

# **Intel® Omni-Path Fabric Suite FastFabric**

**User Guide**

---

***Rev. 5.0***

***December 2016***



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. Learn more at <http://www.intel.com/> or from the OEM or retailer.

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 © 2015–2016, Intel Corporation. All rights reserved.



## Revision History

---

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

Date	Revision	Description
December 2016	5.0	Updates to this document include: <ul style="list-style-type: none"><li>• Document has been restructured and rewritten for usability.</li><li>• Added <a href="#">Cluster Configurator for Intel® Omni-Path Fabric</a> to Preface.</li><li>• Globally, updated filepath from <code>/opt/opa</code> to <code>/usr/lib/opa</code>.</li><li>• Added <a href="#">Intel® Omni-Path Architecture Overview</a> to Overview.</li><li>• Added note to <a href="#">Starting Up the Tools</a> that you must have root privilege to run FastFabric commands.</li></ul>
August 2016	4.0	Document has been updated.
May 2016	3.0	Document has been updated.
February 2016	2.0	Document has been updated.
November 2015	1.0	Document has been updated.



## Contents

---

<b>Revision History.....</b>	<b>3</b>
<b>Preface.....</b>	<b>9</b>
Intended Audience.....	9
Documentation Set.....	9
Cluster Configurator for Intel® Omni-Path Fabric.....	10
Documentation Conventions.....	10
License Agreements.....	11
Technical Support.....	11
<b>1.0 Introduction.....</b>	<b>12</b>
1.1 Document Organization.....	12
<b>2.0 Overview.....</b>	<b>13</b>
2.1 Intel® Omni-Path Architecture Overview.....	13
2.2 FastFabric Overview.....	16
2.2.1 FastFabric Architecture.....	16
2.2.2 FastFabric Capabilities.....	18
<b>3.0 Getting Started.....</b>	<b>24</b>
3.1 Important Note on First-Time Installations.....	24
3.2 Starting Up the Tools.....	24
3.2.1 Accessing the Intel FastFabric OPA Tools Menu.....	24
3.2.2 Accessing the Fabric Performance Monitor.....	25
3.3 Intel FastFabric OPA Tools TUI Overview.....	27
3.4 How to Use the FastFabric TUI.....	27
3.5 Fabric Performance Monitor TUI Overview.....	29
3.6 How to Use the Fabric Performance Monitor TUI.....	30
3.7 Configuration of IPoIB Name Mapping.....	33
3.8 Configuration Files for FastFabric .....	34
3.8.1 FastFabric Configuration File.....	35
3.8.2 Ports List Configuration File.....	35
3.8.3 Chassis List Configuration Files.....	36
3.8.4 Externally-Managed Switch List Configuration File.....	38
3.8.5 Hosts List Configuration Files.....	40
3.8.6 Port Statistics Thresholds Configuration File.....	41
3.8.7 Signal Integrity Thresholds Configuration File.....	41
3.8.8 Fabric Topology Input File.....	42
<b>4.0 Managing the Chassis Configuration.....</b>	<b>44</b>
4.1 Editing the Configuration Files for Chassis Setup.....	46
4.2 Verifying Chassis via Ethernet Ping.....	48
4.3 Updating the Chassis Firmware.....	49
4.4 Setting Up Chassis Basic Configuration.....	49
4.5 Setting Up Password-less ssh/scp.....	51
4.6 Rebooting the Chassis.....	51
4.7 Getting Basic Chassis Configuration.....	52
4.8 Configuring Chassis Fabric Manager.....	52
4.9 Updating the Chassis FM Security Files.....	56



4.10 Getting Chassis FM Security Files.....	57
4.11 Checking the OPA Fabric Status.....	57
4.12 Controlling Chassis Fabric Manager.....	58
4.13 Generating All Chassis Problem Report Information.....	59
4.14 Running a Command on All Chassis.....	60
4.15 Viewing opachassisadmin Result Files.....	60
<b>5.0 Managing the Switch Configuration.....</b>	<b>62</b>
5.1 Editing the Configuration Files for Externally-Managed Switch Setup.....	64
5.2 Generating or Updating Switch File.....	66
5.3 Testing for Switch Presence.....	66
5.4 Verifying Switch Firmware.....	67
5.5 Updating Switch Firmware.....	67
5.6 Setting Up Switch Basic Configuration.....	69
5.7 Rebooting the Switch.....	70
5.8 Reporting Switch Firmware and Hardware Information.....	70
5.9 Getting Basic Switch Configuration.....	71
5.10 Reporting Switch VPD Information.....	72
5.11 Viewing opaswitchadmin Result Files.....	72
<b>6.0 Managing the Host Configuration.....</b>	<b>74</b>
6.1 Editing the Configuration Files for Host Setup.....	76
6.2 Verifying Hosts are Pingable.....	77
6.3 Setting Up Password-Less SSH/SCP.....	78
6.4 Copying /etc/hosts to All Hosts.....	78
6.5 Showing uname -a for All Hosts.....	79
6.6 Installing/Upgrading OPA Software.....	80
6.7 Configuring IPoIB IP Address.....	81
6.8 Building Test Applications and Copying to Hosts.....	81
6.9 Rebooting Hosts.....	82
6.10 Refreshing SSH Known Hosts.....	82
6.11 Rebuilding MPI Library and Tools.....	83
6.12 Running a Command on All Hosts.....	85
6.13 Copying a File to All Hosts.....	86
6.14 Viewing opahostadmin Result Files.....	87
<b>7.0 Verifying the Host.....</b>	<b>88</b>
7.1 Editing the Configuration Files for Host Verification.....	90
7.2 Viewing a Summary of Fabric Components.....	92
7.3 Verifying Hosts Pingable, SSHable, and Active.....	92
7.4 Performing Single Host Verification.....	93
7.5 Verifying OPA Fabric Status and Topology.....	95
7.6 Verifying Hosts See Each Other.....	96
7.7 Verifying Hosts Ping via IPoIB.....	96
7.8 Refreshing SSH Known Hosts.....	97
7.9 Checking MPI Performance.....	98
7.10 Checking Overall Fabric Health.....	99
7.11 Starting or Stopping Bit Error Rate Cable Test.....	100
7.12 Generating All Hosts Problem Report Information.....	101
7.13 Running a Command on All Hosts.....	103
7.14 Viewing opahostadmin Result Files.....	104



