the bottom of the dialog, there are three buttons: 'Reset', 'Cancel', and 'OK'.

Enter a valid email address in the "Sender" field as shown in the following figure. This is the "from" address that will appear in all email notifications.

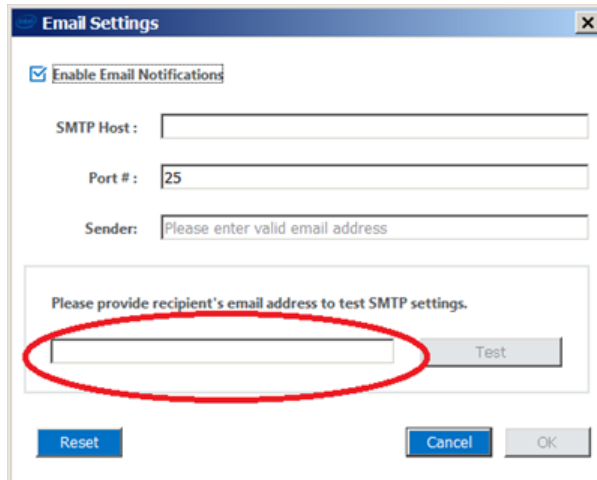
Figure 8. SMTP Sender Selection

The 'Email Settings' dialog box is shown. It has a title bar with a close button. Inside, there is a checkbox labeled 'Enable Email Notifications' which is checked. Below it, the 'SMTP Host' field is empty. The 'Port #' field contains the value '25'. The 'Sender' field is highlighted with a red oval and contains the placeholder text 'Please enter valid email address'. Below these fields, there is a section titled 'Please provide recipient's email address to test SMTP settings.' which contains an empty text box and a 'Test' button. At the bottom of the dialog, there are three buttons: 'Reset', 'Cancel', and 'OK'.

SMTP Server Test

If the Fabric Manager GUI host has a working network connection, you can test the SMTP server settings if you choose to do so. Enter a valid email address for the recipient of the test notification message as shown in the following figure.

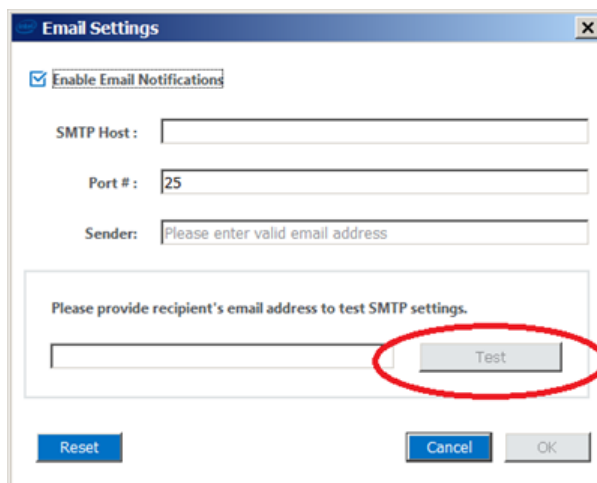
Figure 9. SMTP Test Recipient Selection



The screenshot shows the 'Email Settings' dialog box. At the top, there is a checkbox labeled 'Enable Email Notifications' which is checked. Below this are three input fields: 'SMTP Host:', 'Port #:' (with '25' entered), and 'Sender:' (with the placeholder text 'Please enter valid email address'). Below these fields is a section titled 'Please provide recipient's email address to test SMTP settings.' which contains an empty text input field and a 'Test' button. A red circle is drawn around the 'Test' button. At the bottom of the dialog are three buttons: 'Reset', 'Cancel', and 'OK'.

The **Test** button (shown in the following figure), can be used to send a test notification message to the specified recipient. After all settings in the dialog have been entered, the **Test** button becomes active. Click the button to send the message. If the Fabric Manager GUI is unable to connect to the SMTP server for any reason, the error notification dialog shown in [Figure 5](#) on page 18 will be displayed. Note that you can still save the entered settings, if desired.

Figure 10. SMTP Settings Test



This screenshot is identical to the one in Figure 9, showing the 'Email Settings' dialog box with the 'Test' button highlighted by a red circle.

If the test is successful, the recipient receives the test message at the address specified. Note that the Fabric Manager GUI can provide feedback on server connection status only. It is possible for the Fabric Manager GUI to connect to a valid SMTP server without the test message being received. In this case, verify that the recipient address is correct and that any spam filter is set to pass the message.

Save SMTP Settings

To save SMTP settings, click **OK** as shown in the following figure.


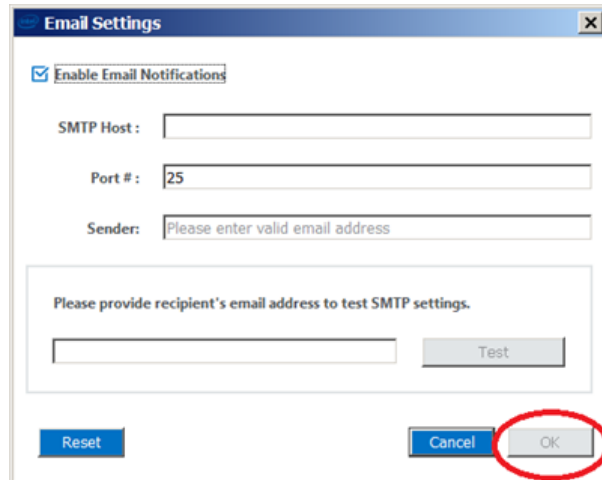


Figure 11. Save SMTP Settings



The 'Email Settings' dialog box contains the following elements:

- ☒ **Enable Email Notifications**
- SMTP Host :** [Text input field]
- Port # :** [Text input field with value 25]
- Sender:** [Text input field with placeholder 'Please enter valid email address']
- Please provide recipient's email address to test SMTP settings.** [Text input field]
- Test** button
- Reset** button
- Cancel** button
- OK** button (circled in red)

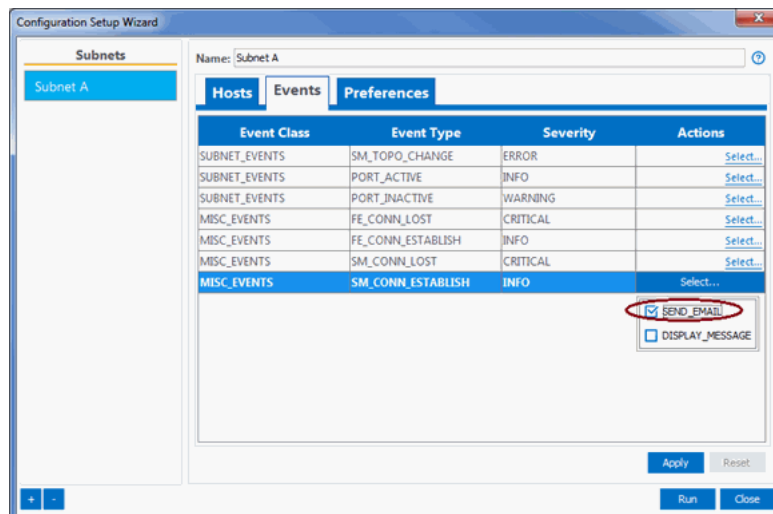
To dismiss the Email Settings dialog without saving, click **Cancel**. To revert to the last saved settings, click **Reset**.

2.2 Email Notification of Events

Select Events

To send an email notification when a particular event occurs, Click **Select...** in the Actions column in the row of the desired event as shown in the following figure. Then click the **SEND_EMAIL** check box. An envelope icon will appear in the Actions column for all selected events.

Figure 12. Email Notification Event Selection



The 'Configuration Setup Wizard' dialog box, 'Events' tab, shows a table of event classes and types. The 'SEND_EMAIL' checkbox is checked for the selected event.

Event Class	Event Type	Severity	Actions
SUBNET_EVENTS	SM_TOPO_CHANGE	ERROR	Select...
SUBNET_EVENTS	PORT_ACTIVE	INFO	Select...
SUBNET_EVENTS	PORT_INACTIVE	WARNING	Select...
MISC_EVENTS	FE_CONN_LOST	CRITICAL	Select...
MISC_EVENTS	FE_CONN_ESTABLISH	INFO	Select...
MISC_EVENTS	SM_CONN_LOST	CRITICAL	Select...
MISC_EVENTS	SM_CONN_ESTABLISH	INFO	Select...

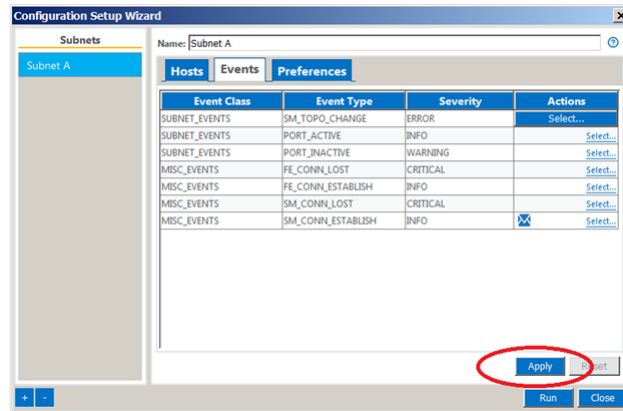
Below the table, the **SEND_EMAIL** checkbox is checked, and the **DISPLAY_MESSAGE** checkbox is unchecked.

Buttons at the bottom: **Apply**, **Reset**, **Run**, **Close**.

Save Events

To save the selected events, click the **Apply** button as shown in the following figure.

Figure 13. Save Events

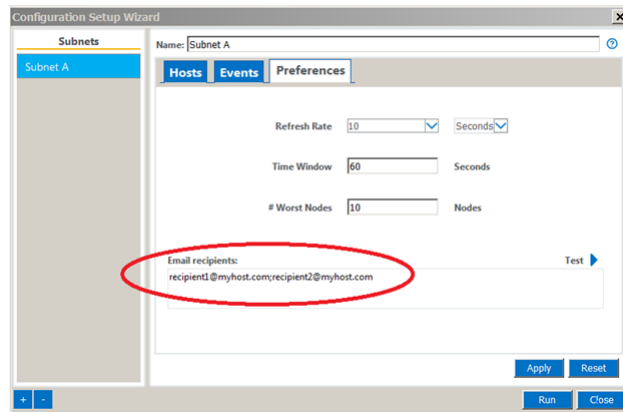


Email Notification Recipients

Email Recipients List

To specify a list of email notification recipients for the subnet, enter one, or more valid email addresses in the **Email recipients** field in the Preferences tab as shown in the following figure. Entries must be separated by a semi-colon.

Figure 14. Email Notification Recipients List




Note that addresses are validated as they are entered. Addresses must be of the form user@host.ext where "ext" is a three character extension (e.g., "com"). A red border surrounds the field until a valid entry is completed.

Email Test

If desired, a test notification message can be sent to all addresses in the recipients list. The Fabric Manager GUI host must have a working network connection and valid SMTP server settings must be configured prior to the test (refer to [Setting up Email Notification](#) on page 17). Click the **Test** button as shown in the following figure.



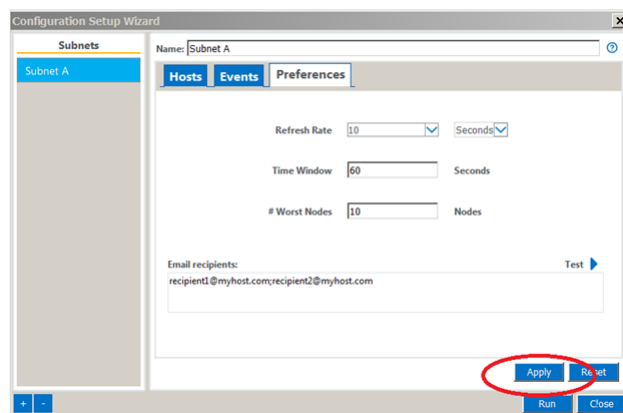
Figure 15. Email Notification Test



Save Recipients

To save the list of email recipients, click the **Apply** button as shown in the following figure.

Figure 16. Save Email Recipients

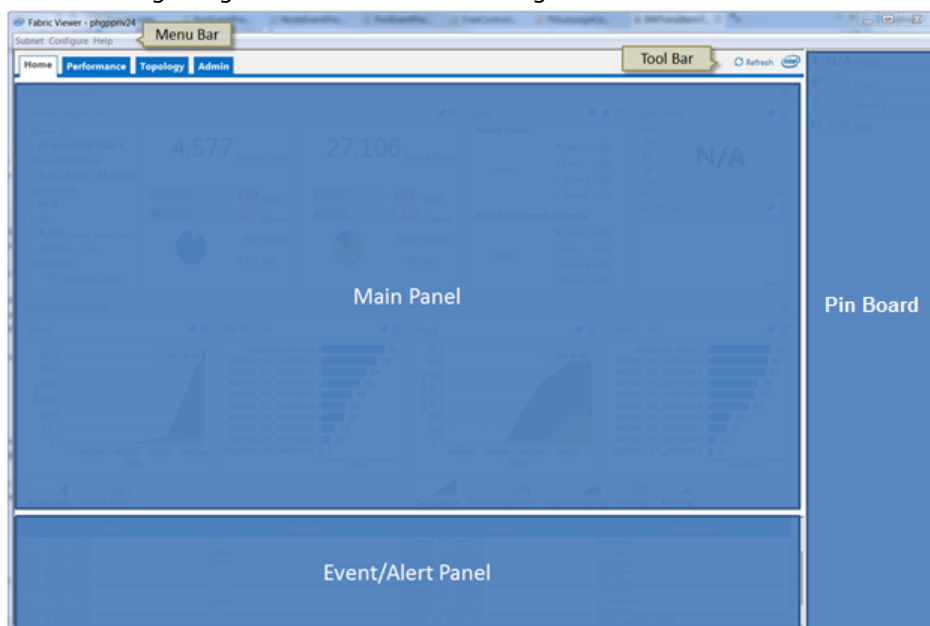


3.0 Fabric Manager GUI Window Navigation

The Fabric Manager GUI window is comprised of the following areas:

- [Menu Bar](#) on page 24
- [Tool Bar](#) on page 25
- [Main Panel](#) on page 25
- [Event Table](#) on page 25
- [Pin Board](#) on page 26

The following image shows the Fabric Manager GUI window with call-outs of the areas.



3.1 Menu Bar

The Menu Bar supplies window and application specific menus that provide access to the following functions:

Subnet Menu

- [Connect to a Subnet](#) on page 25
- Close

Configure Menu

- [Intel® Omni-Path Fabric Suite Fabric Manager GUI Setup Wizard](#) on page 15
- Logging - In the Logging section, you set up the output destination, information level, and output format required for error logging.



- Email Setup - Set up SMTP server used in email notification.

Help Menu

- Online Help
- About

3.1.1 Connect to a Subnet

Subnets may be configured to connect to the fabric either automatically or manually. To connect a subnet automatically upon starting the Fabric Manager GUI, see [Hosts](#) on page 15.



To manually connect a subnet to the fabric:

From the menu bar, click **Subnet, Connect To**, and then the subnet.

Note: If a connection error is received, verify that the host name and port number are correct, and verify that the Fabric Executive is active.

3.2 Tool Bar

The tool bar contains the following buttons.

	Refresh – Refreshes Fabric Manager GUI immediately with the latest data from the Fabric Manager or indicates that a data refresh is taking place.
	Navigation - Navigate to previous or next page and subpage displayed in Fabric Manager GUI.

3.3 Main Panel

The Main Panel contains the content for one of the following selected tabs:

- [Home Tab](#) on page 28
- [Performance Tab](#) on page 35
- [Topology](#) on page 57
- [Admin Tab](#) on page 63

3.4 Event Table

The Event Table displays the latest notices received from the Fabric Manager. By default this table is sorted by notice receive time, the time the Fabric Manager GUI receives a notice, in descending order and therefore shows the latest notice as the first row in the table. This table also shows each notice's severity level with colored text, the source, and the description of the notice. This table is visible across all pages on the **Home** tab, and can be turned on/off by clicking the event summary panel on the pin board.



Figure 17. Event Table Example

Time	Severity	Source	Description
2014-11-14 11:47:21 EST	Critical	SQUIRREL_STL_HF0421:1	Port inactive
2014-11-14 11:47:21 EST	Info	SQUIRREL_STL_HF0448:1	Port active
2014-11-14 11:47:21 EST	Critical	SQUIRREL_STL_HF0448:1	Port inactive
2014-11-14 11:47:21 EST	Error	MOOSE_STL_SW07C980	Subnet topology change
2014-11-14 11:47:21 EST	Critical	SQUIRREL_STL_HF0279:1	Port inactive
2014-11-14 11:47:21 EST	Info	SQUIRREL_STL_HF0282:1	Port active
2014-11-14 11:47:21 EST	Info	SQUIRREL_STL_HF0304:1	Port active
2014-11-14 11:47:21 EST	Info	SQUIRREL_STL_HF0333:1	Port active
2014-11-14 11:47:21 EST	Info	SQUIRREL_STL_HF0392:1	Port active

3.5 Pin Board




The Fabric Manager GUI provides a pin board at the right side of the main window (see the following figure) to monitor cards of interest. Click the pin button  on a card's title bar to make it instantly visible on the pin board, which allows you to monitor performance metrics while exploring other parts of the system. For example, the congestion trend card can be pinned to the pin board and it remains visible even when looking at the performance of a specific port. To remove cards from the pin board, click the  button.

Figure 18. Cards pinned to the Pin Board



The Event Summary is the default pin card that shows the summary of events in a defined time window; it cannot be unpinned. Clicking the Event Summary card toggles the event panel at the bottom of the main window. Below it are user pin cards that are created by clicking the pin button  on a card's title bar. This board holds a maximum of eight pin cards. Each pin card is resizable by dragging the bottom edge. The following table summarizes the available operations on a pin card.

**Table 1. Pin Board Options**





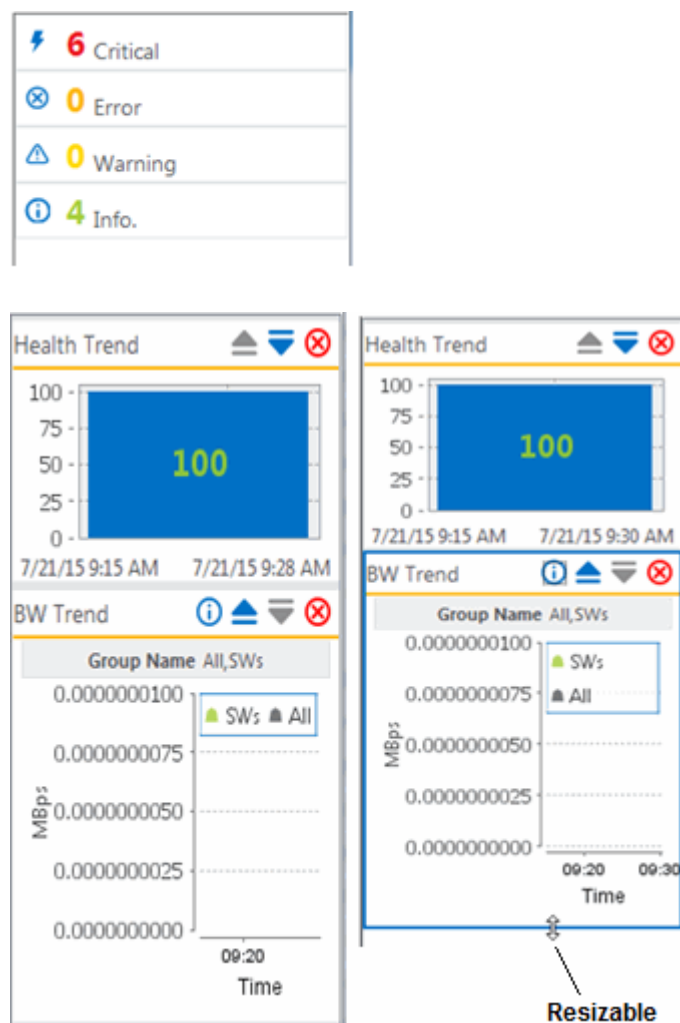
	displays information about this pin card, such as data type, data source, etc.
	moves this pin card up
	moves this pin card down
	unpins this pin card

Figure 19. Pin Board Examples

The previous figure shows the default pin card displaying the summary of events in a defined time window by severity levels. Clicking on this card toggles the event table between visible and invisible states. By default, the Event Table is visible when the Home Page is selected and invisible when other pages are selected.

4.0 Home Tab

The Home tab provides the "big picture" of a fabric using two sections: Subnet Summary and Subnet Performance. See the topics below for details on each section of the Home tab.

- [Subnet Summary \(Home\)](#) on page 28
- [Subnet Performance](#) on page 31

4.1 Subnet Summary (Home)

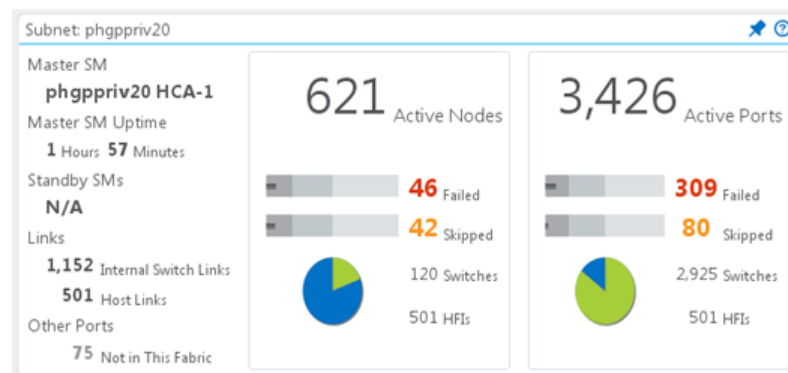
This section of the Home tab summarizes what happened in the past. It includes the following panels:

- [Subnet Statistics](#) on page 28
- [Subnet Status](#) on page 29
- [Health Trend](#) on page 29
- [Worst Nodes](#) on page 31

4.1.1 Subnet Statistics

Subnet Statistics lists basic information about a subnet in the first column: Master SM name and its up time; Standby SMs' names; Summary of Links (number of internal switch links and host links); and Other ports (ports that are not in the subnet). The second and third columns show the statistical summary of nodes and ports separately. The top of each column displays the number of active nodes/ports, followed by the numbers of no response and skipped nodes/ports, and then the distribution of device types (specifically Switch and HFI).

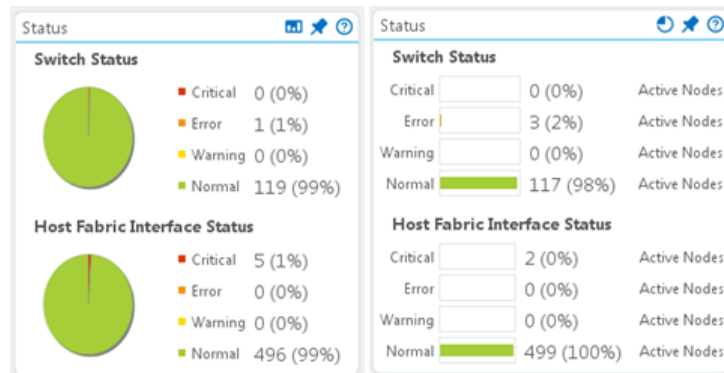
Figure 20. Subnet Statistics Example



4.1.2 Subnet Status

Subnet Status shows the current status of switches and HFIs. It tells a user that in the past time window (defined in setup wizard), how many switches/HFIs are under critical, error, warning, or normal status and are measured based on notices from the Fabric Manager. If one node has multiple notices within a time window, it is counted by the notice with the highest severity level. The map between notice and severity is customizable with the setup wizard. Choose either pie chart or bar chart with the icon on the right side of the title bar.


Figure 21. Subnet Status Example



4.1.3 Health Trend

Health Trend displays the health history of a subnet in an area chart.

Figure 22. Health Trend Example




The card shown in the previous figure displays the health history of a subnet in an area chart. On the chart is the current health score which is calculated with the following equation:

$$Health\ Score = \sum_j (S_j \times W_j) \times 100$$

where:

- S_j = the score for attribute j
- W_j = the weight for attribute j



The score for an attribute can be calculated using the following formula:

$$S_j = \frac{V_j}{\sum_j (B_j \times W_j)}$$

where:

- V_j = the current value for attribute j
- B_j = the current baseline for attribute j


Currently, the attributes, their source and their weight in the formula are as follows:

Attribute	Source	Weight
Number of switches	ImageInfo	Average number of active ports per switch + 1
Number of HFIs	ImageInfo, FabricInfo	Average number of active ports per HFI + 1
Number of Inter Switch Links	FabricInfo	3
Number of HFI links	FabricInfo	2
Number of Active Ports	ImageInfo	1
Number of Non-degraded Inter Switch Links	FabricInfo	3
Number of Non-degraded HFI links	FabricInfo	2

Weights represent the relative importance of one attribute against the others. For example, consider the small fabric in the following image: two 48-port switches connected by four ISLs and serving eight nodes. In this configuration, a switch down will have 5 times more impact in the health score than a HFI link down (switch weight = 4 node-ports + 4 ISLs + port 0 + 1 = 10 against weight HFI link = 2). However, the impact is even greater if you consider that a switch down may also affect other attributes. In this example, the health score goes down to 53% when you bring down one of the switches but only goes down to 93% when you bring down a HFI link. The end result is that even small changes in the fabric configuration are not diluted when the size of the fabric is big: in simulations with a fabric of over 8K nodes, a port being brought down is reflected with the non-perfect score of 99%.



Figure 23. Health Score Example



You can see the current values of each attribute and their baselines by hovering the mouse over the score; a tooltip appears to help you determine the source of a non-perfect score. The baseline values for each attribute are taken during initialization of the application and whenever the number of switches or HFIs increases in the fabric.

4.1.4 Worst Nodes

Worst Nodes lists the nodes with low health scores in a subnet as measured by a node's status. Normal, warning, error, and critical status levels have the scores of 100, 80, 30, and 10 respectively. Click on a worst node to display more information about the node and view performance, topology, and other data.

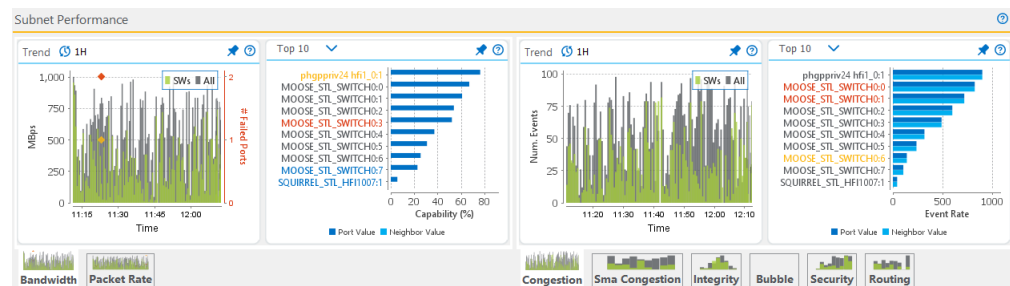
4.2 Subnet Performance

The Subnet Performance displays a statistical overview of the fabric to provide insight into the hot spots. It contains two groups of performance: Utilization and PM Counter Categories.

- [Utilization Group](#) on page 32
- [PM Counter Categories Group](#) on page 33

Refer to *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*, in the section titled "Counter Classification" for more information on how Performance and "error" groups work, how thresholds are set, weights settings, adding PM groups, etc).

Figure 24. Subnet Performance Example





4.2.1 Utilization Group

The utilization group provides visual representation of a fabric's performance in terms of bandwidth and packet rate. Each metric is summarized by the following:

Trend Chart	The historical values measured by all devices and all switches in a fabric, showing the trend and how switches contribute to the metric. The value represents the aggregated metric on all data traffic types: internal, send and transmit.
Top N Chart	The CURRENT top N worst ports measured by the selected metric. The parameter N can be changed with the Setup Wizard, or in this chart.
Histogram Chart	The CURRENT value distribution against ports, i.e., how many ports fall in a value range.

The combination of these charts provides mixed insights about the history, the present, and the global and local performance over the ports. On the bottom of the charts panel are the trend thumbnails of available metrics.

To bring the metric to the front, click on a thumbnail . To switch between the Top *N* and Histogram charts, use the drop-down menu on the title bar.

4.2.1.1 Trend (Performance)

The Trend chart shows the historical values measured by all devices and all switches in a fabric. The value represents the aggregated metric on all data traffic types: internal, external, send, and transmit (received). See [General Summary \(Performance\)](#) on page 37 for information about the different traffic types.

The second axis at the right side of the chart represents the number of no response ports within the fabric. If an issue happens when PA collects data, a mark at the top of the trend chart indicates that the data may be inaccurate. Move your mouse over the mark to display detailed information about the issue type and number of no response ports.

4.2.1.2 Top *N* Chart

The Top *N* chart shows the *current* top N worst ports measured by the selected metric. The parameter can be changed with the Setup Wizard, or in this chart.

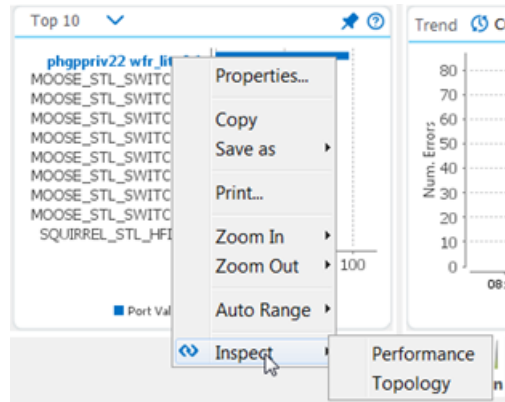
The port labels on the left side are color coded with issue states. Move your mouse over a label to show a description of the state, indicated by the following colors:

- Dark Gray – no issue
- Blue – PMA ignore
- Yellow – PMA no response
- Red – Topology no response

Right-click a port to open a context menu to drill down to view this port's performance or topology.



Figure 25. Top N Example



4.2.1.3 Histogram

The Histogram chart shows the *current* value distribution against ports, i.e., how many ports fall in a value range.

4.2.2 PM Counter Categories Group

The PM Counter Categories Group is similar to the Utilization Group except the following metrics are applied:

- Congestion
- SMA Congestion
- Integrity
- Bubble
- Security
- Routing

Congestion

These counters reflect possible errors that indicate traffic congestion in the fabric.

When congestion or a packet that has seen congestion is detected, one of these counters will be incremented and then depending on the issue reported, the packet will just have to wait or in an extreme case, it may be dropped.

SMA Congestion

These counters reflect congestion in the fabric specific to communication between the Subnet Manager and Subnet Manager Agents using the management VL (VL 15).

The category is calculated exactly as the Congestion category using the same weights and the correct VL15 utilization counters.



Integrity

These counters reflect errors in the Physical (PHY) and Link Layers, as well as errors in Firmware. The typical cause is a hardware problem such as a poor connection, marginal cable, incorrect length/model cable for signal rate, or damaged/broken hardware (for example, bad connectors).

When a bad packet is detected, one of these counters are incremented and the Link Layer discards the packet.

During the link training sequence, assorted errors may be observed. This is a normal part of the link training and clock synchronization process. Hence, errors observed as part of rebooting nodes or moving cables should not be considered a problem.

Bubble

These counters occur when an unexpected idle *flit* is transmitted or received. the term *flit* refers to a Flow Control Digit or Flit, the smallest unit of information on which flow control may be performed. Intel® Omni-Path Fabric packets are divided into flits of 64 bits for transmission across a link. The flit excludes any headers; the 64 bits is the payload size.

The transmit port will send idle flits until it can continue sending the rest of the packet. The category is calculated as follows:

1. The maximum value between the sum of the XmitWastedBW and XmitWaitData or the neighbor's PortRcvBubble.
2. Divide the previous value by the port's utilization to provide context.

Security

These counters reflect possible security problems in the fabric.

Security problems occur if a PKey or SLID violation occurs at the port during the ingress or egress of a packet.

Routing

These counters reflect possible routing issues. When a routing issue occurs, the offending packet is dropped.


A typical cause of this error is the routing to a wrong egress port or an improper Service Channel (SC) mapping.



5.0 Performance Tab

The Performance tab displays the detailed information about a subnet's performance. On the left of the page is a resource tree panel that organizes subnet devices by different categories which allows you to select resources. The main panel shows the performance of a selected resource and its content changes based on the type of resource selected.

Figure 26. Performance Tab Example



See the topics below for details on each section of the Performance tab.

- [Resource Tree](#) on page 35
- [Tree Search](#) on page 36
- [Summarized Performance](#) on page 36
- [Node Performance](#) on page 39
- [Port Performance](#) on page 50

5.1 Resource Tree

The Resource Tree displays and organizes resources of the following types:

- **Device Type:** Organize resources by device types, Switch and HFI
- **PmPortGroups:** Organize resources by PM Port Groups
- **Virtual Fabric:** Organize resources by defined virtual fabric



Although the resources are organized differently in the Resource Tree, there are only three types of resources, as described below. Click a resource to display the performance information in the right panel.

- **Device Set:** A set of device nodes, such as a subnet, a device group, a virtual fabric, or a type of device
- **Device Node:** A fabric node that can be either a switch or HFI
- **Device Port:** A port in a fabric node

5.2 Tree Search

Tree Search enables you to search a node by name, node LID, or node GUID. The search results are displayed in a tree and selecting a node from the results tree populates the node from the Resource Tree.

5.3 Summarized Performance

Summarized performance shows aggregated performance on a set of device nodes. Similar to the Home tab, it includes two sections: Subnet Summary and General Performance.

[Subnet Summary \(Home\)](#) on page 28

[General Summary \(Performance\)](#) on page 37

5.3.1 Subnet Summary (Performance)

The Subnet Summary section displays the following:

[Statistics \(Performance Subnet Summary\)](#) on page 36

[Events \(Performance Subnet Summary\)](#) on page 36

5.3.1.1 Statistics (Performance Subnet Summary)


The Statistics section displays the basic information about a selected resource: the resource statistics summary, i.e., total number of active nodes/ports, their distribution against device types, and the number of ports not in the subnet.

5.3.1.2 Events (Performance Subnet Summary)

The Event section displays past performance statistics based on recent events. The pie chart on the left summarizes the events distribution in a defined time window. This chart is very similar to the Status chart on Home tab, where nodes with informational or no events are categorized as normal nodes. Nodes without events are represented in gray. The bar chart on the right shows how the events distribute at each time point in the past.