<b>8.0 Monitoring Fabric Performance.....</b>	<b>106</b>
8.1 Viewing the Fabric Performance Monitoring Summary Screen.....	106
8.2 Viewing the PM Configuration.....	108
8.3 Viewing Image Information.....	110
8.4 Viewing Bandwidth Statistics.....	111
8.5 Viewing Error Statistics.....	114
8.6 Viewing Configuration Information.....	117
8.7 Viewing Focus Information.....	119
8.8 Viewing Port Statistics.....	122
8.9 Navigating PM Sweeps.....	125
8.10 Bookmarking a Sweep.....	126
8.11 Using the opatop Command Line Options.....	127



## Figures

1	Intel® OPA Building Blocks.....	13
2	Intel® OPA Fabric.....	14
3	Intel® OPA Fabric and Software Components.....	15
4	FastFabric Architecture.....	17
5	Fabric Performance Monitor TUI Screen (Example).....	30
6	Fabric Performance Monitoring TUI Navigation.....	32



## Tables

1	FastFabric Methods.....	18
2	FastFabric OPA Fabric Monitoring Menu Descriptions.....	26
3	Fabric Performance Monitor TUI Descriptions.....	30
4	FastFabric OPA Chassis Setup/Admin Menu Descriptions.....	45
5	FastFabric OPA Switch Setup/Admin Menu Descriptions.....	63
6	FastFabric OPA Host Setup Menu Descriptions.....	75
7	FastFabric OPA Host Verification/Admin Menu Descriptions.....	89
8	Performance Impact.....	99
9	Summary Screen Field Descriptions.....	107
10	PM Configuration Field Descriptions.....	109
11	Image Information Field Descriptions.....	111
12	Bandwidth Statistics Field Descriptions.....	113
13	Bandwidth Statistics Field Descriptions.....	116
14	Configuration Information Field Descriptions.....	118
15	Focus Information Field Descriptions.....	121
16	Port Statistics Field Descriptions.....	123





## 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.

## Documentation Set

The complete end user publications set for the Intel® Omni-Path product includes the following items.

- Hardware Documents:
  - *Intel® Omni-Path Fabric Switches Hardware Installation Guide*
  - *Intel® Omni-Path Fabric Switches GUI User Guide*
  - *Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide*
  - *Intel® Omni-Path Edge Switch Platform Configuration Reference Guide*
  - *Intel® Omni-Path Fabric Managed Switches Release Notes*
  - *Intel® Omni-Path Fabric Externally-Managed Switches Release Notes*
  - *Intel® Omni-Path Host Fabric Interface Installation Guide*
- Software Documents:
  - *Intel® Omni-Path Fabric Software Installation Guide*
  - *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*
  - *Intel® Omni-Path Fabric Suite FastFabric User Guide*
  - *Intel® Omni-Path Fabric Host Software User Guide*
  - *Intel® Omni-Path Fabric Suite Fabric Manager GUI Online Help*
  - *Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide*



- *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*
- *Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide*
- *Intel® Omni-Path Fabric Performance Tuning User Guide*
- *Intel® Omni-Path Host Fabric Interface Platform Configuration Reference Guide*
- *Intel® Omni-Path Fabric Software Release Notes*
- *Intel® Omni-Path Fabric Manager GUI Release Notes*
- *Intel® Omni-Path Storage Router Design Guide*
- *Building Lustre\* Servers with Intel® Omni-Path Architecture Application Note*
- *Intel® Omni-Path Fabric Staging Guide*

Documents are available at the following URLs:

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

## 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 11 for more information.  
For more information, visit [www.intel.com](http://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 [www.intel.com](http://www.intel.com) for additional detail.



## 1.0 Introduction

---

This manual provides instructions for using the Intel® Omni-Path Fabric Suite FastFabric, a 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. All of the functions that the TUI menus perform can also be performed using CLI commands. This manual focuses on using the TUI menus. Information on the CLI commands can be found in *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*. To aid in learning the commands, the TUI shows each CLI command as it executes it.

Throughout this document and the FastFabric Tools, "chassis" refers to internally-managed switches and "switches" refers to externally-managed switches.

*Note:* This manual assumes that you have already installed the Intel® Omni-Path Software as prescribed in the *Intel® Omni-Path Fabric Software Installation Guide*

### 1.1 Document Organization

This manual is organized as follows:

- This **Introduction** provides an overview of this document and its structure.
- **FastFabric Overview** provides an overview of the FastFabric architecture and capabilities.
- **Getting Started** provides instructions and information for starting up and using the FastFabric TUI.
- **Managing the Chassis Configuration** provides instructions for setting up and managing the internally-managed switches or chassis.
- **Managing the Switch Configuration** provides instructions for setting up and managing the externally-managed switches.
- **Managing the Host Configuration** provides instructions for setting up and installing the fabric software on all hosts.
- **Verifying the Host** provides instructions for verifying hosts and the fabric.
- **Monitoring Fabric Performance** provides instructions for monitoring the performance, congestion, and error information of a fabric.




## 2.0 Overview

This section provides an overview of the Intel® Omni-Path Architecture and Intel® Omni-Path Fabric Suite FastFabric.

### 2.1 Intel® Omni-Path Architecture Overview

Intel® Omni-Path Architecture (Intel® OPA) 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. These building blocks are shown in the following figure.


**Figure 1. Intel® OPA Building Blocks**



For software applications, Intel® OPA maintains consistency and compatibility with existing Intel® True Scale Fabric and InfiniBand\* APIs utilizing the open source OpenFabrics Alliance\* (OFA) software stack on Linux\* distribution releases.