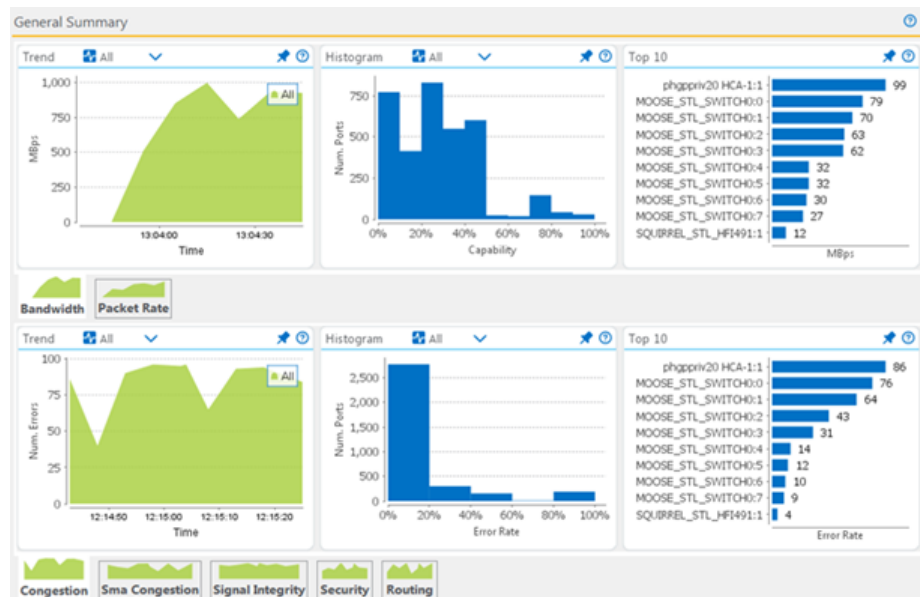
Figure 27. Events Example



5.3.2 General Summary (Performance)

This section is very similar to the Subnet Performance section on the Home tab. It organizes charts line-by-line to allow a user to look at the details of the charts. It also supports different data traffic types via the drop-down menu on the title bar.

Figure 28. General Summary Example



The supported traffic types include:

- Internal - internal data traffic within a device set
- Send - data sent from a device set
- Transmit - data received from the outside
- External - the sum of send and transmit, i.e., the total data traffic with the outside
- All - the sum of internal, send, and transmit

The General Summary contains the following sections:

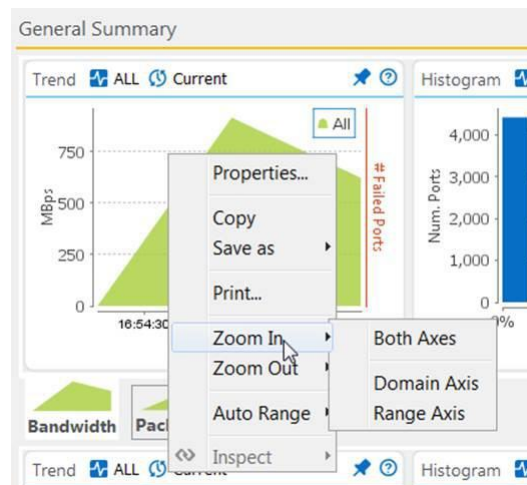
- [Trend \(Performance\)](#) on page 32
- [Top N Chart](#) on page 32
- [Histogram](#) on page 33

The combination of the [Trend \(Performance\)](#) on page 32, [Top N Chart](#) on page 32, and [Histogram](#) on page 33 charts provide mixed insights about the history, the present, and the global and local performance over the ports.

On the bottom of the charts panel are the trend thumbnails of available metrics. Click on a thumbnail to bring the metric to the front.

To zoom an area of the chart, choose one of these methods:

- Using your mouse, draw a rectangle on the chart where you want to zoom.
- OR
- Right-click over the rectangle and select **Zoom In** or **Zoom Out**, then choose **Both Axes**, **Domain Axis**, or **Range Axis**.



To return to the original scale after zooming in/out, click and drag the mouse upward.

5.3.2.1 Trend (Performance)

The Trend chart shows the historical values measured by all devices and all switches in a fabric. The value represents the aggregated metric on all data traffic types: internal, external, send, and transmit (received). See [General Summary \(Performance\)](#) on page 37 for information about the different traffic types.

The second axis at the right side of the chart represents the number of no response ports within the fabric. If an issue happens when PA collects data, a mark at the top of the trend chart indicates that the data may be inaccurate. Move your mouse over the mark to display detailed information about the issue type and number of no response ports.

5.3.2.2 Top N Chart

The Top *N* chart shows the *current* top N worst ports measured by the selected metric. The parameter can be changed with the Setup Wizard, or in this chart.

The port labels on the left side are color coded with issue states. Move your mouse over a label to show a description of the state, indicated by the following colors:


- Dark Gray – no issue



- Blue – PMA ignore
- Yellow – PMA no response
- Red – Topology no response

Right-click a port to open a context menu to drill down to view this port's performance or topology.

Figure 29. Top N Example



5.3.2.3 Histogram

The Histogram chart shows the *current* value distribution against ports, i.e., how many ports fall in a value range.

5.4 Node Performance

Node performance of a device node is displayed in three sections:

- [Performance \(Node\)](#) on page 39
- [Connectivity \(Node\)](#) on page 43
- [Node Properties](#) on page 46

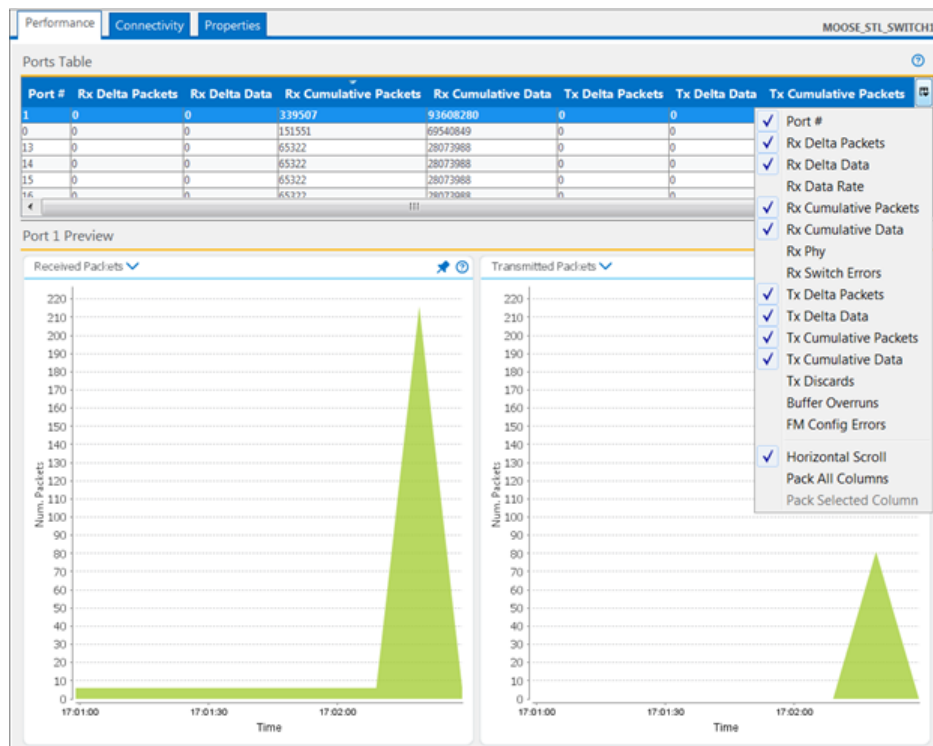
5.4.1 Performance (Node)

The Performance section includes the following:

- [Ports Table](#) on page 40
- [Received Data/Packets Rate](#) on page 42
- [Transmitted Data/Packets Rate](#) on page 42



Figure 30. Performance (Node) Example



5.4.1.1 Ports Table


The Ports Table lists each port's performance data. A preview panel displays the historical traffic data of a selected port, such as received packets/data and transmitted packets/data. Data supporting the table and charts are the port counters belonging to the selected node. Use the  button on the top-right corner of the ports table to select which port counters fields to display.

Table 2. Ports Table Data

Option	Description
Port #	Port number.
LinkQualityIndicator	This is a status indicator, similar to the signal strength bar display on a mobile phone, that enumerates link quality as a range of 0-5, with 5 being very good. Values in the lower part of the range may indicate hardware problems such as port, cable, etc. that surface as signal integrity issues, leading to performance and other problems.
RcvPktsRate	Receive Packets Rate in Packets per second (Pps).
RcvDataRate (MBps)	Receive Data Rate in Megabytes per second (MBps).
RcvData (MB)	Receive Data in Megabytes (MB).
RcvPkts	The total number of received fabric data packets.
MulticastRcvPkts	The total number of multicast and collective packets received.

continued...



Option	Description
	This counter includes all valid packets and all packets with a header up to and including the DLID, where the DLID is within the configured range for multicast or collectives. Packets within the configured multicast or collective address space are counted, even if later checks determine the packet is unroutable or exceeds the SwitchInfo.MulticastFDBCap, SwitchInfo.CollectiveFDBCap, configured SwitchInfo.MulticastFDBTop, or configured SwitchInfo.CollectiveFDBTop.
RcvErrors	This counter indicates the total number of packets containing an error that were received by the port, including physical errors and malformed packets. It may indicate possible misconfiguration of a port, either by the SM or (more likely) by user intervention (e.g., using a tool such as <code>opaportconfig</code>).
RcvConstraintErrors	This counter is incremented when partition key or source LID violations are detected in a received packet, indicating a possible security issue or misconfiguration of device security settings.
RcvSwitchRelayErrors	This counter indicates the number of packets that were dropped due to internal routing errors. It is indicative of the possible misconfiguration of a switch by the SM.
RcvRemotePhysicalErrors	This counter indicates the number of downstream effects of signal integrity problems. It is indicative of an SI issue in the upstream path.
RcvFECN	When a device receives a packet with the FECN (Forward Explicit Congestion Notification) bit set to one, this counter is incremented.
RcvBECN	When a device receives a packet with the BECN (Backward Explicit Congestion Notification) bit set to one, this counter is incremented.
RcvBubble	This counter indicates the total number of "flit times" where one or more packets have started to be received, but the receiver received idle flits from the wire.
XmitPktsRate	Transmitted Packets Rate in Packets per second (Pps).
XmitDataRate (MBps)	Transmitted Data Rate in Megabytes per second (MBps).
XmitData (MB)	The total number of transmitted fabric data in Megabytes.
XmitPkts	The total number of fabric packets transmitted. This counter includes all fabric packet head flits transmitted with and without errors (such as <code>PktBadHead</code>).
MulticastXmitPkts	The total number of multicast and collective packets transmitted
XmitDiscards	The number of packets dropped due to one of the following errors: <ul style="list-style-type: none"> • Switch lifetime Limit exceeded • Switch head-of-queue lifetime limit exceeded • Output port not in active state • Packet length exceeded maximum fabric packet size for MTU for VL • Flow control disabled and insufficient credits available • SC2VL_t mapping invalid for given SC
XmitConstrainErrors	This counter is incremented when partition key or source LID violations are detected in a packet attempting to be transmitted, indicating a possible security issue or misconfiguration of device security settings.
XmitWait	This counter indicates the amount of time (in "flit times") any virtual lane had data but was unable to transmit (for reasons such as no credits available, or that the link was busy sending non-data packets such as link layer retraining or flow control).
XmitTimeCong	This counter indicates the total number of "flit times" that the counter was in a congested state.
continued...	



Option	Description
XmitWastedBW	This counter indicates the number of "flit times" where one or more packets have been started but the transmitters are forced to send idles due to bubbles.
XmitWaitData	This counter indicates the number of "flit times" where one or more packets have been started but interrupted due to bubbles in the ingress stream.
LocalLinkIntegrityErrors	This counter indicates the number of retries initiated by the link transfer layer. It may be indicative of low signal quality, or may be due to long or low quality cables.
FMConfigErrors	This counter indicates inconsistencies of low level SMA configuration on both sides of the link. It is indicative of the possible misconfiguration of a port, either by the SM, or (more likely) by user intervention (by using a tool such as <code>opaportconfig</code>).
ExcessiveBufferOverruns	This counter, associated with credit management, indicates an input buffer overrun. It may indicate possible misconfiguration of a port, either by the SM or (more likely) by user intervention (e.g. using a tool such as <code>opaportconfig</code>).
SwPortCongestion	This switch-only counter indicates the number of packets that were discarded as unable to transmit due to flow control issues.
MarkFECN	This counter indicates the total number of packets that were marked FECN by the transmitter due to congestion.
LinkErrorRecovery	This counter indicates the number of times the link has successfully completed the link error recovery process. If LQI is fluctuating toward low values AND this counter is increasing, it may be indicative of a bad link. Indication of a more severe signal quality problem.
LinkDowned	This counter indicates the total number of times the port has failed the link error recovery process and downed the link. A large number of occurrences of these events can cause disruptions to fabric traffic.
UncorrectableErrors	This counter indicates the number of unrecoverable internal device errors. It is indicative of a severe hardware defect or data corruption on the wire.

5.4.1.2 Performance (Charts)

Click the clock icon to choose the time option for the chart displays. **Current** show current values in a time window of $100 \times \text{refreshRate}$. **1H**, **2H**, **6H** and **24H** show the history values in the last 1, 2, 6, or 24 hours respectively.

5.4.1.2.1 Received Data/Packets Rate

The chart displays received traffic data through the selected port in the ports table either in Data Rate in Byte per second (Bps) or in Packets Rate in Packet per second (Pps) by a user selection.

5.4.1.2.2 Transmitted Data/Packets Rate

The chart displays transmitted traffic data through the selected port in the ports table either in Data Rate in Byte per second (Bps) or in Packets Rate in Packet per second (Pps) by a user selection.



5.4.2 Connectivity (Node)

The Connectivity section summarizes each port's link status in a table. Each active port is followed by its neighboring port, and inactive ports are presented in gray. Data shown in the table are the port counters for each involved port and the table columns are configurable with the top-right corner button.

Figure 31. Connectivity Example

Performance Connectivity Properties MOOSE_STL_SWITCH119								
Node Name	Node GUID	Port #	Link State	Physical Link State	Active Link Speed	Link Speed Supported	# Link Down	
MOOSE_STL_SWITCH119	0x0002b31800000077	1	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH73	0x0002b31800000049	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	2	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH75	0x0002b31800000046	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	3	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH77	0x0002b31800000044	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	4	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH79	0x0002b3180000004f	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	5	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH81	0x0002b31800000051	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	6	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH83	0x0002b31800000053	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	7	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH85	0x0002b31800000055	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	8	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH87	0x0002b31800000057	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	9	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH89	0x0002b31800000059	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	10	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH91	0x0002b3180000005b	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	11	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH93	0x0002b3180000005d	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	12	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH95	0x0002b3180000005f	24 (Neighbor)	Active	Link Up	25 Gbps	12.5 Gbps or 25 Gbps	0	
MOOSE_STL_SWITCH119	0x0002b31800000077	13	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	14	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	15	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	16	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	17	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	18	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	19	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	20	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	21	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	22	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	23	Inactive				N/A	
MOOSE_STL_SWITCH119	0x0002b31800000077	24	Inactive				N/A	

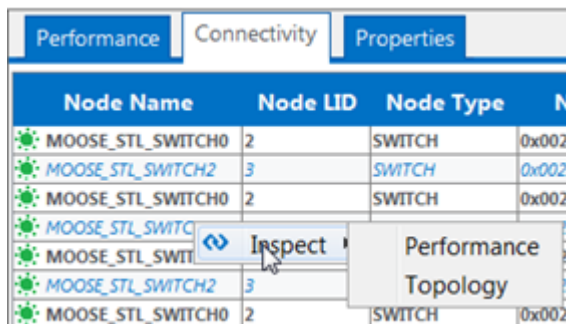
Note: If link speed and link width are displayed with red color, it indicates the following situations:

- Slow speed – active speed/width is lower than configured
- Configured slow speed – enabled speed is lower than supported
- Mismatched speed - active speed/width does not match neighboring port

Right-click on a port to open a menu to quickly jump to this port's performance page or topology page.

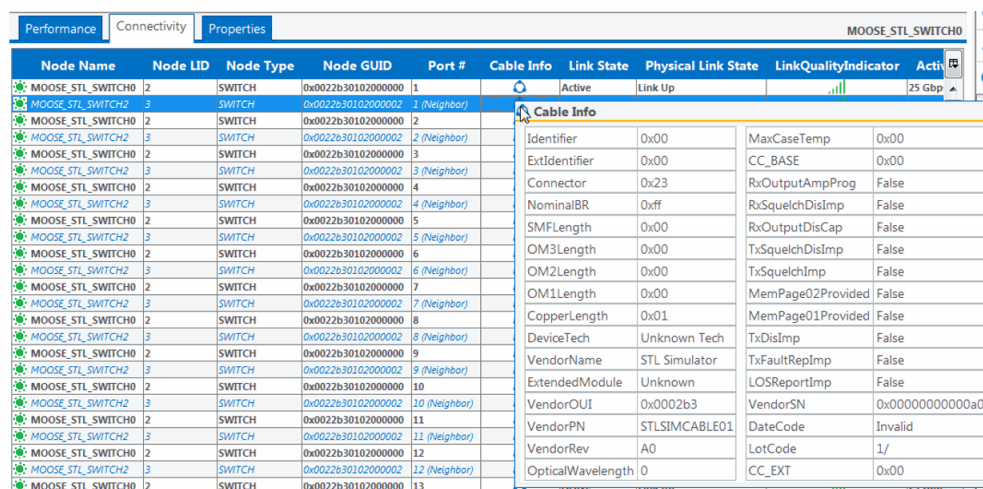


Figure 32. Jump to Performance or Topology



Left-click on a Cable Info column to directly display cable information about this port.

Figure 33. Cable Information Jump



The fields in the table below can be selected from the drop-down list at the upper-right corner of the section.

Table 3. Connectivity Data

Option	Description
Node Name	Node Name is typically assigned by the system administrator based on the desired naming convention. It is typically the same or derived from the Linux hostname for the server. Once selected by the sysadmin, the value persists across OS reboot.
Node GUID	GUID of the HFI or Switch.
Port #	The link port number this SMP came on in.
Link State	Port State
Physical Link State	Physical Port State
continued...	



Option	Description
LinkQualityIndicator	This is a status indicator, similar to the signal strength bar display on a mobile phone, that enumerates link quality as a range of 0-5, with 5 being very good. Values in the lower part of the range may indicate hardware problems such as port, cable, etc. that surface as signal integrity issues, leading to performance and other problems.
Link Width	The possible values for link width: 1x, 2x, 3x, 4x.
Link Width Enabled	The set of link widths that the LNI protocol negotiates. The LNI protocol uses only LW.E to negotiate link width during LNI.
Link Width Supported	The link widths the port can negotiate to during LNI. In some implementations firmware/driver and/or local device settings may restrict this value further.
Active Link Speed	The link speed active value of this port.
Link Speed Enabled	The link speed enabled value of this port.
Link Speed Supported	The link speed supported value of this port.
RcvData (MB)	Receive Data Rate in Megabytes per second (MBps).
RcvPkts	The total number of received fabric data packets.
MulticastRcvPkts	The total number of multicast and collective packets received. This counter includes all valid packets and all packets with a header up to and including the DLID, where the DLID is within the configured range for multicast or collectives. Packets within the configured multicast or collective address space are counted, even if later checks determine the packet is unroutable or exceeds the SwitchInfo.MulticastFDBCap, SwitchInfo.CollectiveFDBCap, configured SwitchInfo.MulticastFDBTop or configured SwitchInfo.CollectiveFDBTop.
RcvErrors	This counter indicates the total number of packets containing an error that were received by the port, including physical errors and malformed packets. It may indicate possible misconfiguration of a port, either by the SM or (more likely) by user intervention (e.g., using a tool such as opaportconfig).
RcvConstraintErrors	This counter is incremented when partition key or source LID violations are detected in a received packet, indicating a possible security issue or misconfiguration of device security settings.
RcvSwitchRelayErrors	This counter indicates the number of packets that were dropped due to internal routing errors. It is indicative of the possible misconfiguration of a switch by the SM.
RcvRemotePhysicalErrors	This counter indicates the number of downstream effects of signal integrity problems. It is indicative of an SI issue in the upstream path.
RcvFECN	When a device receives a packet with the FECN (Forward Explicit Congestion Notification) bit set to one, this counter is incremented.
RcvBECN	When a device receives a packet with the BECN (Backward Explicit Congestion Notification) bit set to one, this counter is incremented.
RcvBubble	This counter indicates the total number of "flit times" where one or more packets have started to be received, but the receiver received idle flits from the wire.
XmitData (MB)	The total number of transmitted fabric data in Megabytes.
XmitPkts	The total number of fabric packets transmitted. This counter includes all fabric packet head flits transmitted with and without errors (such as PktBadHead).
MulticastXmitPkts	The total number of multicast and collective packets transmitted
continued...	



Option	Description
XmitDiscards	The number of packets dropped due to one of the following errors: <ul style="list-style-type: none">• Switch lifetime Limit exceeded• Switch head-of-queue lifetime limit exceeded• Output port not in active state• Packet length exceeded maximum fabric packet size for MTU for VL• Flow control disabled and insufficient credits available• SC2VL_t mapping invalid for given SC
XmitConstrainErrors	This counter is incremented when partition key or source LID violations are detected in a packet attempting to be transmitted, indicating a possible security issue or misconfiguration of device security settings.
XmitWait	This counter indicates the amount of time (in "flit times") any virtual lane had data but was unable to transmit (for reasons such as no credits available, or that the link was busy sending non-data packets such as link layer retraining or flow control).
XmitTimeCong	This counter indicates the total number of "flit times" that the counter was in a congested state.
XmitWastedBW	This counter indicates the number of "flit times" where one or more packets have been started but the transmitters are forced to send idles due to bubbles.
XmitWaitData	This counter indicates the number of "flit times" where one or more packets have been started but interrupted due to bubbles in the ingress stream.
LocalLinkIntegrityErrors	This counter indicates the number of retries initiated by the link transfer layer. It may be indicative of low signal quality, or may be due to long or low quality cables.
FMConfigErrors	This counter indicates inconsistencies of low level SMA configuration on both sides of the link. It is indicative of the possible misconfiguration of a port, either by the SM, or (more likely) by user intervention (by using a tool such as <code>opaportconfig</code>).
ExcessiveBufferOverruns	This counter, associated with credit management, indicates an input buffer overrun. It may indicate possible misconfiguration of a port, either by the SM or (more likely) by user intervention (e.g., using a tool such as <code>opaportconfig</code>).
SwPortCongestion	This switch-only counter indicates the number of packets that were discarded as unable to transmit due to flow control issues.
MarkFECN	This counter indicates the total number of packets that were marked FECN by the transmitter due to congestion.
LinkErrorRecovery	This counter indicates the number of times the link has successfully completed the link error recovery process. If LQI is fluctuating toward low values AND this counter is increasing, it may be indicative of a bad link. Indication of a more severe signal quality problem.
LinkDowned	This counter indicates the total number of times the port has failed the link error recovery process and downed the link. A large number of occurrences of these events can cause disruptions to fabric traffic.
UncorrectableErrors	This counter indicates the number of unrecoverable internal device errors. It is indicative of a severe hardware defect or data corruption on the wire.

5.4.3 Node Properties

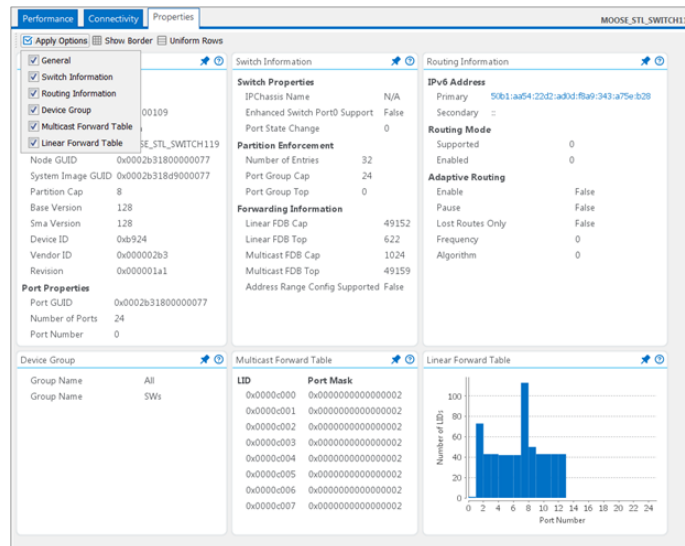
The Node Properties section displays the properties for a selected node. The tool bar contains the following buttons:

- Apply Options – select which properties to display



- Show Border – turns on/off borders on property items
- Uniform Rows – turns on/off alternative row colors

Figure 34. Node Properties Example



5.4.3.1 General Properties (Node Performance)

The General panel includes device information for a host or switch. This information is not editable by the user.

Table 4. Device Information

Property	Description
Node State	The current state of the node.
LID	The Local Identifier (LID) is an address assigned at an Endpoint by the Fabric Manager, unique within the Fabric, used for forwarding Packets.
Node Name	Node Name is typically assigned by the system administrator based on the desired naming convention. It is typically the same or derived from the Linux hostname for the server. Once selected by the sysadmin, the value should persist across OS reboot.
Type	Node Type: UNKNOWN, HFI, SWITCH, ROUTER, OTHER.
Node GUID	GUID of the HFI or switch.
System Image GUID	GUID of system, same GUID for all OPA devices in a single system.
Partition Cap	Number of entries in the Partition Table for end ports.
Base Version	Supported MAD Base Version.
Sma Version	Supported Subnet Management Class (SMP) Version.
Device ID	Device ID information as assigned by device manufacturer.
Vendor ID	Device vendor, per IEEE.
Revision	Device revision, assigned by manufacturer.

**Table 5. Port Properties**

Property	Description
Port GUID	The globally unique identifier (GUID) for the node.
Number of Ports	The number of ports on the node.
Port Number	The link port number this SMP came on in.

5.4.3.2 Switch Information (Properties)

Fabric management switch node information, including Switch Properties, Partition Key Enforcement and Forwarding Information.

Table 6. Switch Properties

Property	Description
Enhanced Switch Port 0 Support	When set to true, indicates switch port 0 supports enhanced functions (TCA port). When set to false, indicates switch port 0 is a base switch port 0.
Life Time Value	Sets the time a packet can live in the switch.
Port State Change	It is set to one anytime the PortState component in the PortInfo of any ports transitions from Down to Initialize, Initialize to Down, Armed to Down, or Active to Down as a result of link state machine logic. Changes in Portstate resulting from SubnSet() do not change this bit. This bit is cleared by writing one, writing zero is ignored.

Table 7. Partition Key Enforcement

Property	Description
Number of Entries	Partition Enforcement Cap. It specifies the number of entries in the partition enforcement table per physical port. Zero indicates that partition enforcement is not supported by the switch.
Port Group Cap	The max size of the port group table.
Port Group Top	The current size of the port group table.

Table 8. Forwarding Information

Property	Description
Linear FDB Cap	Number of entries supported in the Linear Unicast Forwarding Table (starting at LID=0x0000 going up). LinearFDBCap= 0 indicates that there is no Linear Forwarding Table.
Linear FDB Top	Indicates the top of the linear forwarding table. Packets received with unicast DLIDs greater than this value are discarded by the switch. A valid LinearFdbTop is less than LinearFdbCap. This component applies only to switches that implement linear forwarding tables and is ignored by switches that implement random forwarding tables.
Multicast FDB Cap	Number of entries supported in the Multicast Forwarding Table (starting at LID=0xC000 going up).
Multicast FDB Top	Indicates the upper bound of the range of the multicast forwarding table. Packets received with MLIDs greater than MulticastFDBTop are considered to be outside the range of the Multicast Forwarding Table. A valid MulticastFDBTop is less than MulticastFDBCap + 0xC000. This component applies only to switches that implement the optional multicast forwarding service. A switch shall ignore the
continued...	



Property	Description
	MulticastFDBTop component if it has the value zero. The initial value for MulticastFDBTop shall be set to zero. A value of 0xBFFF means there are no MulticastForwardingTable entries.
Address Range Config Supported	A flag to indicate if address range for multicast and collectives can be configured.

5.4.3.3 Routing Information (Performance Properties)

Routing information for a switch, including IP Address, Routing Mode and Adaptive Routing.

Table 9. IP Address

Property	Description
IPv6 Address	IP Address to allow query of the IP address by which additional switch chassis management functions may be accessed.
IPv4 Address	IP Address to allow query of the IP address by which additional switch chassis management functions may be accessed.

Table 10. Routing Mode

Property	Description
Supported	Supported routing mode. Routing mode enables alternate routing table modes
Enabled	Enabled routing mode. Routing mode enables alternate routing table modes.

Table 11. Adaptive Routing

Property	Description
Enable	Enables/Disables Adaptive Routing.
Pause	Temporarily disables AR.
Lost Routes Only	Indicates that AR should only be done for failed links.
Frequency	Value expands to $2^F \times 64\text{ms}$.
Algorithm	0 = Random, 1 = Greedy, 2 = Random Greedy.
Threshold	<p>This is the threshold, above which, adaptive routing will trigger to start offloading DLIDs to alternate ports with no or low or no congestion. The thresholds are in terms of the distribution over the available buffering for a given port VL. Values 1-4 are considered high congestion at the current egress port. Values 5-7 are considered medium congestion. The weight values are as follows:</p> <ul style="list-style-type: none"> 0 = switch default tuning 1 = has consumed nearly all of the buffering(or Tags) for a VL (~100%) 2 = 90% 3 = 80% 4 = 70% 5 = 65% 6 = 60% 7 = 55%



5.4.3.4 Device Group

Device Group properties include information regarding the device groups that have been set up on a host or switch. This information is available for hosts and switches and is not editable by the user.

5.4.3.5 Multicast Forwarding

Multicast forwarding properties include information regarding the number of ports followed by the list of port numbers on a switch to which packets of a specific MLID value are forwarded. The MLID is the destination address for the packets, and packets sent to the MLID address will be sent to the ports listed for that MLID. This information is only available for switches and is not editable by the user.

5.4.3.6 Linear Forwarding

Linear forwarding properties include information regarding the port numbers on a switch to which packets of a specific LID value are forwarded. A port number of 255 indicates that the switch does not accept packets containing the specified LID. This information is only available for switches and is not editable by the user.

5.4.3.7 SC2SL Mapping Table

This mapping table converts the SC from a received packet back into the Service Level the SC is associated with.


5.5 Port Performance

Port performance of a device node is displayed in three sections:

- Performance
- Connectivity
- Properties



Figure 35. Port Performance Example



Performance Tab

The performance tab (shown above) includes two sections: [Performance \(Port\)](#) on page 51 and [Counters](#). The performance section displays received and transmitted packets and data to facilitate comparison between received and transmitted data. The counters section displays unit and error counters for the selected port.

Connectivity Tab

This section displays the same content as the [Connectivity \(Node\)](#) on page 43.

Properties Tab

This section displays Port Performance properties. The Port Performance properties include Device Information, Link, Link Connection, Port Capability, Virtual Lane, Diagnostics, Partition Key Enforcement, Management, Flit Control, Port Error Actions, Miscellaneous, MTU by VL, HoQLife By VL, VL Stall Count By VL, QSFP Interpreted Cable Information, SC2VLT Mapping Table, SC2VLNT Mapping Table, and Link Down Error Log.

5.5.1 Performance (Port)

The performance subpage includes two sections: Performance and Counters. The performance section displays received and transmitted packets and data side-by-side with the same scale so you can compare received and transmitted data. The counters section displays utilization and error counters for the selected port.



5.5.1.1 Charts (Port Performance)

See [Performance \(Charts\)](#) on page 42.

5.5.1.2 Counters

The counters section displays receive, transmit, and other counters in the Port Counters card. When a counter field is selected from the card, the top-right chart displays history trend data for the selected field; the bottom-right card displays the description of the field. The descriptions are based on [Appendix A Port Counters Overview](#). Field names can be slightly different.

Hide/Show Border – a toggle button that turns on/off borders on the display

Alternating/Uniform Rows – a toggle button that turns on/off alternate row colors on the display

5.5.1.2.1 Port Counters

This section displays the Receive and Transmit data shown below. Refer to [Appendix A Port Counters Overview](#) for more information.

Table 12. Receive Counters

Data Displayed	Description
RcvData	The total number of fabric packet flits received. This counter includes all fabric packet flits received with and without errors (such as <code>PktBadTail</code>). Only includes flits actually received (for example, <code>PktLen</code> in header may exceed actual flits received in certain error situations, in which case only actual flits received are counted. In IB mode: The total number of data octets, divided by 4, received.
RcvPkts	The total number of fabric packets received. This counter includes all fabric packet head flits received with and without errors (such as <code>PktBadHead</code>).
MulticastRcvPkts	The total number of multicast and collective packets received. This counter includes all valid packets and all packets with a header up to and including the DLID, where the DLID is within the configured range for multicast or collectives. Packets within the configured multicast or collective address space are counted, even if later checks determine the packet is unroutable or exceeds the <code>SwitchInfo.MulticastFDBCap</code> , <code>SwitchInfo.CollectiveFDBCap</code> , configured <code>SwitchInfo.MulticastFDBTop</code> or configured <code>SwitchInfo.CollectiveFDBTop</code> .
RcvErrors	This counter indicates the total number of packets containing an error that were received by the port, including physical errors and malformed packets. It may indicate possible misconfiguration of a port, either by the SM or (more likely) by user intervention (e.g., using a tool such as <code>opaportconfig</code>).
RcvConstraintErrors	This counter is incremented when partition key or source LID violations are detected in a received packet, indicating a possible security issue or misconfiguration of device security settings.
RcvSwitchRelayErrors	This counter indicates the number of packets that were dropped due to internal routing errors. It is indicative of the possible misconfiguration of a switch by the SM.
continued...	



Data Displayed	Description
RcvRemotePhysicalErrors	This counter indicates the number of downstream effects of signal integrity problems. It is indicative of an SI issue in the upstream path.
RcvFECN	When a device receives a packet with the FECN (Forward Explicit Congestion Notification) bit set to one, this counter is incremented.
RcvBECN	When a device receives a packet with the BECN (Backward Explicit Congestion Notification) bit set to one, this counter is incremented.
RcvBubble	This counter indicates the total number of "flit times" where one or more packets have been started to be received, but receiver received idle flits from the wire.

Table 13. Transmit Counters

Data Displayed	Description
XmitData	The total number of fabric packet flits transmitted. Does not include idle nor other LF command flits. This counter includes all fabric packet flits transmitted without and with errors (such as PktBadTail). Only includes flits actually transmitted (e.g., PktLen in header may exceed actual flits transmitted in certain error situations, in which case only actual flits transmitted are counted).
XmitPkts	The total number of fabric packets transmitted. This counter includes all fabric packet head flits transmitted with and without errors (such as PktBadHead).
MulticastXmitPkts	The total number of multicast and collective packets transmitted
Transmit Discards	The number of packets dropped due to one of the following errors: <ul style="list-style-type: none"> Switch lifetime Limit exceeded Switch head-of-queue lifetime limit exceeded Output port not in active state Packet length exceeded maximum fabric packet size for MTU for VL Flow control disabled and insufficient credits available SC2VL_t mapping invalid for given SC
XmitConstraintErrors	This counter is incremented when partition key or source LID violations are detected in a packet attempting to be transmitted, indicating a possible security issue or misconfiguration of device security settings.
XmitWait	This counter indicates the amount of time (in "flit times") any virtual lane had data but was unable to transmit (for reasons such as no credits available, or that the link was busy sending non-data packets such as link layer retraining or flow control).
XmitTimeCong	This counter indicates the total number of "flit times" that the counter was in a congested state.
XmitWastedBW	This counter indicates the number of "flit times" where one or more packets have been started but the transmitters is forced to send idles due to bubbles.
XmitWaitData	This counter indicates the number of "flit times" where one or more packets have been started but interrupted due to bubbles in the ingress stream.