The following figure shows a sample Intel® OPA-based fabric, consisting of different types of nodes and servers.

**Figure 2. Intel® OPA Fabric**

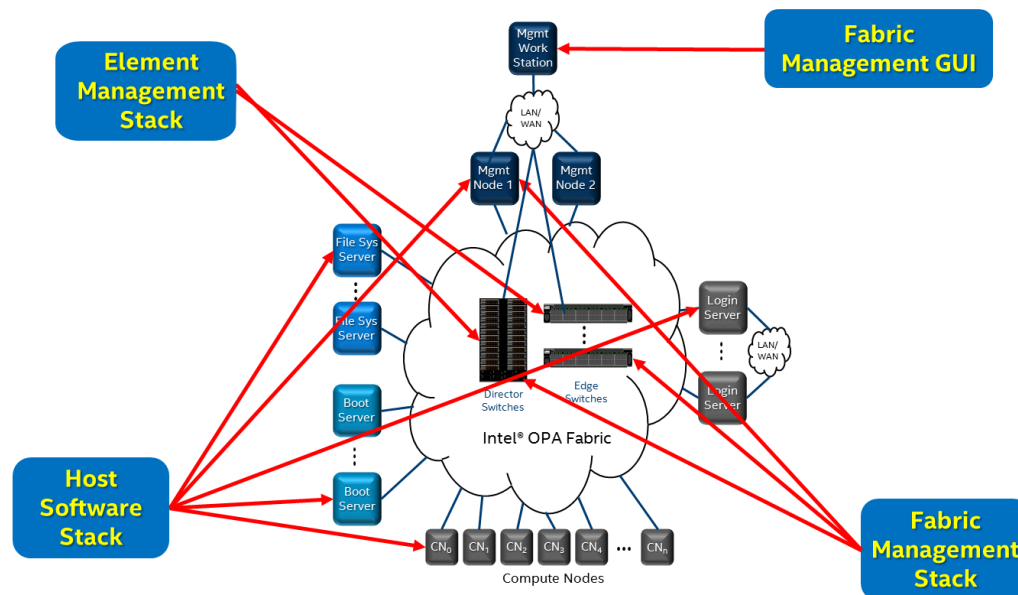


### Software Components

The key software components and their usage models are shown in the following figure and described in the following paragraphs.



Figure 3. Intel® OPA Fabric and Software Components



## Software Component Descriptions

**Element Management Stack**

- Runs on an embedded Intel processor included in managed Intel® OP Edge Switch 100 Series and Intel® Omni-Path Director Class Switch 100 Series switches.
- Provides system management capabilities, including signal integrity, thermal monitoring, and voltage monitoring, among others.
- Accessed via Ethernet\* port using command line interface (CLI) or graphical user interface (GUI).

User documents:

- *Intel® Omni-Path Fabric Switches GUI User Guide*
- *Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide*

**Host Software Stack**

- Runs on all Intel® OPA-connected host nodes and supports compute, management, and I/O nodes.
- Provides high performance, highly scalable MPI implementation via PSM2 and extensive set of upper layer protocols.
- Includes Boot over Fabric mechanism for configuring a server to boot over Intel® Omni-Path using the Intel® OP HFI Unified Extensible Firmware Interface (UEFI) firmware.

User documents:

- *Intel® Omni-Path Fabric Host Software User Guide*

*continued...*



Software Component Descriptions
<ul style="list-style-type: none"><li>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</li></ul>
<b>Fabric Management Stack</b> <ul style="list-style-type: none"><li>Runs on Intel® OPA-connected management nodes or embedded Intel processor on the switch.</li><li>Initializes, configures, and monitors the fabric routing, QoS, security, and performance.</li><li>Includes a toolkit for configuration, monitoring, diagnostics, and repair.</li></ul> User documents: <ul style="list-style-type: none"><li>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</li><li>Intel® Omni-Path Fabric Suite FastFabric User Guide</li><li>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</li></ul>
<b>Fabric Management GUI</b> <ul style="list-style-type: none"><li>Runs on workstation with a local screen and keyboard.</li><li>Provides interactive GUI access to Fabric Management features such as configuration, monitoring, diagnostics, and element management drill down.</li></ul> User documents: <ul style="list-style-type: none"><li>Intel® Omni-Path Fabric Suite Fabric Manager GUI Online Help</li><li>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</li></ul>

## 2.2 FastFabric Overview

Intel® Omni-Path Fabric Suite FastFabric is a set of fabric management tools designed to simplify and optimize common fabric management tasks. FastFabric includes the following capabilities:

- Monitoring and diagnostic tools
- Fabric deployment and verification
- Switch management
- Host management