5.5.1.2.2 Other Counters

This section displays the data shown in the following table. Refer to [Appendix A Port Counters Overview](#) for more information.

Table 14. Other Counters

Data Displayed	Description
LocalLinkIntegrityErrors	This counter indicates the number of retries initiated by the link transfer layer. It may be indicative of low signal quality, or may be due to long or low quality cables.
FMConfigErrors	This counter indicates inconsistencies of low level SMA configuration on both sides of the link. It is indicative of the possible misconfiguration of a port, either by the SM, or (more likely) by user intervention (by using a tool such as <code>opaportconfig</code>).
ExcessiveBufferOverruns	This counter, associated with credit management, indicates an input buffer overrun. It may indicate possible misconfiguration of a port, either by the SM or (more likely) by user intervention (e.g. using a tool such as <code>opaportconfig</code>).
SwPortCongestion	This switch-only counter indicates the number of packets that were discarded as unable to transmit due to flow control issues.
MarkFECN	This counter indicates the total number of packets that were marked FECN by the transmitter due to congestion.
LinkErrorRecovery	This counter indicates the number of times the link has successfully completed the link error recovery process. If LQI is fluctuating toward low values AND this counter is increasing, it may be indicative of a bad link. Indication of a more severe signal quality problem.
LinkDowned	This counter indicates the total number of times the port has failed the link error recovery process and downed the link. A large number of occurrences of these events can cause disruptions to fabric traffic.
UncorrectableErrors	This counter indicates the number of unrecoverable internal device errors. It is indicative of a severe hardware defect or data corruption on the wire.
LinkQualityIndicator	This is a status indicator, similar to the signal strength bar display on a mobile phone, that enumerates link quality as a range of 0-5, with 5 being very good. Values in the lower part of the range may indicate hardware problems such as port, cable, etc. that surface as signal integrity issues, leading to performance and other problems.

5.5.2 Connectivity (Port Performance)

See [Connectivity \(Node\)](#) on page 43.

5.5.3 Port Properties

Port properties are not editable by the user.

To view the port properties:

1. From the main screen, click the **Performance** tab.
2. From the left resource panel, click the **Device Types** drop-down arrow to reveal the hierarchical view.



3. For any node, click + (plus sign) to expand the tree and select a port for a host or a switch.
4. Click the **Properties** tab.

5.5.3.1 Device Information (Performance Properties)

This panel displays general port information.

5.5.3.2 Link (Performance Properties)

This panel displays link information for this port.

5.5.3.3 Link Connection (Performance Properties)

This panel displays link connection information that includes the link connected to, neighbor mode, and port mode.

5.5.3.4 Port Capability

This panel displays the port capabilities that are supported for the current port.

5.5.3.5 Virtual Lane

This panel displays the information about the virtual lanes of the port.

5.5.3.6 Diagnostics (Performance)

This panel displays General diagnostic information for the current port.

5.5.3.7 Partition Key Enforcement

This panel displays the information about packet partition key enforcement for the current port. Choices are true, false or N/A (not applicable).

5.5.3.8 Management (Performance)

This panel displays subnet manager and management key information for the current port.

5.5.3.9 Flit Control (Performance)

This panel displays flit control information that includes the interleave and preemption.

5.5.3.10 Port Error Actions

This panel displays a port error action mask.

5.5.3.11 Miscellaneous (Performance Properties)

This panel displays miscellaneous information such as IP Address, Buffer Units and Subnet.



5.5.3.12 MTU By VL

This panel displays RW/HS-E Neighbor MTU values per VL.

5.5.3.13 HoQLife By VL

This panel displays the time a packet can live at the head of a VL queue.

5.5.3.14 VL Stall Count By VL

This panel displays the number of sequential packets dropped that causes the port to enter the VLStalled state.

5.5.3.15 QSFP Interpreted Cable Information

This panel displays QSFP port type cable information.

5.5.3.16 SC2VLT Mapping Table

This panel displays the Service Channel (SC) to Virtual Lane Transmit (VLT) mapping table.

5.5.3.17 SC2VLNT Mapping Table

This panel displays the Service Channel (SC) to Virtual Lane Non-transmit (VLNT) mapping table.

5.5.3.18 Link Down Error Log


This panel displays the Link Down Reason and Neighbor Link Down Reason in time that indicates the reason the port transitioned from Link states Init, Armed, Active to Down.

6.0 Topology

The Topology tab displays the topology for selected resources. On the left of the page is a resource tree panel that organizes subnet devices by different categories, which allows you to select resources. The main panel shows the topology of a selected resource and its content changes based on the type of selected resource.

The main panel displays the topology and related information. The left section showing the overall topology outline and selected resources corresponds to the topology of selected resources shown at the right. The bottom section displays detailed information for the currently selected resources and its contents change based on the type of resource selected.

Figure 36. Topology Example



See the topics below for details on each section of the Topology tab:

- [Resource Tree](#) on page 35
- [Outline](#) on page 58
- [Selected Resources](#) on page 58
- [Local Topology](#) on page 58

6.1 Resource Tree (Topology)

The Resource Tree panel is the same as on the Performance tab except that it allows multiple selections on resources with the same type, for example, multiple nodes, multiple ports or multiple device groups, virtual fabrics. However, when a user selects among different resource types, only one resource is allowed.

6.2 Outline

The Outline section shows the topology outline for a subnet, highlighting the selected resources. Use the **Enlarge** button at the bottom to open a pop up a window of the topology outline for a closer look at the topology.

6.3 Selected Resources

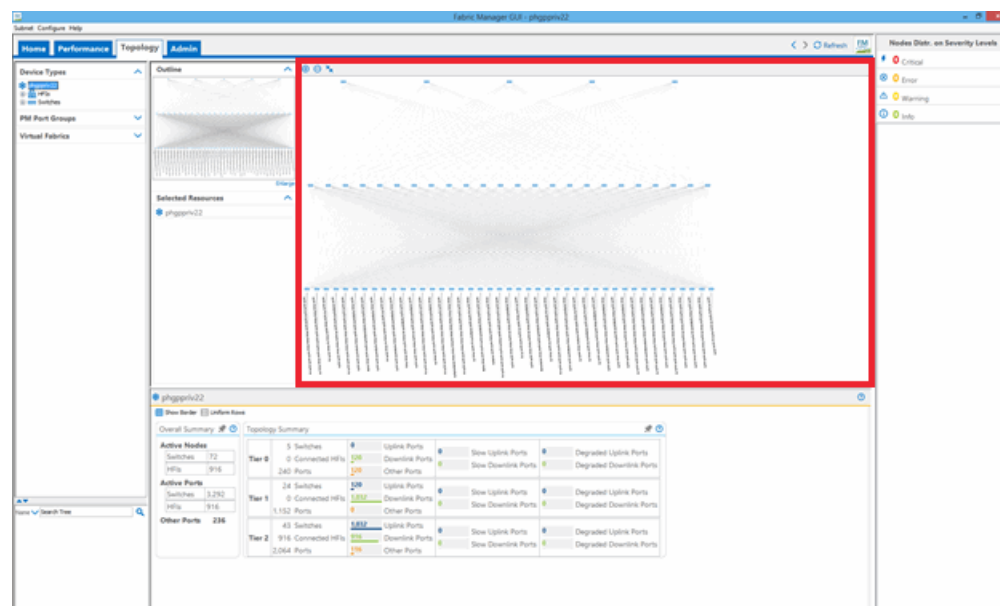
This section displays the selected resources when you select multiple resources that cannot be shown on the same screen as the device tree panel.

6.4 Local Topology

This section displays a small and clear topology of selected resources for a closer look. Although it shows only local topology, it shares the same topology architecture as the outline graph to indicate how the local topology relates to the subnet topology overall.

Buttons at the top can be used to zoom in/out of the graph or fit the graph into the window. Each element shown on the local topology is clickable for a single selection. Hover the mouse over a device node to show the name of the node and to highlight its connections on the topology graph to see how it is connected.

Figure 37. Local Topology Example





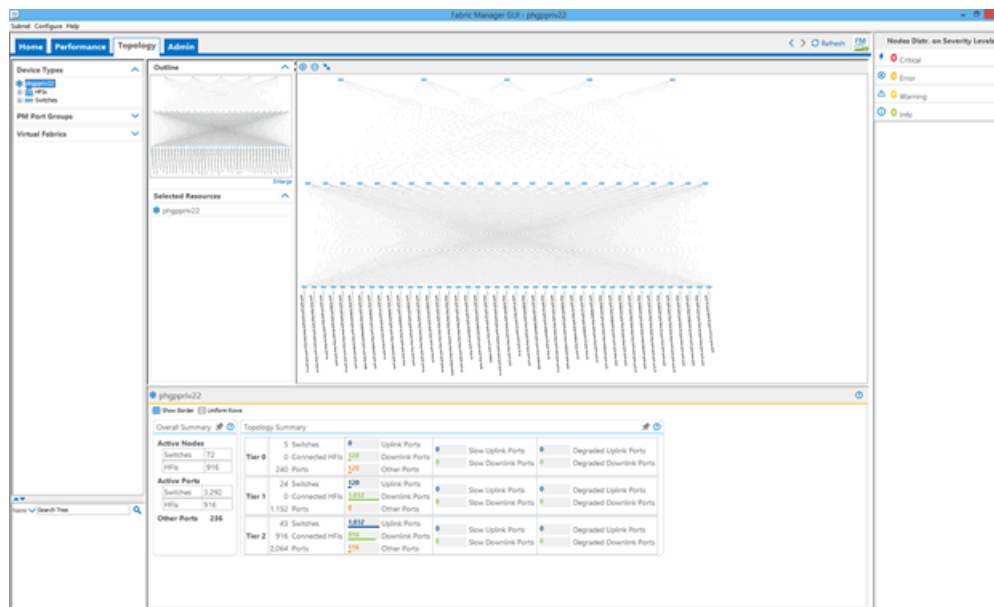
The content of this panel depends upon the type of resources selected, as described in the topics linked below.

- [Device Node](#) on page 60
- [Links](#) on page 61
- [Routes](#) on page 61

6.4.1 Device Set (Topology)

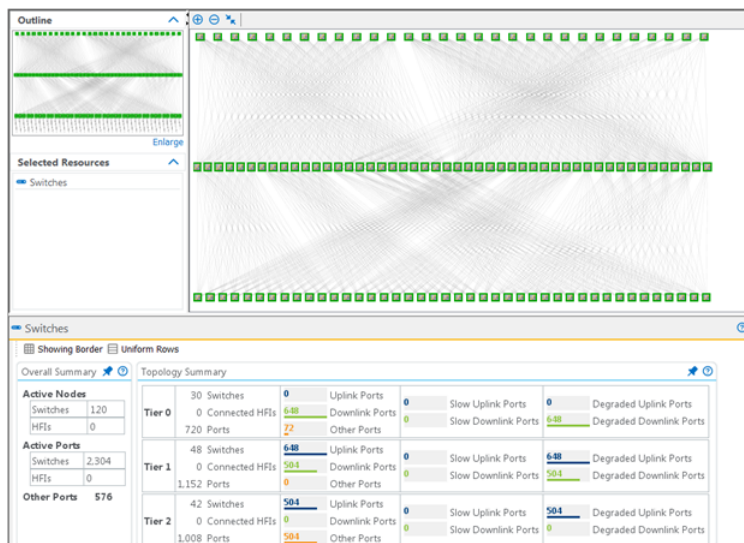
Selection on a whole subnet is a special case of the device set selection. This panel shows the whole topology of the subnet.

Figure 38. Device Set Example



Selections on one or more device types, device groups, or virtual fabrics are also treated as a device set. This panel shows the topology among all involved devices, and each device is highlighted with a green border.

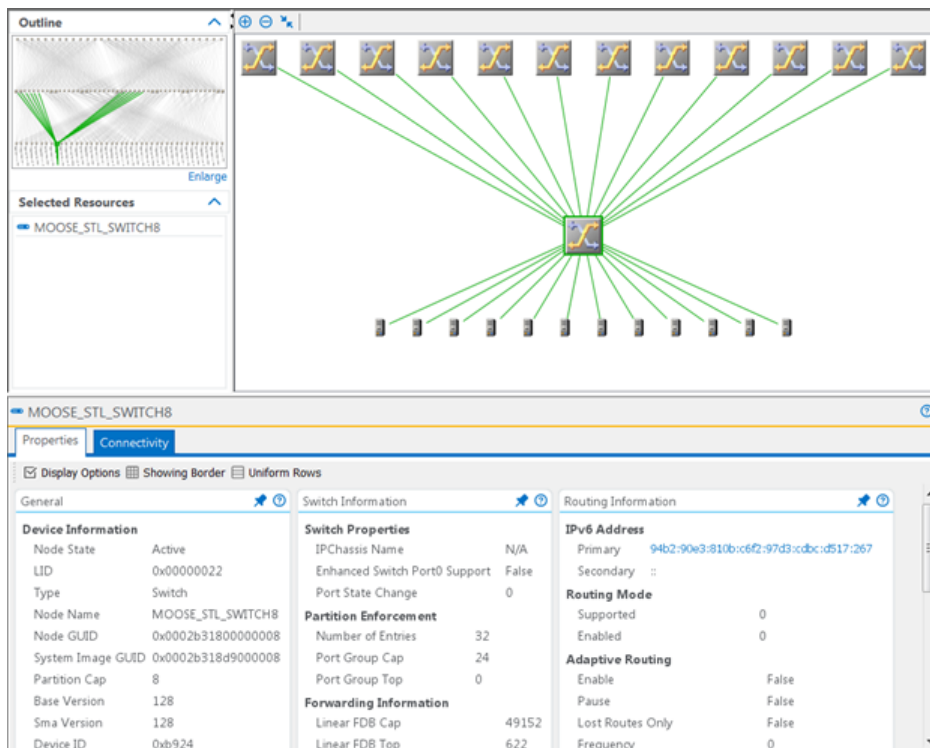
Figure 39. Device Groups Example



6.4.2 Device Node

When a device node is selected, this panel shows the connections of the node. The selected node is highlighted with a green border to distinguish it from its neighbors.

Figure 40. Device Node Example





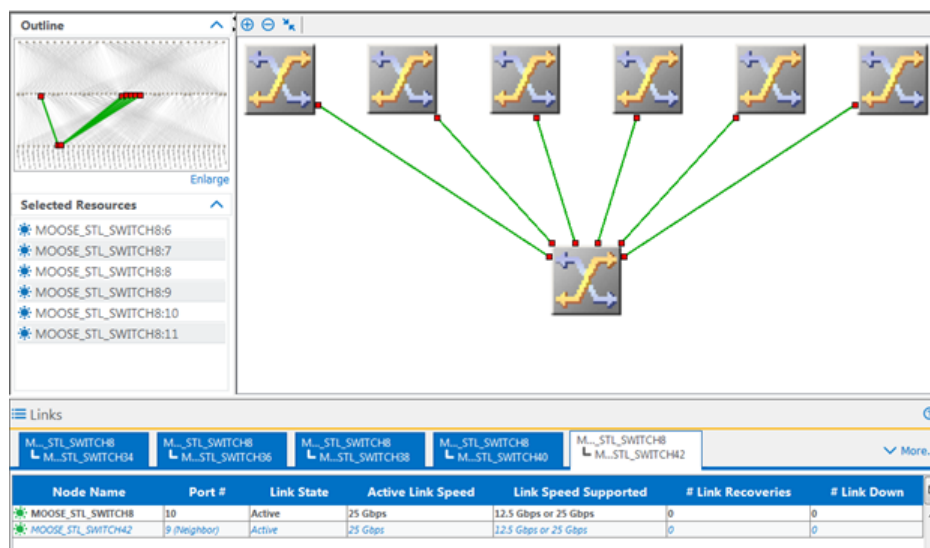
Detailed Information Panel

When a device node is selected this panel displays this node's properties and connectivity table, which is the same data represented with different views. For more detailed information see [Connectivity \(Node\)](#) on page 43 and [Node Performance](#) on page 39.

6.4.3 Links

When you select one or more ports, this panel shows and highlights each port's connection.

Figure 41. Links Example



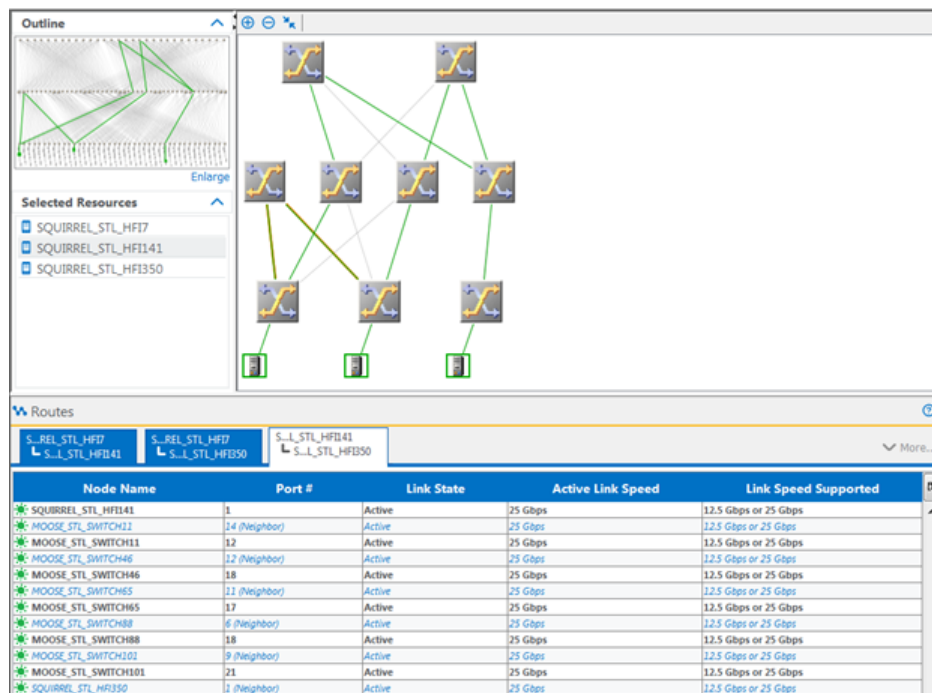
Detailed Information Panel

When multiple ports are selected, this panel shows each port's connectivity table in a tabbed pane. Use the **More** button on the top-right corner to display additional links. When the number of links is less than 5, the button is disabled.