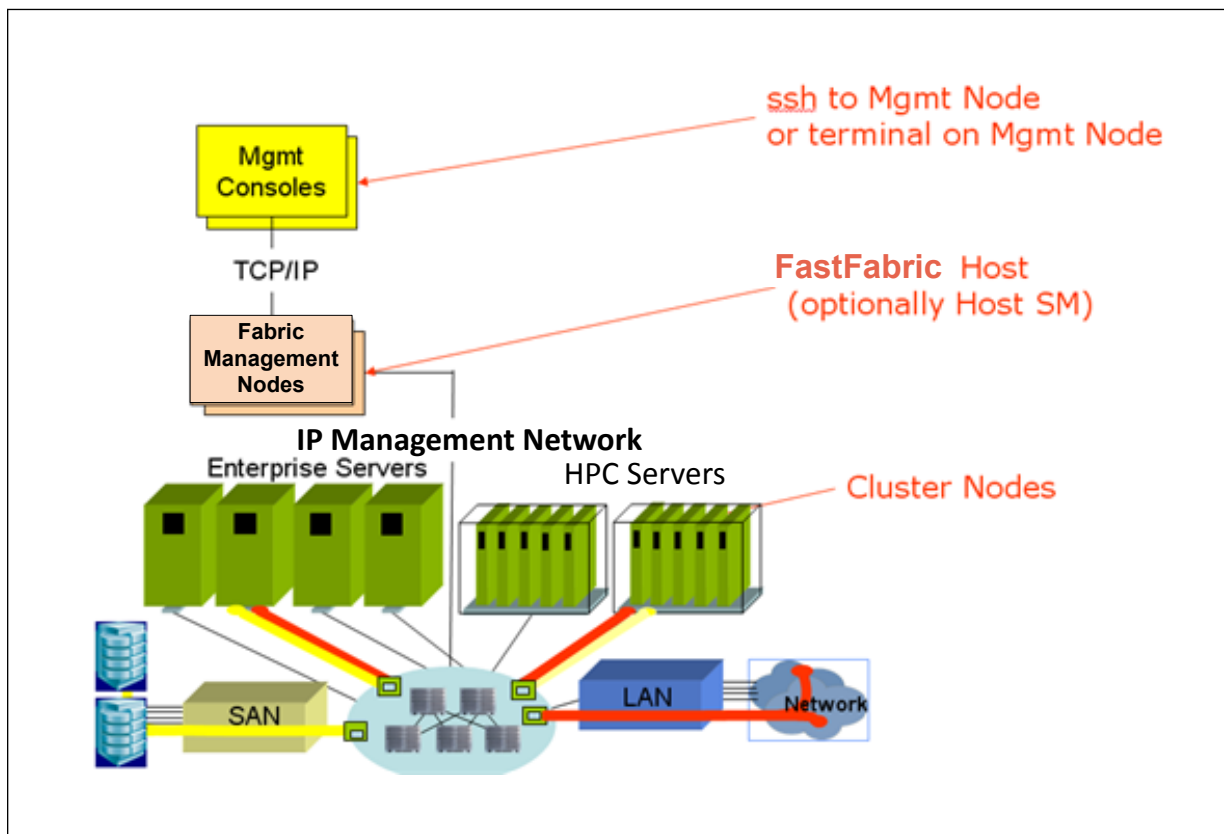
FastFabric consists of a hierarchy of commands and tools. In order to simplify learning and use, these tools all have similar command line arguments. Many of the FastFabric tools are designed to be easily extended via scripting or exporting data into other formats, such as spreadsheets.

The higher level tools allow you to focus on the names assigned to devices, and avoid the need to figure out LIDs or remember GUIDs for basic operations. As such, Intel recommends that you establish a naming convention for the cluster, and assign names to all the hosts and switches in the cluster.

### 2.2.1 FastFabric Architecture

FastFabric is typically installed on one or more Fabric Management Nodes. The Fabric Management Node must be connected to the rest of the cluster through the Intel® Omni-Path Fabric and a management network. The management network may be the primary Internet Protocol over InfiniBand\* (IPoIB) network or Ethernet\*. The management network is used for FastFabric host setup and administration tasks. It may also be used for other aspects of server administration or operation. Refer to the following figure for a high-level block diagram of the FastFabric architecture.



**Figure 4. FastFabric Architecture**

Depending on cluster size and design, the Fabric Management node may also be used as the master node for starting Message Passing Interface (MPI) jobs. It may also be used to run an Intel® Omni-Path Fabric Suite Fabric Manager and other management software. Refer to the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* for details and what combinations are valid.

**Note:** When IPoIB is used as the management network, FastFabric is not able to install host software or configure IPoIB. However in this configuration, FastFabric is able to support host software upgrades, verification, and all its other features.

If remote access to FastFabric is desired, set up remote access to the Fabric Management Node using the Intel® Omni-Path Fabric Suite Fabric Manager GUI, ssh, Telnet, X-Windows, VNC or any other mechanism that will allow the remote user to access a Linux\* Command Line shell. Typically FastFabric is used only by cluster administrators.

### 2.2.1.1 How FastFabric Works

FastFabric manages two types of switching devices that are managed by the "Chassis Setup/Admin" and "Externally Managed Switch Setup/Admin" menus.

The Chassis menu allows management of switching devices that are termed "internally managed." These include both edge and director class switching devices that have one or more management cards in place. The management card provides an environment



that exposes various TCP/IP services. This includes a command line interpreter login shell environment, with which FastFabric communicates. The device has an active Ethernet connection for LAN connectivity. The user is instructed to build a list of chassis in a "chassis" file, listing either the IP addresses or host names of the chassis to be managed; FastFabric provides tools to help discover such devices in the fabric and construct such a file. Communication with these devices is primarily out-of-band.

The Externally Managed Switch menu allows management of switching devices that are edge switches without management cards. Consequently, there is no environment present to provide TCP/IP services. There is no active Ethernet connection on the device. Therefore, communication to these devices must be accomplished in-band, via Intel® OPA management protocols designed specifically for this purpose. The user is instructed to build a list of switches in a "switches" file, listing the GUIDs of the switches to be managed; FastFabric provides tools to help discover such devices in the fabric and construct such a file.

FastFabric consists of a variety of tools to administrate hosts, chassis and externally managed switches. Depending on the tool, the method of accessing and administering the target devices may differ.

The following table describes the access methods that FastFabric uses.

**Table 1. FastFabric Methods**

Method	Examples
Inband access	Fabric performance, error and congestion monitoring. Fabric topology reports, SA database queries, fabric error and link speed analysis, tools for externally managed switches, etc.
Log in through a management network	Host setup and installation, tools for internally managed chassis, etc.
MPI job startup (can be inband or through a management network)	Verify MPI performance, running sample MPI benchmarks, host-to-switch cable test.

Tools that log into other hosts will do so in a password-less manner using ssh. Tools that log into internally managed chassis can also use ssh. Chassis tools can prompt for a single password for all chassis, use password-less ssh, or can be pre-configured with the password. These approaches permit the tools to operate with minimal user interaction, and for this reason reduce the time to perform operations against many hosts or chassis.

After initial installation, FastFabric can be configured to use IPoIB instead of the management network.

*Note:* IPoIB cannot be used to reconfigure IPoIB or install new hosts.

## 2.2.2 FastFabric Capabilities

### 2.2.2.1 FastFabric Command Hierarchy

FastFabric provides numerous powerful commands. These commands can be best understood as a hierarchy of capabilities permitting operations at high, mid and low levels.



### 2.2.2.1.1 Monitoring and Diagnostics

At the highest level, FastFabric provides an interactive Textual User Interface (TUI), called `opafastfabric`. The TUI provides an easy and efficient way to perform fabric deployment and verification, and diagnosis of typical fabrics. The TUI is structured in the typical sequence of operations for fabric verification. All of the functions that the TUI performs are also available using command line interface (CLI) commands. To aid in learning the commands, the TUI shows each CLI command as it executes it.

Other high level tools can provide an initial view of fabric status and health. These include the Intel® Omni-Path Fabric Suite Fabric Manager GUI, the interactive cluster errors and performance display tool (`opatop`), and the tools to verify cluster status as compared to a previous baseline (`opaallanalysis` and its sub-tools: `opalinkanalysis`, `opafabricanalysis`, `opachassisanalysis`, `opaesmanalysis`, and `opahostsmanalysis`).

When analyzing the fabric at a mid-tier of information, the next tier of tools include: `opafabricinfo`, `opareports`, `opareport`, `opaextractbadlinks`, `opaextractlink`, `opaextractsellinks`, and `opaextractstat2`. These tools provide very powerful ways to query the fabric. The `opaextract*` family of tools are all scripts that take advantage of `opareport` to generate delimited files that can be easily parsed or exported into spreadsheets for offline analysis. These scripts can also be good samples for the creation of site-specific sysadmin scripts.

At the next level of lower analysis there are additional tools. These provide direct access to more of the raw fabric information, such as port counters, LIDs, and other configured parameters. Tools in this tier include: `opashowallports`, `opaextractlids`, `opaextracterror`, `opaextractperf`, `opaextractstat`. Many of these tools are scripts that are also built on top of `opareport`. `opareport` is a foundational tool in FastFabric that provides a rich set of fabric analysis capabilities, and can provide both high level and very detailed output.

At the lowest level of analysis are tools that can access the management protocols directly. This can permit all the details of a given port or device to be viewed or analyzed. Typically, these tools only need to be used when debugging subtle issues. These tools include: `opahfirev`, `opashowmc`, `opafirmware`, `opasaquery`, `opapaquery`, `opafequery`, `opasmaquery`, `opaportinfo`, and `opapmaquery`.

### 2.2.2.1.2 Benchmark and Stress Tests

FastFabric includes a number of benchmarks and stress tests. These can be found in `/usr/lib/opa/src/mpi_apps` and `/usr/mpi/*/*/tests`. The `opacabletest` tool also provides a simple way to create high stress on all links in the fabric to aid in the verification of fabric stability.

In addition, other existing Intel® Omni-Path benchmarks and test programs may also be used to exercise the `libfabric` and `verbs` interfaces.

### 2.2.2.2 Host and Switch Management

FastFabric includes tools to manage both internally- and externally-managed switches, as well as hosts. These tools are in addition to the fundamental operational controls that the Fabric Manager provides for all devices in the fabric. Many of these capabilities are also available in the `opafastfabric` TUI.

For externally-managed switches, `opagenswitches` can assist in generating a list of the devices currently in the fabric, and `opaswitchadmin` provides the primary control and query functions to manage firmware, check status, and reboot.

For internally-managed switches, `opagenchassis` can assist in generating a list of the devices currently in the fabric (`opagenesmchassis` will generate the list of those currently running an embedded subnet manager (ESM)), and `opachassisadmin` provides the primary control and query functions to manage firmware, check status, and root. In addition `opapingall`, `opacmdall`, and `opasetupssh` can provide direct access to the switch CLI.

For hosts, `opahostadmin` provides typical control and query functions to manage host software and configuration. `opafindgood` and `opaverifyhosts` can provide analysis of the host status. In addition, `opapingall`, `opacmdall`, `opascpall`, `opadownloadall`, `opauploadall`, and `opasetupssh` are tools that are included to perform basic ssh and scp operations against the hosts.

### 2.2.2.3 Topology Analysis

FastFabric includes a rich set of topology analysis and verification capabilities. This can start with a pre-assembly description of the cluster design, from which `opaxlattopology` or `opaxlattopology_cust` can generate a `topology.xml` file for use by FastFabric and the Fabric Manager.

`opareport` has a number of reports for verifying the topology (`-o verify*`) and can do analysis of the current routing for credit loops, degree of path balance, and so forth. In addition, reports such as `opareport -o links`, `opaextractlink` and `opaextractsellinks` can provide an in-depth view of the fabric connectivity and design.

### 2.2.2.4 Link and Port Management

From a fabric perspective, a fabric consists of numerous fabric ports. The Fabric Manager controls and configures these ports, but FastFabric also includes a rich set of tools to analyze ports and perform some basic control functions, such as bouncing a port.

`opainfo` is an easy-to-use tool that can provide the primary status of the current host's ports. For controlling ports, `opaportconfig` can control a single port, while `opaenableports` and `opadisableports` can use the output from `opaextractbadlinks` or `opaextractsellinks` to disable a list of ports. `opadisablehosts` can also disable the ports on a list of hosts. `opaenableports` can reenale ports. For switches, `opaswdisableall` can disable unused ports on switches, while `opaswenableall` can re-enable them. To get all the low-level details of port status and configuration, `opaportinfo`, `opasmaquery` and `opapmaquery` may be used.

### 2.2.2.5 Focused Fabric Feature Analysis

Tools and reports are available to provide in-depth analysis of various fabric features.