6.4.4 Routes

When multiple device nodes are selected, this panel shows the local topology of the selected nodes. It also highlights the routes among the nodes to show how they connect to each other.

Figure 42. Routes Example



Detailed Information Panel

When multiple device nodes are selected, this panel shows the pairwise route among the selected nodes in a tabbed pane. Each tab shows the route as a connectivity table (see [Connectivity \(Node\)](#) on page 43). It shows the whole path from one port to another and each port's performance data.



7.0 Admin Tab

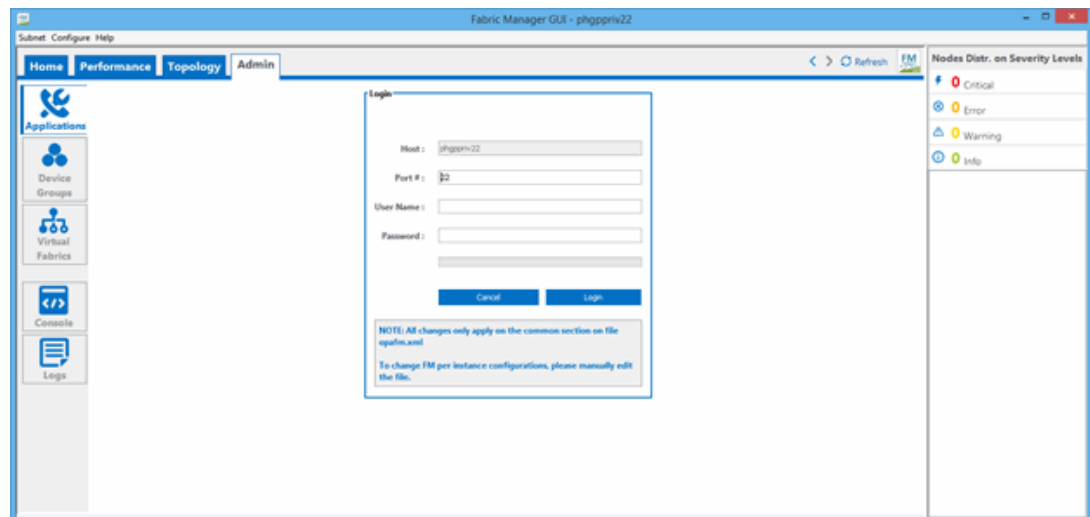
The Admin tab provides tools to manage a subnet. The control panel on the left lists all available functions. The right side displays the corresponding content for a selected function.

The first time you select one of the icons on the control panel, you must enter a valid username and password to log in, which gives you access to the administrative features. After the initial login, admin features, there is no need to enter the credentials again.

The following functions are described in the linked topics:

- [Applications Management](#) on page 63
- [Device Group Management](#) on page 65
- [Virtual Fabric Management](#) on page 65
- [Interactive Console](#) on page 66
- [Log View](#) on page 69

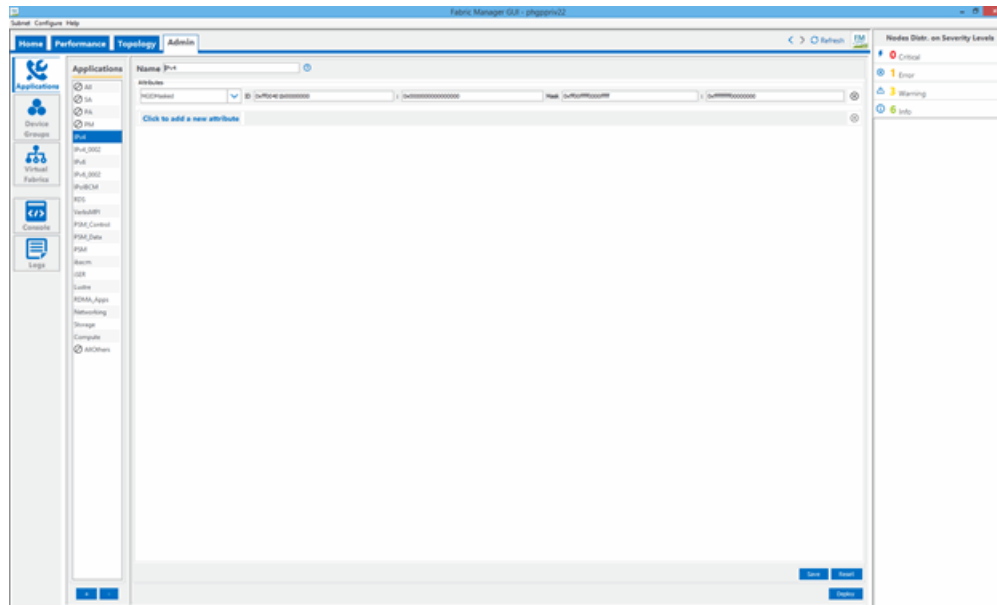
Figure 43. Admin Tab



7.1 Applications Management

The left panel lists the application names with two buttons on the bottom that allow you to add or remove one application.

Figure 44. Applications Management Example



The right panel displays and allows you to edit the configuration of one application. An application's configuration is made up of a set of attributes that are represented as attribute bars.


The **Click to add a new attribute** button allows you to select an attribute type from a drop-down list and specify values for the attribute. Use the **Save** and **Reset** buttons to save the new configuration or reset the configuration to original values.

Fabric Manager GUI validates the configuration before you save or remove an application, checking for duplicate names, invalid value ranges, for example a min value that is larger than the max value, and application references. An application with reference to others cannot be removed or renamed until you resolve any reference issues.

The **Deploy** button on the bottom-right corner deploys the applications to the master Subnet Manager in the following way:

1. Fabric Manager GUI makes a copy of current configuration file `opafm.xml` with a name such as `opafm.xml.<timestamp>.fv`.
2. Fabric Manager GUI maintains the number of backup files to 16 by removing older files.
3. Fabric Manager GUI replaces `opafm.xml` with the new configurations.

To see the new configurations take effect, restart and relaunch Fabric Manager GUI.

Some applications are predefined and they cannot be edited or removed. These applications are indicated with  on the view only mode on the right panel.

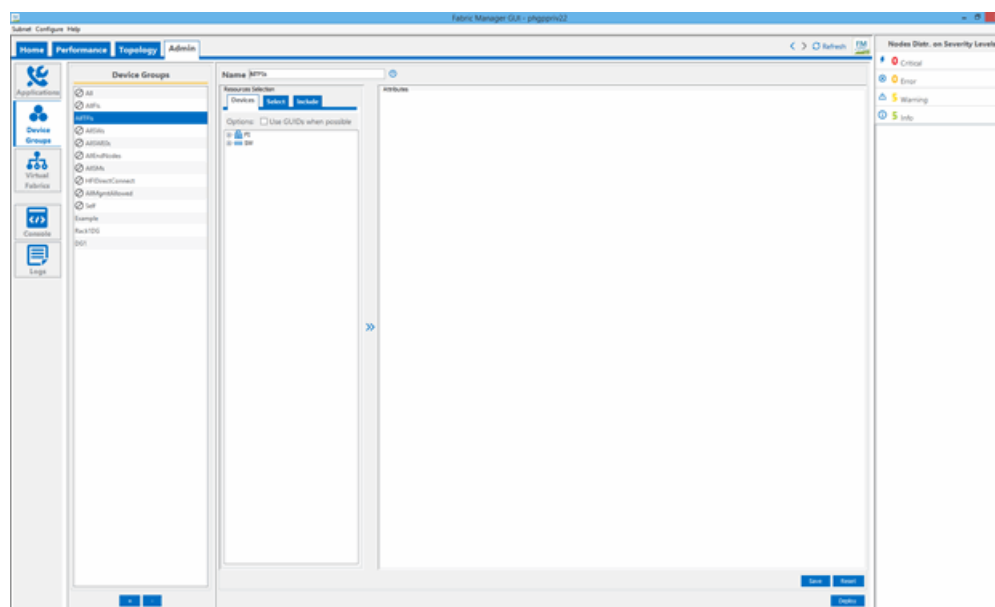


7.2 Device Group Management

Device Groups management is similar to Applications management, but it has a different way to add or change an attribute. On the middle center is a resource selection panel that displays current resource selections and allows you to change them.

The **Devices** tab shows selected devices that can be described either by name or GUID. The **Select** tab shows the selection on special types of devices. The **Include** tab allows a Device Group to include another one.

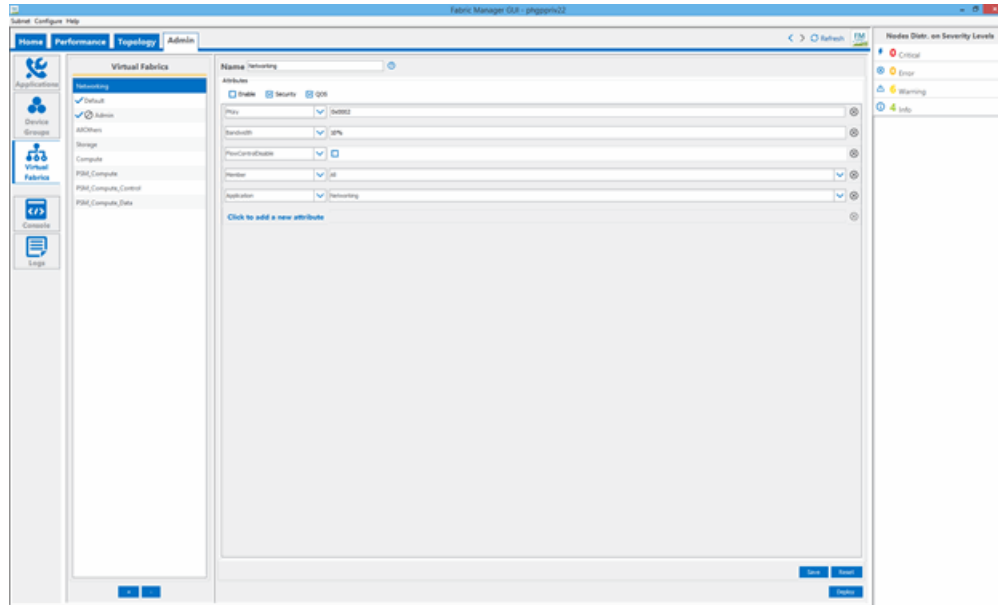
Figure 45. Device Group Management Example



7.3 Virtual Fabric Management

This management tool is the same as [Applications Management](#) on page 63 except that groups are replaced with virtual fabrics.

Figure 46. Virtual Fabric Management Example



On the right panel, you can specify basic administrator policies, for example, whether to enable one virtual fabric and whether it should provide security control and/or QOS control.

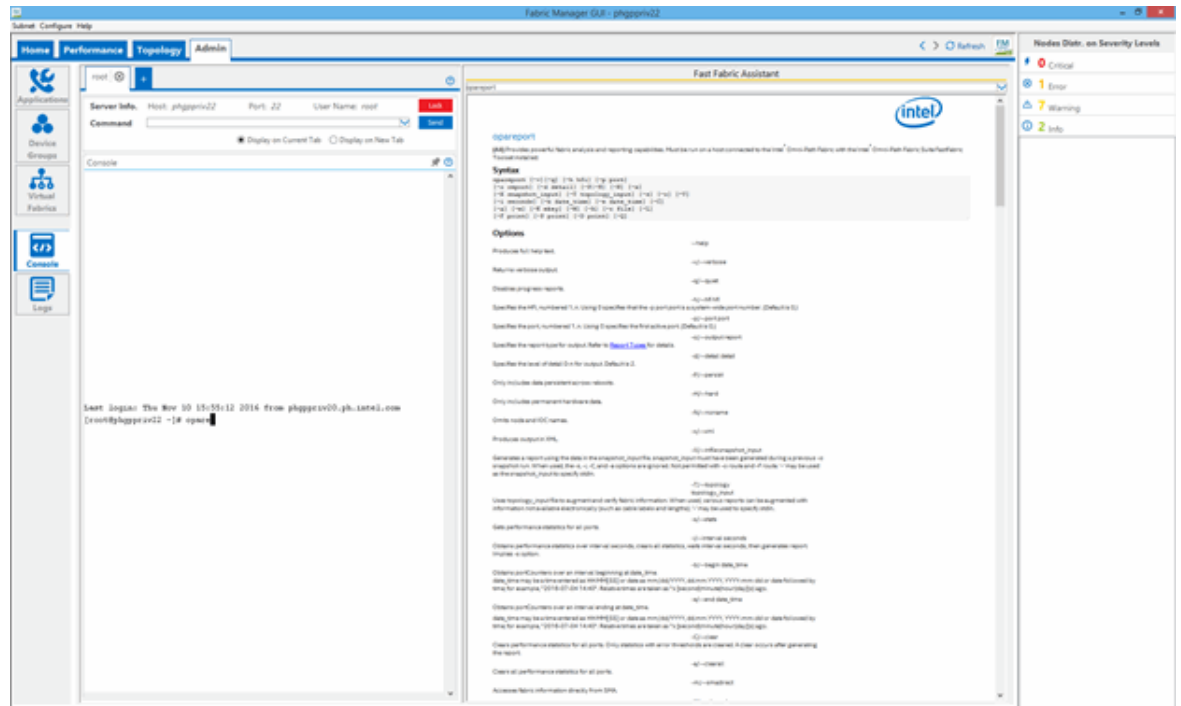
The administrator policies attributes can be used only once and the used attribute is identified with an icon to indicate that it is not available.

7.4 Interactive Console

When you select the Console icon on the control panel (or the "+" tab) Fabric Manager GUI displays a new console screen where you must enter a username, password, host, and port (SSH uses port 22 by default). After logging in, the Interactive Console window changes as shown in the following figure.



Figure 47. Interactive Console



The Interactive Console provides:

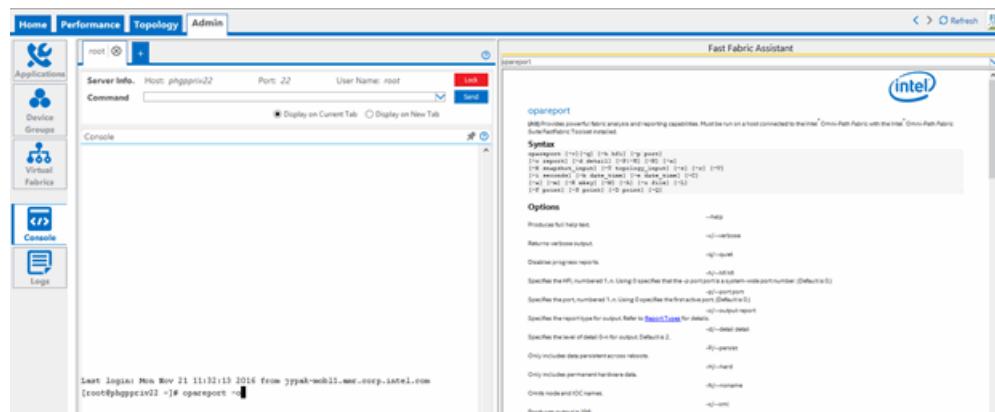
- An Integrated Management Environment to allow a power user to manage a subnet via consoles within the Fabric Manager GUI.
- A Dynamic Online Help for CLI tools. It may be difficult for users to remember the arguments for a CLI tool. This special console monitors your actions and is able to automatically identify and show the corresponding help for the command as you are typing.

For example, when you type `opare`, the console analyses the input, selects the `opareport` command, and then displays the help content for this command with arguments and their meanings. If you continue to type `opareport -o`, the panel scrolls to the available output types for argument `-o`.

The Interactive Console also manages consoles by tabs and each command can be issued either in the same console or a new console. Issuing commands in a new console enables you to compare results from similar commands by switching tabs.

To close a console, click the icon on the tab. To log into a console, click the icon (+).

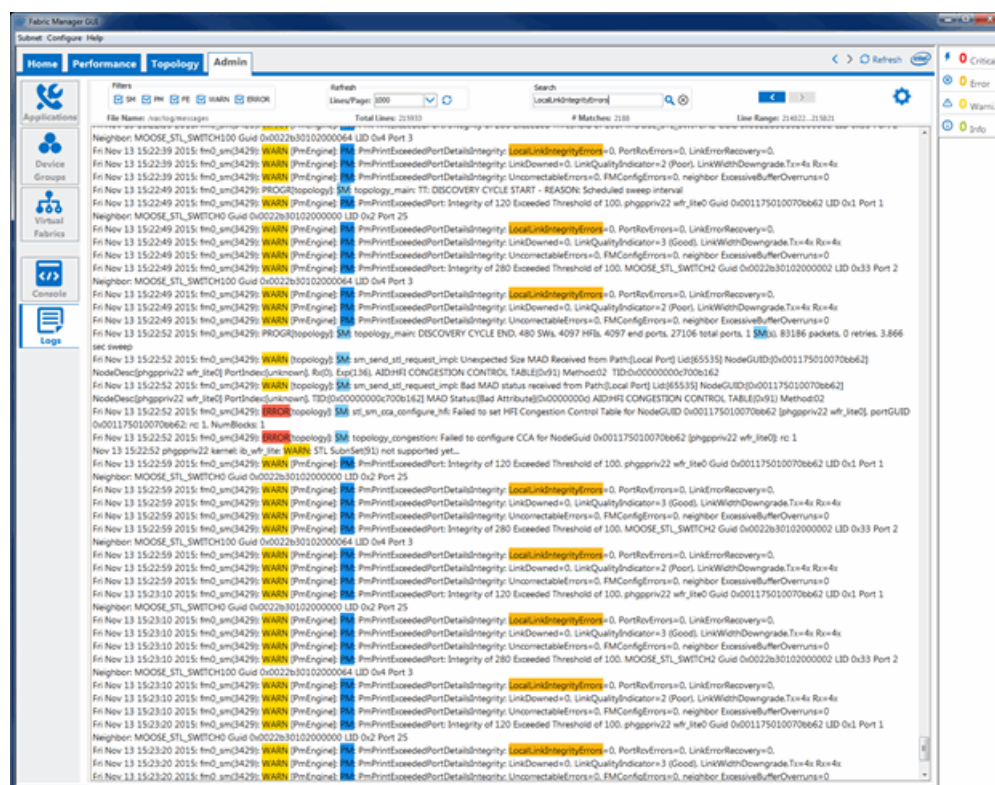
[illegible]




7.5 Log View

The SM Log Viewer shown in the following figure shows the error log from the remote SM host. At the top of the page controls are provided to change the log text shown. A status line displays the file name, number of lines, number of matched search results, and range of lines in the file being viewed. The **File Name** field shows the path of the log file on the remote host. The tooltip to display the entire path is available for file paths that exceed the name field width. The main text portion of the window displays the lines of text taken from the log file (see Lines/Page below).


Figure 49. Log Viewer




Filters

Check boxes  are used to filter lines of text containing SM, PM, FE, WARN, and ERROR. When a checkbox is selected, all lines containing the selected filter are displayed and highlighted in a different color; lines that do not contain the filtered text are not displayed.


Refresh






The Lines/Page control  allows you to select the number of lines to display on each page as the file is traversed. The **Total Lines** field shows the total number of lines in the log file and is updated as lines are added.

Note:

After you enter or select a number of lines, click the  to refresh the viewer with the new lines.

Search


The Search feature  is accessible through the actions described below. The **# Matches** field below it indicates the number of search results found.

- Search Field
 - Click the search field to highlight text.
 - Double-clicking on the search field removes highlights from existing text and places a cursor at the end of the text.
 - Enter text to search for and then click  to perform search.
 - Click  to remove highlights from search results and to clear the search field.
- Marked Search
 - In the main text window, double-click on a word to highlight all instances of the word.
 - Selected word also appears in the search field.
 - Click  to remove highlights from search results and to clear the search field.
- Selected Search
 - Copy/Paste
 1. While holding down the left mouse button, drag the mouse pointer across a word, phrase, line, or multiple lines to select the desired search entry.
 2. Right-click on the text area to display the context menu and select **Copy**.
 3. Right-click in the search field to display the context menu and select **Paste**.
 4. Click  or press **ENTER** to perform search.
 5. Click  to remove highlights from search results and to clear the search field.







Note: You can also use **Ctrl+C** to copy to the clipboard, and **Ctrl+V** to paste from the clipboard.


— Highlighting

1. While holding down the left mouse button, drag the mouse pointer across a word, phrase, line, or multiple lines to select the desired search entry.
2. Right-click on the text area to display the context menu and select **Highlight** to highlight all instances of the selected text.
3. Click  to remove highlights from search results and to clear the search field.

Paging

The Paging control   allows you to move backward and forward through the log file in line increments specified in the Lines/Page field. Highlighting of text imposed by the Filter, Highlight, or Search functions remains in effect as the page view changes. The **Line Range** field, located below the   controls, shows the range of lines currently being viewed. This range remains unchanged if lines are filtered and show a subset of lines designated by Lines/Page or range.

Configuration

The Configuration control  allows you to log out of the Log Viewer and redefine the host credentials and log file name. When the Configuration control is clicked the screen shown in the figure below is displayed. Two modes are available when logging into the Log Viewer:

- Auto Config – Only the user name and password fields are available and it is assumed that the Host, Port, and Log File path information is correct.
- Custom Config – In this mode, all fields are editable and the user must provide the necessary login information.

With this feature, the log on any host can be viewed regardless of the host to which the subnet window is connected.



Figure 50. Configuration Login Screen

Login

☒ Auto Config ☐ Custom Config

Host :

Port # :

User Name :

Password :

Log File :

NOTE: The ESM Syslog cannot be accessed directly from a switch!
See the Intel Omni-Path Fabric Switches CLI Reference Guide to configure the remote Syslog, and the Intel Omni-Path Fabric Suite Fabric Manager User Guide for further information about how to use this screen.



8.0 Troubleshooting Guide

8.1 Overview

This troubleshooting guide is intended to provide assistance to system administrators and users to detect and correct system anomalies brought on by hardware failures or incorrect hardware/software configuration.

Assumptions

It is assumed that the user has some knowledge of High-Performance Computing (HPC) systems and is familiar with the following:

- The Intel® Omni-Path Fabric Suite FastFabric Tool Suite
- Navigation of the Linux command line
- Protocols: Secure Shell, Secure Copy, and Secure FTP (SFTP)
- General administration of a network cluster

About This Guide

This guide provides steps that can be taken to prevent or correct errors encountered when using the Fabric Manager GUI. It is not meant as an exhaustive troubleshooting guide for fabric-related anomalies but provides direction in correcting improper functioning of the Fabric Manager GUI.

Who Should Use This Guide

This guide is intended for system administrators responsible for the configuration and maintenance of HPC systems.

Typographical Conventions

The typographical conventions used throughout the Troubleshooting Guide are shown in the following table.

Table 15. Typographical Conventions

Convention	Description
Normal	Indicates content describing anomalies and how to correct them.
Shaded block	Indicates a command typed at the command line.
Bold	Indicates key Fabric Manager GUI components such as menu options.
Square Brackets []	Optional Items.
Angle Brackets < >	User-supplied values.



8.2 General Error Prevention

The number of errors encountered when using the Fabric Manager GUI can be minimized by adhering to a set of commonly ascribed best practices. While it is possible to encounter errors even when measured precautions have been observed, Intel highly recommends that these practices be used as a first line of defense against commonly known errors.


8.3 Best Practices

By following these best practices, many common errors can be averted.

- Keep an Ssh Terminal Open

When running the Fabric Manager GUI, it is helpful to always have an Ssh Terminal available to verify host connectivity, analyze OPAFM configuration, use Intel® Omni-Path Fabric Suite FastFabric Tools to verify expected behavior, etc. Even if it is not possible to connect to the subnet, an Ssh session can still be opened using a **CONSOLE** feature under the **ADMIN** tab.

- Use the Online Help System

The Fabric Manager GUI has an Online Help System that should be used as a guide to the proper use and understanding of features being explored. Click the  on each panel to obtain specific information about that panel, or select **ONLINE HELP** from the **HELP** menu for information about setting up access to a subnet.

- Use the Fast Fabric Assistant

The **CONSOLE** feature mentioned above has the additional ability of providing real-time help through the FastFabric Assistant, which provides command line syntax for commands in the Intel® Omni-Path Fabric Suite FastFabric Tool Suite.

- Take advantage of the SM Log

The Fabric Manager GUI provides a window into the SM Log file where errors and other activities are logged by the Intel® Omni-Path Fabric Suite Fabric Manager. This can be an invaluable resource in detecting the root cause of errors encountered when running the Fabric Manager GUI. Under the **ADMIN** tab, select the **LOGS** tab and log into the host. Lines of text from the log file are provided beginning with the most recent 100 lines to a maximum of 1000 lines of text. Use the filters to narrow the text being displayed by device type, warnings, or errors.

8.4 Application Log

There are several application log files which, if configured, can provide much guidance in determine the root cause of system failures. These logs are located in different directories depending on whether it is installed under Windows or Linux.

The log files names are:

- fmgui.log
- hibernate.log
- opadbmgr.log
- opafecdriver.log



The paths for each OS are:

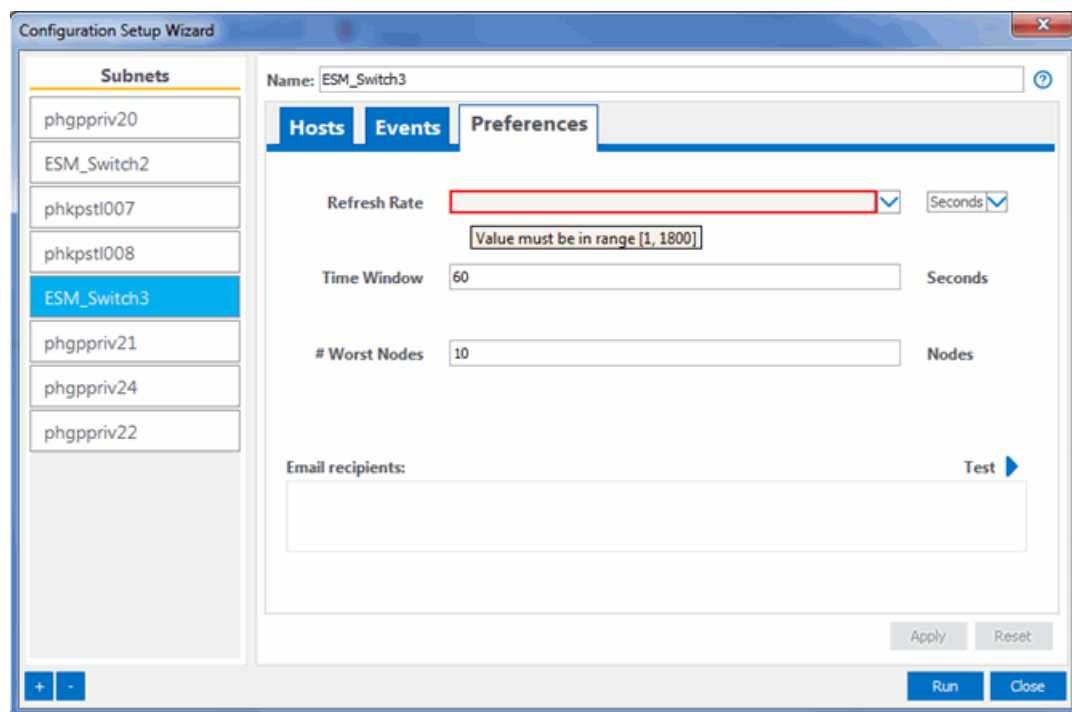
- **Windows:** C:\Users\\AppData\Roaming\Intel\FabricManagerGUI\logs
- **Linux:** \$HOME/.Intel/FabricManagerGUI/logs

Log file names and location are configurable in the logging configuration file at C:\Users\\AppData\Roaming\Intel\FabricManagerGUI\logconfig.xml

8.5 Troubleshooting Q & A

In general, text fields throughout the Fabric Manager GUI provide real-time validation. If the type of information provided is not appropriate or left blank, the field is highlighted in light red with a red border as shown in the following example. Mouse-over the field to display a tool tip about how to provide the correct information.

Figure 51. Troubleshooting Example



Q1. Why does Fabric Manager GUI fail to connect to Fabric Manager?

A1. Diagnose the connection issue using the following steps.

1. Check whether the FE is configured properly. You can use the `opafequery` tool from the IFS package to do this.
This tool is an easy way to diagnose the connectivity to the FE without using Fabric Manager GUI. You can run the tool on the same server where the Fabric Manager and FE are running. If `opafequery` fails, refer to the *Intel® Omni-Path*



Fabric Suite Fabric Manager User Guide in the section "Controlling FM Startup" to configure and start FE properly. Ensure IPV6 is enabled on the FM/FE host because it is required by FE.

2. Check that the network is configured properly.

If the opafequery test in Step 1 works properly, there are issues with the network connectivity.

- a. Check the firewall to ensure that the FE port is reachable.
- b. If Step a. passes, the issue may be more complicated.
 - i. If you are using both ipv4 and ipv6, try to focus on ipv4 only.
 - ii. If you are using only ipv6, you might try looking at the Java system property "java.net.preferIPv4Stack" and see if it is set to "true". Check on the machine running the FM GUI first, look at JAVA_TOOL_OPTIONS or _JAVA_OPTIONS settings. If not, perhaps your organization is enforcing this policy and you might try working with your system or network administrator to reach a resolution.
3. If you have verified that the Fabric Manager GUI machine can communicate with the FE port, then you should send the fmgui log file to Intel Fabric Manager GUI support.

Q2. Why is the number of active nodes/ports wrong on the Home screen?

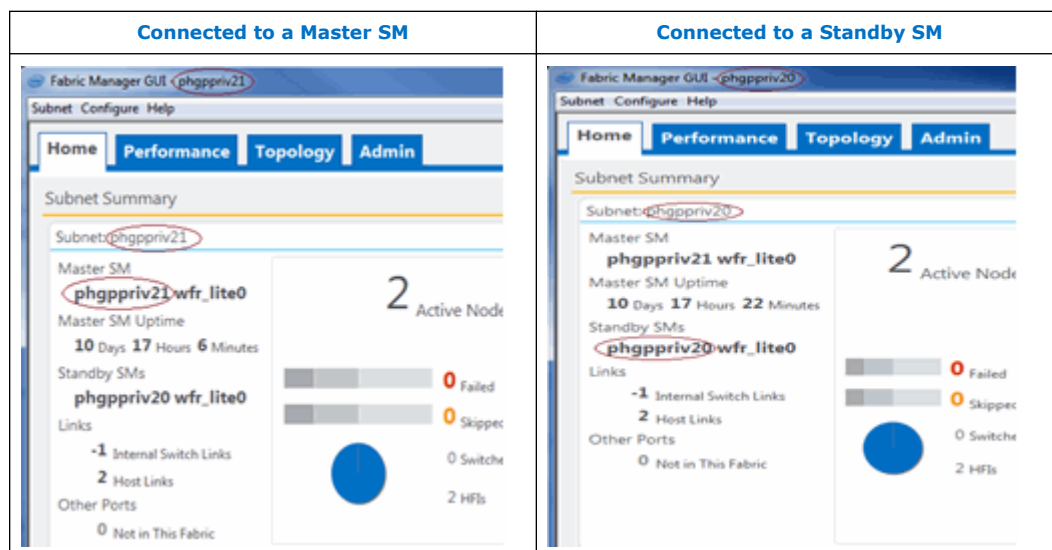
A2. It is possible that the Fabric Manager GUI is configured to connect to the wrong host. Double-check the hostname/IP address to confirm the connection to the expected host.

Q3. On the **HOME** page, why does the **MASTER SM** field show the name of a different host than the one I am connected to?

A3. On the **HOME** page, notice that the name of the host you are connected to is shown in a number of places:


- In the title bar at the top of the screen.
- In the Subnet: field at the top of the **SUBNET SUMMARY** section.
- Either in the **MASTER SM** or **STANDBY SMs** fields.

If you have connected to a Standby SM, the name in the **MASTER SM** field will be different from the name in the title bar and **Subnet:** field.

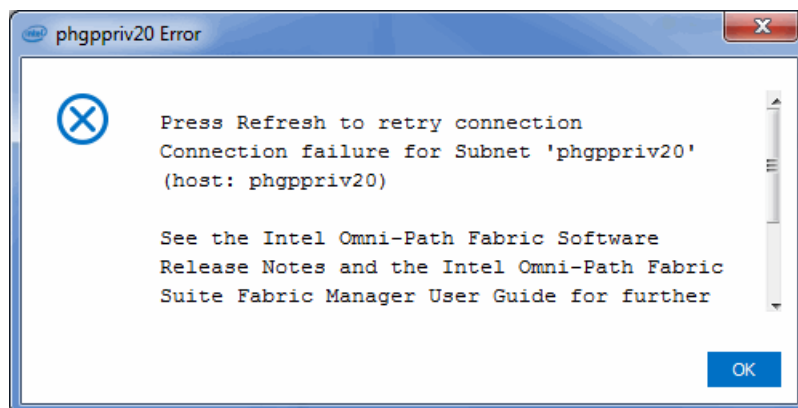


Q4. The Fabric Manager GUI fails to connect to a subnet after a long time and then displays an error message. How can I fix this?

A4. There are many reasons why a remote subnet fails to connect. The Fabric Manager GUI repeatedly attempts to connect to a subnet through the *failover* process. By traversing the list of subnet hosts, provided in the Subnet Wizard, the Fabric Manager GUI attempts to obtain a response from the Subnet Manager (SM) and subsequently the Performance Manager (PM). During the connection process the following dialog is displayed:



If no response is received from the SM, or if the SM responds but the PM does not, then the failover process has failed and the following dialog is displayed:



If this happens, then it is necessary to determine the root cause of the problem. Here are a few things to look at:

1. Check whether the host is reachable.

Open an Ssh session using the **CONSOLE** feature under the **ADMIN** tab and then type the following command:

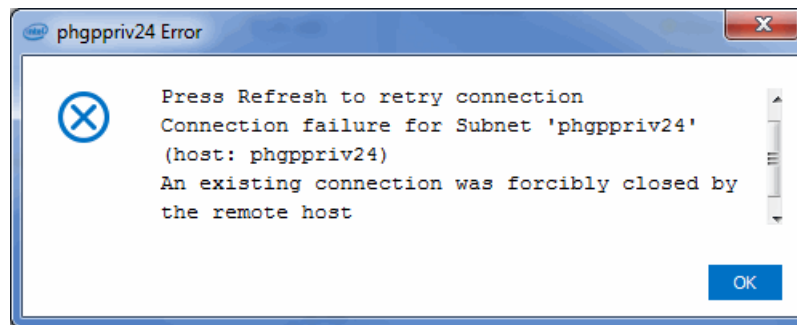
```
ping <ip/hostname>
```

If there is no response, determine the root cause and try again.