Link quality, signal integrity, security errors, routing errors, and other issues can be analyzed using the following:

- `opafastfabric` TUI



- Intel® Omni-Path Fabric Suite Fabric Manager GUI
- `opatop`
- `opaallanalysis`
- `opaextractbadlinks`
- `opareport` (such as `-o errors`, `-o slow*`, and `-o mis*` reports)
- `opashowallports`

The details of Quality of Service (QoS) configuration and operation can be reviewed using the following:

- Intel® Omni-Path Fabric Suite Fabric Manager GUI
- `opasaquery -o vfinfo`
- `opasaquery -o path`
- various `opareport` options (such as `-ovfinfo`, `-o vfmember`, and `-o bfrctrl`)
- all the low-level details can be reviewed using `opasmaquery opareport -V -o comps -d 10`, and assorted `opasaquery` reports that show SL, SC and VL tables

Fabric routing can be analyzed using various `opareport` options, such as:

- `-o portusage`
- `-o treepathusage`
- `-o pathusage`
- `-o portgroups`
- `-o validateroutes`
- `-o validatepgs`
- `-o validatecreditloops`
- `-o linear`
- `-o mcast`

To view or analyze Link Width Downgrade (LWD), the following commands can be used:

- `opareport -o slowlinks`
- `opafabricanalysis` (uses `opareport -o slowlinks`)
- `opaextracterror` (uses `opareport -o comps`, shows main error counters)
- `opaextractperf` (uses `opareport -o comps`, shows per port counters)
- `opalinkanalysis slowlinks`

### 2.2.2.6 Scripting and Integration Enablement

Various additional tools can facilitate extending FastFabric, or integrating it with other tools. Among these are the XML processing tools (`opaxmlextract`, `opaxmlfilter`, and `opaxmlindent`), which can permit the XML output formats from `opareport` and/or the `opafm.xml` file itself to be easily parsed and analyzed in other scripts. The `opaextract*` scripts can provide samples of how to effectively use these tools.

`opagetvf` and `opagetvf_env` provide an easy way to extract key virtual fabric parameters to aid job scheduler and job launch integration with virtual fabrics.

### 2.2.2.7 Scripting on Top of FastFabric

Intel® Omni-Path Fabric Suite FastFabric was designed to make the scripting of OEM or site-specific tools easy to use. However, to ensure forward compatibility, scripts should be created using tools and arguments that are documented in the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

A number of the tools, such as the `opa*analysis` set of tools, are designed for easy use through exit code checks. These tools can easily be scripted to be run, and then, on bad exit codes, to issue emails or other forms of alerts to system administrators. Such mechanisms can be scheduled for regular execution by way of cron jobs. The file that is created by these tools can then be analyzed by the system administrators.


`opareport` is a powerhouse tool that provides a wide range of fabric data-gathering and analysis capabilities. The best way to script with this tool is to take advantage of its `-x` option to output XML. That output can then be easily parsed by `opaxmlextract` to extract sets of fields into delimited formats that can then be easily parsed by scripts, or exported to external tools such as spreadsheets. The `opa*extract` set of scripts are all built on top of `opareport -x` and `opaxmlextract`. These scripts can provide a great starting point by copying them and then creating new variations to meet your unique needs.

Intel recommends against creating scripts that attempt to directly parse `opareport -o` snapshot output. This format cannot be guaranteed to be forward-compatible with future FastFabric software releases. Most of the information in an `opareport -o` snapshot is also available in a forward-compatible format via `opareport -x -o comps -d 10 -s`. The remainder can be found in other `opareport` output by using different options. See the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for documentation on `opareport` and other tools.

Intel also recommends against creating scripts that attempt to parse the human-readable output formats produced by the tools. Intel reserves the right to refine these formats in future FastFabric software releases, and therefore, these formats cannot be guaranteed to be forward-compatible.

### 2.2.2.8 Customer Support Data Gathering

Detailed information about the current fabric status and configuration can be quickly obtained using `opacapture` for a single node, or `opacaptureall` for multiple nodes, to aid customer support.



### 2.2.2.9 Other Tools and Capabilities

In addition, a number of the non-Infiniband\*-specific OpenFabrics Alliance\* (OFA) tools will continue to function on an Intel® Omni-Path Fabric, and can provide additional information. Among these are `ibv_devinfo` (note that MTU will not correctly report MTUs beyond 4K), `ibstat` (note that some of the extended physical states for Intel® Omni-Path ports, such as offline, will not be reported properly), `ibsrpdm`, and `ibv_devices`.

The `ibacm` Distributed Subnet Administrator Provider (DSAP) plugin may be queried and debugged using `opa_osd_dump`, `opa_osd_exercise`, `opa_osd_perf`, and `opa_osd_query`.

`opapacketcapture` can provide traces of many of the verbs packets, including all the management packets, to enable offline analysis and debug in Wireshark using the Intel® Omni-Path dissector.



## 3.0 Getting Started

---

This section provides instructions and information for getting started with the Intel® Omni-Path Fabric Suite FastFabric tools.

### 3.1 Important Note on First-Time Installations

**This user guide is not an installation guide. If you are installing and configuring the fabric for the first time, you must refer to the *Intel® Omni-Path Fabric Software Installation Guide*.**

### 3.2 Starting Up the Tools

**Note:** To run the Intel® Omni-Path Fabric Suite FastFabric tools described in this manual, you must have root privileges.

#### 3.2.1 Accessing the Intel FastFabric OPA Tools Menu

The Intel FastFabric OPA Tools menu allows you to configure and manage the Intel® Omni-Path Fabric.

##### Using the `opafastfabric` Command

To start up the Intel FastFabric OPA Tools menu from the command prompt, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter `opafastfabric`.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

##### From the Intel OPA Software Menu

To start up the Intel FastFabric OPA Tools menu from the Intel OPA Software main menu, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter `opaconfig`.





The Intel OPA [version] Software main menu is displayed.

```
Intel OPA X.X.X.X.X Software

1) Show Installed Software
2) Reconfigure OFA IP over IB
3) Reconfigure Driver Autostart
4) Generate Supporting Information for Problem Report
5) FastFabric (Host/Chassis/Switch Setup/Admin)
6) Uninstall Software

X) Exit
```

3. At the cursor, type 5.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

### 3.2.2 Accessing the Fabric Performance Monitor

The Fabric Performance Monitor allows you to monitor performance, congestion, and error information in a fabric.

#### Using the opatop Command

To start up the Fabric Performance Monitor from the command prompt, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opatop**.

The Fabric Performance Monitor Summary screen is displayed.

```
opatop: Img: 10s @ Wed Sep 14 11:29:52 2016, Live
Summary: SW:      0 Ports: SW:      0 HFI:      2      Link:      1
          SM:      1 Node Fail:      0 Skip:      0 Port Fail:      0 Skip:      0
          AvgMBps  MinMBps  MaxMBps  AvgKPps  MinKPps  MaxKPps
0 All      Int      0      0      0      0      0      0
  Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
1 HFIs     Int      0      0      0      0      0      0
  Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
2 SWs      No ports in group

Master-SM: LID: 0x0001 Port: 1 Priority: 0 State: Master
           Name: phcppriv10 hfi1_0
           PortGUID: 0x0011750101575300
Secondary-SM: none
```



```
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

### From the Intel FastFabric OPA Tools Menu

To start up the Fabric Performance Monitor menu from the Intel FastFabric OPA Tools menu, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opafastfabric**.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

3. At the cursor, type 5.

The FastFabric OPA Fabric Monitoring menu is displayed.

```
FastFabric OPA Fabric Monitoring Menu

0) Fabric Performance Monitoring          [ Skip ]

P) Perform the Selected Actions           N) Select None
X) Return to Previous Menu (or ESC)
```

**Table 2. FastFabric OPA Fabric Monitoring Menu Descriptions**

Menu Item	Description
0) Fabric Performance Monitoring	Allows you to access the TUI that monitors the performance, congestion, and error information about a fabric. Associated CLI Command: <code>opatop</code>

4. Type 0 to toggle to the [Perform] option.

5. Type P to perform the operation.

The Fabric Performance Monitor information is displayed.

```
opatop: Img: 10s @ Fri Sep 16 11:35:24 2016, Live
Summary: SW: 0 Ports: SW: 0 HFI: 2 Link: 1
SM: 1 Node Fail: 0 Skip: 0 Port Fail: 0 Skip: 0
      AvgMBps  MinMBps  MaxMBps  AvgKPps  MinKPps  MaxKPps
0 All      Int      0      0      0      0      0      0
      Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
1 HFIs      Int      0      0      0      0      0      0
      Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
2 SWs      No ports in group
```



```
Master-SM: LID: 0x0001 Port: 1 Priority: 0 State: Master
          Name: phcppriv10 hfil_0
          PortGUID: 0x0011750101575300
Secondary-SM: none
```

```
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

### 3.3 Intel FastFabric OPA Tools TUI Overview

The Intel FastFabric OPA Tools TUI allows you to perform common fabric management tasks including setting up and managing the chassis, switches, and hosts.

**Note:** Additional CLI options and details about FastFabric command tools are available and documented in the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

The following is an example of the Intel FastFabric OPA Tools main menu.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

### 3.4 How to Use the FastFabric TUI

The FastFabric TUI menus are set up for ease of use. The submenus are designed to present operations in the order they would typically be used during an installation.

**Note:** All FastFabric TUI menu alpha-based options are case-insensitive.

#### Selecting Menu Items and Performing Operations

1. From the Intel FastFabric OPA Tools main menu, select the target menu item (0-4<sup>1</sup>).

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

---

<sup>1</sup> For menu item 5, refer to [How to Use the Fabric Performance Monitor TUI](#) on page 30.



The target menu is displayed as shown in the example below:

```
FastFabric OPA Chassis Setup/Admin Menu
Chassis File: /etc/sysconfig/opa/chassis
Setup:
0) Edit Config and Select/Edit Chassis File [ Skip ]
1) Verify Chassis via Ethernet Ping [ Skip ]
2) Update Chassis Firmware [ Skip ]
3) Set Up Chassis Basic Configuration [ Skip ]
4) Set Up Password-Less SSH/SCP [ Skip ]
5) Reboot Chassis [ Skip ]
6) Get Basic Chassis Configuration [ Skip ]
7) Configure Chassis Fabric Manager (FM) [ Skip ]
8) Update Chassis FM Security Files [ Skip ]
9) Get Chassis FM Security Files [ Skip ]
Admin:
a) Check OPA Fabric Status [ Skip ]
b) Control Chassis Fabric Manager (FM) [ Skip ]
c) Generate All Chassis Problem Report Info [ Skip ]
d) Run a Command on All Chassis [ Skip ]
Review:
e) View opachassisadmin Result Files [ Skip ]

P) Perform the Selected Actions N) Select None
X) Return to Previous Menu (or ESC)
```

2. Type the key corresponding to the target menu item (0-9, a-d) to toggle the Skip/Perform selection.

More than one item may be selected.

3. Type **P** to perform the operations that were selected.

**Notes:**

- If more than one menu item is selected, the operations are performed in the order shown in the menu. This is the typical order desired during fabric setup.

- If you want to perform operations in a different order, you must select the first target menu item, type **P** to perform the operation, then repeat this process for the next menu item operation to be performed, and so on.

4. Type **N** to clear all selected items.
5. Type **X** or press **Esc** to exit this menu and return to the Main Menu.

### Aborting Operations

While multiple menu items are performing, you have an opportunity to abort individual operations as they come up. After each operation completes and before the next operation begins, you are prompted as shown below:

```
Hit any key to continue...
```

- Press **Esc** to stop the sequence of operations return to the previous menu.  
Any unperformed operations are still highlighted in the menu. To complete the selected operations, type **P**.
- Press any other key to perform the next selected menu item being performed.  
This prompt is also shown after the last selected item completes, providing an opportunity to review the results before the screen is cleared to display the menu.