2. Check whether the FE is running.

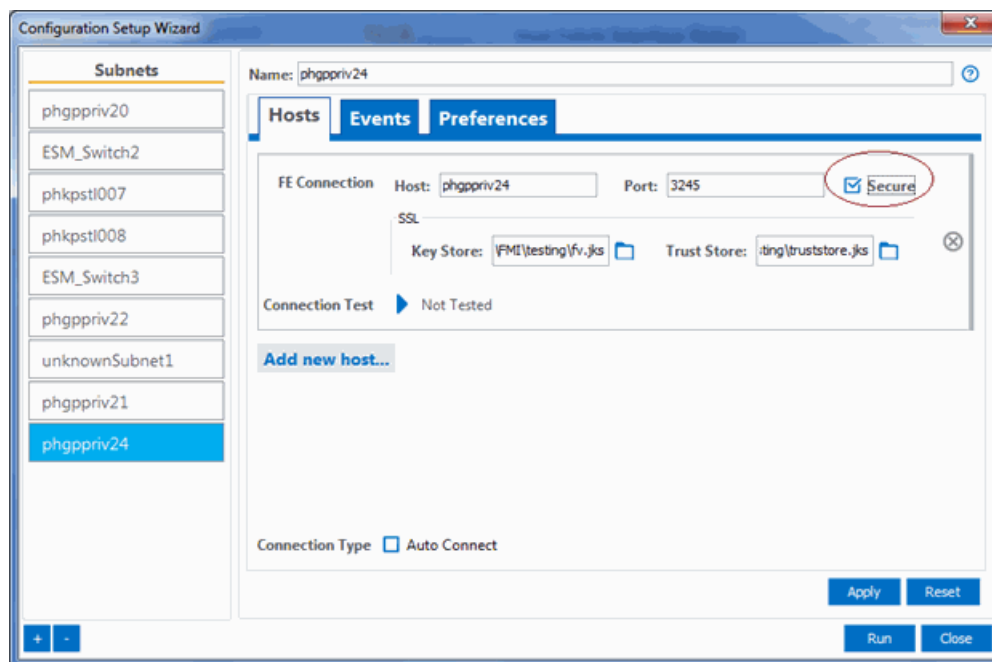
For more information about configuring the Fabric Manager to initialize the FE, refer to the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* in the section titled "Controlling FM Startup."

Q5. When I try to connect to a subnet I get an error indicating that the connection was forcibly closed by the remote host. What can I do?




A5. This error is displayed when an attempt is made to connect to a secure fabric without providing the necessary security certificates to establish the connection.

1. Follow the procedure in "SSL Key Creation for Fabric Manager GUI" appendix in the *Intel Omni-Path Fabric Suite Fabric Manager User Guide*.
2. Open the **SETUP WIZARD** under the **CONFIGURE** menu. Select the subnet to be configured and click on the **SECURE** checkbox as shown in the following figure. Use the file browser  icons to locate the Key and Trust Store files, click **APPLY**, and then **RUN**.



3. Enter the password for the certificates on the new dialog to establish a connection to the remote host.



Q6. While trying to open an Ssh session to a remote host using the **CONSOLE** feature under the **ADMIN** tab, the following error was reported. Why?

```
SSH_MSG_DISCONNECT: 2 Too many authentication failures for invaliduser
```

A6. This error is displayed when:

- Invalid credentials are provided on the console login panel; check the credentials and try again.
- Invalid or valid credentials are provided, but the host is unreachable.

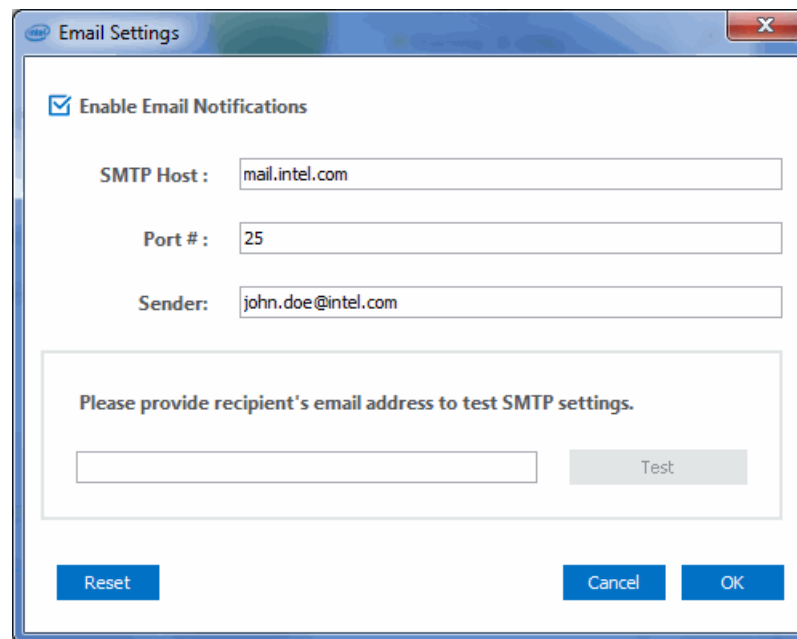
Q7. While trying to open an Ssh session to a remote host using the **CONSOLE** feature under the **ADMIN** tab, the following error was reported:

```
Unable to connect to remote host phgppriv22: java.net.ConnectException:
Connection refused: connect
```


A7. This error is displayed when the Ssh port is incorrect. By default the Ssh port number is 22 unless specifically reconfigured by the system administrator.


Q8. In the Preferences Wizard I put an email recipient in the list and clicked on the Test button, but no email was sent.

A8. To send emails, the SMTP server must be running on the remote host. Then the Fabric Manager GUI must be configured to point to the SMTP server. Under the configure menu, select **Email Setup** and enter the SMTP Host, Port #, and Sender's email address.



Q9. On the **LOGGING** window under the **CONFIGURATION** menu, how can I get more information about filling out the **OUTPUT FORMAT** field?

A9. Click on the Help  icon to the right of the **OUTPUT FORMAT** field to display the formatting help window, which provides detailed information about the formatting symbols.



App. Log Level:

Output Format: [Formatting Help](#)

Maximum File Size: MB

Maximum # Files:

File Location:

Preview

To preview a sample message with formatting, enter a format string below. The format string entered will be applied to message "This is a sample log message".

Enter Format String:

Sample Formatted Output

Conversion Pattern

Conversion Character	Effect
%c	Used to output the category of the logging event. The category conversion specifier can be optionally followed by precision specifier, that is a decimal constant in brackets. If a precision specifier is given, then only the corresponding number of right most components of the category name will be printed. By default the category name is printed in full. For example, for the category name "a.b.c" the pattern %c(2) will output "b.c".
%C	Used to output the fully qualified class name of the caller issuing the logging request. This conversion specifier can be optionally followed by precision specifier, that is a decimal constant in brackets. If a precision specifier is given, then only the corresponding number of right most components of the class name will be printed. By default the class name is output in fully qualified form. For example, for the class name "org.apache.xyz.SomeClass", the pattern %C(1) will output "SomeClass"
%d	Used to output the date of the logging event. The date conversion specifier may be followed by a date format specifier enclosed between braces. For example, %d(%H:mm:ss,SSS) or %d(%d %MM %YY %H:mm:ss,SSS). If no date format specifier is given then ISO8601 format is assumed.
%F	Used to output the file name where the logging request was issued.
%l	Used to output location information of the caller which generated the logging event. The location information depends on the JVM implementation but usually consists of the fully qualified name of the calling method followed by the callers source file name and line number between parentheses.
%L	Used to output the line number from where the logging request was issued.
%m	Used to output the application supplied message associated with the logging event.
%M	Used to output the method name where the logging request was issued.
%n	Outputs the platform dependent line separator character or characters.
%p	Used to output the priority of the logging event.
%r	Used to output the number of milliseconds elapsed from the construction of the layout until the creation of the logging event.
%t	Used to output the name of the thread that generated the logging event.
%x	Used to output the NDC (nested diagnostic context) associated with the thread that generated the logging event.
%%	The sequence %% outputs a single percent sign.



Appendix A Port Counters Overview

Each port in an Intel® Omni-Path Fabric maintains a set of port counters to indicate both traffic and error counts. These counters can be grouped into the categories described in this section. Each port stops incrementing when the max value is reached, irrespective of counter size. Most of the counters are 64-bits in size. Exceptions are noted.

A.1 Utilization

These counters reflect the normal utilization of the port and Virtual Lane when present.

Several of these counters are used during the calculation of Congestion, SMA Congestion, and the Bubble Categories. The Utilization metrics provide a way of giving some of the other counters context by comparing them to the amount of data or packets that were transmitted or received.

A.1.1 PortXmitData (TxD) and PortVLXmitData[n]

These counters indicate the total number of fabric packet flits transmitted. This does not include idle nor other LF command flits.

A.1.2 PortRcvData (RxD) and PortVLRcvData[n]

These counters indicate the total number of fabric packet flits received.

A.1.3 PortMulticastXmitPkts (MTxP)

This counter indicates the number of multicast and collective packets transmitted.

A.1.4 PortMulticastRcvPkts (MRxP)

This counter indicates the number of multicast and collective packets received.

A.2 Link Integrity

These counters reflect errors in the Physical (PHY) and Link Layers, as well as errors in firmware. In some cases, these errors are benign and can be ignored. However in other cases, excessive link integrity errors can indicate a hardware problem such as a poor connection, marginal cable, incorrect length/model cable for signal rate, or damaged/broken hardware, such as bad connectors.

When a bad packet is detected, one of these counters is incremented and the Link Layer may either discard or replay the packet.



During the link training sequence, assorted errors may be observed. This is a normal part of the link training and clock synchronization process. Hence, errors observed as part of rebooting nodes or moving cables should not be considered a problem.

The category is calculated as a weighted sum of the counters in the group. With the exception of ExcessiveBufferOverflowErrors, the counters in this group report on the receive side of the link. However, the counter can indicate a problem on either side of the link.

A.2.1 Link Quality Indicator (LQI)

This is a status indicator, similar to the signal strength bar display on a mobile phone, that enumerates link quality as a range of 0-5, with 5 being very good. Values in the lower part of the range may indicate hardware problems with components such as ports and cables that surface as signal integrity issues, leading to performance and other problems. The LQI gives you an instantaneous view of a link's quality on every hardware port.

Table 16. Link Quality Values and Description

Link Quality Value	Description
5	Working at or above preferred link quality, no action needed.
3	Working on low end of acceptable link quality, recommended corrective action on next maintenance window.
2	Working below acceptable link quality, recommend timely corrective action.
1	Working far below acceptable link quality, recommend immediate corrective action.
0	Link down

A.2.2 LocalLinkIntegrityErrors (LLI) Counter

This counter indicates the number of retries initiated by a link transfer layer receiver.

The retry rate is represented by the Link Quality Indicator. A link that is meeting performance requirements has a Link Quality of 5, which corresponds to 1000 or fewer replays per second.

A.2.3 PortRcvErrors (RxE) Counter

This counter indicates the total number of packets containing an error that were received by the port, including Link Layer protocol violations and malformed packets. It indicates possible misconfiguration of a port, either by the SM or by user intervention. It can also indicate hardware issues or extremely poor link signal integrity.

A.2.4 ExcessiveBufferOverflowErrors (EBO) Counter

This counter, associated with credit management, indicates an input buffer overrun. It indicates possible misconfiguration of a port, either by the SM or by user intervention. It can also indicate hardware issues or extremely poor link signal integrity.



A.2.5 LinkErrorRecovery (LER) Counter

This counter indicates the number of times the link has successfully completed the link error recovery process.

Link Quality Indicator is the primary indicator for link quality to use. This counter is factored into the value reported for Link Quality Indicator. This counter may be non-zero for a properly functioning link.

A.2.6 LinkDowned (LD) Counter

This counter indicates the total number of times the port has failed the link error recovery process and downed the link. These events can cause disruptions to fabric traffic.

A.2.7 UncorrectableErrors (Unc) Counter

This counter indicates the number of unrecoverable device errors. This may indicate a defect in the reporting device.

A.2.8 FMConfigErrors Counter (FMC)

This counter reports inconsistent configurations of the low-level Subnet Management Agent (SMA) on either side of the link. It indicates possible misconfiguration of a port, either by the SM or by user intervention.

A.3 Congestion

These counters reflect possible errors that indicate traffic congestion in the fabric.

When congestion or a packet that has seen congestion is detected, one of these counters is incremented and then depending on the issue reported, the packet must wait. In an extreme case, the packet may time out and be dropped.

The category is calculated as a weighted sum of the counters in the context of the utilization counters. With the exception of PortRcvFECN, the counters are all reported on the transmit side of the link. In addition, PortRcvBECN is only taken if the local node is an HFI. However, the counter could indicate a problem on either side of the link.

A.3.1 CongDiscards (CD) Counter

Note: Formerly known as "SwPortCongestion".

This switch-only counter indicates the number of packets that were discarded as unable to transmit due to timeouts.

A.3.2 PortRcvFECN (RxF) Counter

When a device receives a packet with the Forward Explicit Congestion Notification (FECN) bit set to one, this counter is incremented.



A.3.3 PortRcvBECN (RxB) Counter

When a device receives a packet with the Backward Explicit Congestion Notification (BECN) bit set to one, this counter is incremented.

A.3.4 PortMarkFECN (MkF) Counter

This counter indicates the total number of packets that were marked Forward Explicit Congestion Notification (FECN) by the transmitter due to congestion.

A.3.5 PortXmitTimeCong (TxTC) Counter

This counter indicates the total number of *flit times* that the port was in a congested state for any data VL.

A.3.6 PortXmitWait (TxW) Counter

This counter indicates the amount of time (in *flit times*) any virtual lane had data but was unable to transmit due to no credits available.

A.4 SMA Congestion

These counters reflect congestion in the fabric specific to communication between the Subnet Manager and Subnet Manager Agents using the management VL (VL 15).

The category is calculated exactly as the Congestion category using the same weights and the correct VL15 utilization counters.

A.4.1 PortVLXmitWait[15] (VLTxW[15]) Counter

This counter behaves the same as PortXmitWait, but it is restricted to VL 15, which carries only SM traffic.

A.4.2 VLCongDiscards[15] (VLCD[15]) Counter

Note: Formerly known as "SwPortVLCongestion".

This counter behaves the same as CongDiscards, but it is restricted to VL 15, which carries only SM traffic.

A.4.3 PortVLRcvFECN[15] (VLRxF[15]) Counter

This counter behaves the same as PortRcvFECN, but it is restricted to VL 15, which carries only SM traffic.

A.4.4 PortVLRcvBECN[15] (VLRxB[15]) Counter

This counter behaves the same as PortRcvBECN, but it is restricted to VL 15, which carries only SM traffic.



A.4.5 PortVLXmitTimeCong[15] (VLTxTC[15]) Counter

This counter behaves the same as PortXmitTimeCong, but it is restricted to VL 15, which carries only SM traffic.

A.4.6 PortVLMarkFECN[15] (VLMkF[15]) Counter

This counter behaves the same as PortMarkFECN, but it is restricted to VL 15, which carries only SM traffic.

A.5 Bubble

These counters occur when an unexpected idle flit is transmitted or received.

The transmit port sends idle flits until it can continue sending the rest of the packet. The category is calculated as follows:

1. The maximum value between the sum of the XmitWastedBW and XmitWaitData or the neighbor's PortRcvBubble.
2. Then divide the previous value by the port's utilization to provide context.

A.5.1 PortXmitWastedBW (WBW) Counter

This counter indicates the number of *flit times* where one or more packets have been started but the transmitters are forced to send idles due to bubbles in the ingress stream. Also, the VLs that have data to be sent are not permitted to preempt the currently transmitting VL.

A.5.2 PortXmitWaitData (TxWD) Counter

This counter indicates the number of *flit times* where one or more packets have been started but interrupted due to bubbles in the ingress stream.

A.5.3 PortRcvBubble (RxBb) Counter

This counter indicates the total number of *flit times* where one or more packets have started to be received, but the receiver received idle flits from the wire.

A.6 Security

These counters reflect possible security problems in the fabric.

Security problems can occur if a PKey or SLID violation occurs at the port during the ingress or egress of a packet.

The category is calculated as the sum of the neighbor's PortRcvConstraintErrors and the local port's PortXmitConstraintErrors.

A.6.1 PortRcvConstraintErrors (RxCE)

This counter is incremented when partition key or source LID violations are detected in a received packet, indicating a possible security issue or misconfiguration of device security settings.



A.6.2 PortXmitConstraintErrors (TxCE)

This counter is incremented when partition key violations are detected in a packet attempting to be transmitted, indicating a possible security issue or misconfiguration of device security settings.

A.7 Routing

These counters reflect possible routing issues. When a routing issue occurs, the offending packet is dropped.

A typical cause of this error is the routing to a wrong egress port or an improper Service Channel (SC) mapping. These errors can be a side effect of a port or device going down while traffic was still in flight to or through the given port or device.

A.7.1 PortRcvSwitchRelayErrors (RxSR)

This counter indicates the number of packets that were dropped due to internal routing errors. It indicates possible misconfiguration of a switch by the SM.

A.8 Other

These counters do not fit into any of the previous categories.

A.8.1 PortRcvRemotePhysicalErrors (RxRP)

This counter indicates the number of downstream effects of signal integrity (SI) problems. It indicates an SI issue in the upstream path.

This counter was not included as it does not directly indicate the link that had the issue, so it can be misleading.

A.8.2 PortXmitDiscards (TxDc)

This counter indicates the number of packets dropped due to several reasons including timeouts and improper packet lengths.

Note: This counter is a super set that includes Congestion Discards counter.