## Submenu Configuration Files

On each FastFabric submenu, item 0 permits a different file to be selected and edited (using the editor selected by the EDITOR environment variable). It also permits reviewing and editing of the `opafastfabric.conf` file. The `opafastfabric.conf` file guides the overall configuration of FastFabric and describes cluster-specific attributes of how FastFabric operates. It is discussed in greater detail in [Configuration Files for FastFabric](#) on page 34.

At the top of each FastFabric submenu screen beneath the title, the directory and configuration file containing the components on which to operate are shown.

In the example below, the configuration file is noted in bold.

```
FastFabric OPA Host Setup Menu
Host File: /etc/sysconfig/opa/hosts
Setup:
0) Edit Config and Select/Edit Host File      [ Skip ]
1) Verify Hosts Pingable                     [ Skip ]
2) Set Up Password-Less SSH/SCP              [ Skip ]
```

**Note:** During the execution of each menu selection, the actual FastFabric command line tool being used is shown. This can be used as an educational aid to learn the command line tools.

The example snippet below shows how the CLI is displayed in the TUI execution.

```
Performing Chassis Admin: Verify Chassis via Ethernet Ping
Executing: /usr/sbin/opapingall -C -p -F /etc/sysconfig/opa/chassis
```

## 3.5 Fabric Performance Monitor TUI Overview

The Fabric Performance Monitor TUI allows you to monitor performance, congestion, and error information about a fabric.

The following is an example of the (`opatop`) Fabric Performance Monitor TUI.

```
opatop: Img: 10s @ Wed Sep 14 11:29:52 2016, Live
Summary: SW:      0 Ports: SW:      0 HFI:      2      Link:      1
          SM:      1 Node Fail:      0 Skip:      0 Port Fail:      0 Skip:      0
          AvgMBps  MinMBps  MaxMBps  AvgKPps  MinKPps  MaxKPps
0 All      Int      0          0          0          0          0          0
  Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
1 HFIs     Int      0          0          0          0          0          0
  Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
2 SWs      No ports in group

Master-SM: LID: 0x0001 Port: 1 Priority: 0 State: Master
          Name: phcppriv10 hfil_0
          PortGUID: 0x0011750101575300
Secondary-SM: none
```

```
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

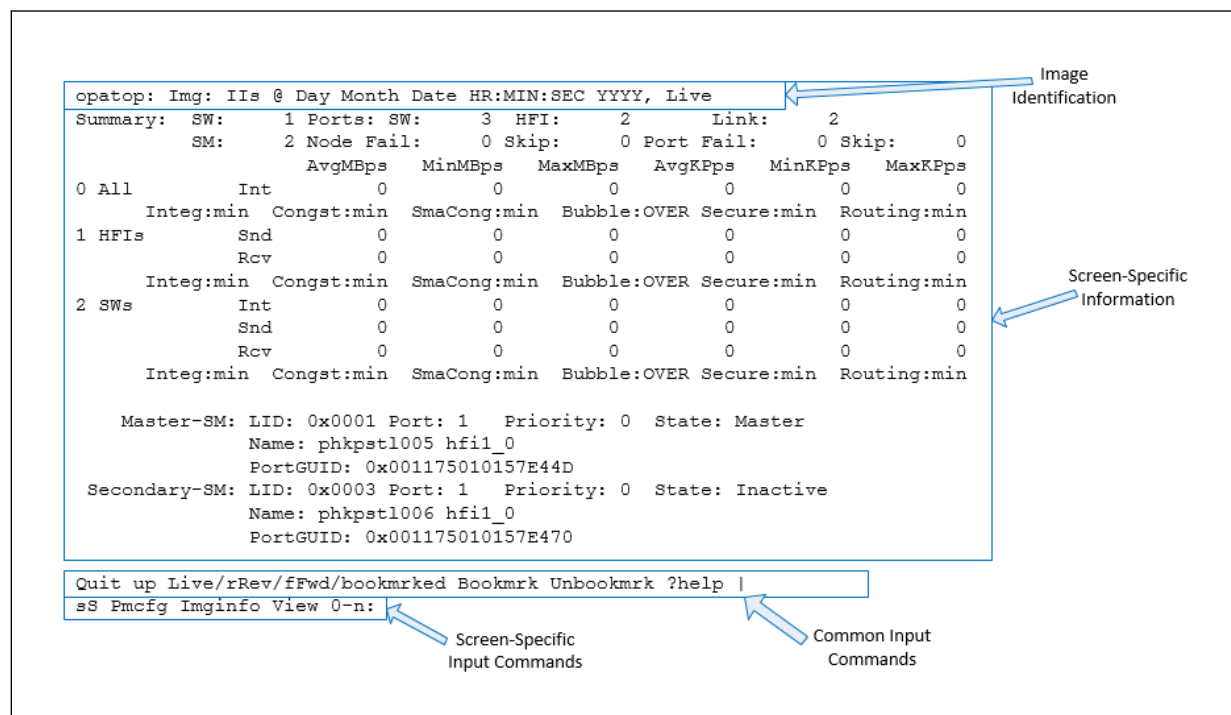
### 3.6 How to Use the Fabric Performance Monitor TUI

The Fabric Performance Monitor TUI allow you to view and interact with live performance data.

#### Reading the TUI Screens

The figure below shows the major sections common to all Fabric Performance Monitor TUI screens.

**Figure 5. Fabric Performance Monitor TUI Screen (Example)**



**Table 3. Fabric Performance Monitor TUI Descriptions**

Section of Screen	Description
opatop	Refers to the CLI command that initiates to the Fabric Performance Monitoring TUI. <b>NOTE:</b> opatop may be used interchangeably with Intel® Fabric Performance Monitoring TUI within this manual.
Image Identification	Displays the following image (Img) information: <ul style="list-style-type: none"> <li>Image interval (II): The time over which this image data is relevant. <ul style="list-style-type: none"> <li>For in-memory images, this value is equal to the PM Sweep Interval.</li> <li>For images stored on disk (Short Term History), the interval is equal to the sum of all the intervals for each image compounded into the composite (disk) image.</li> </ul> </li> </ul>

*continued...*



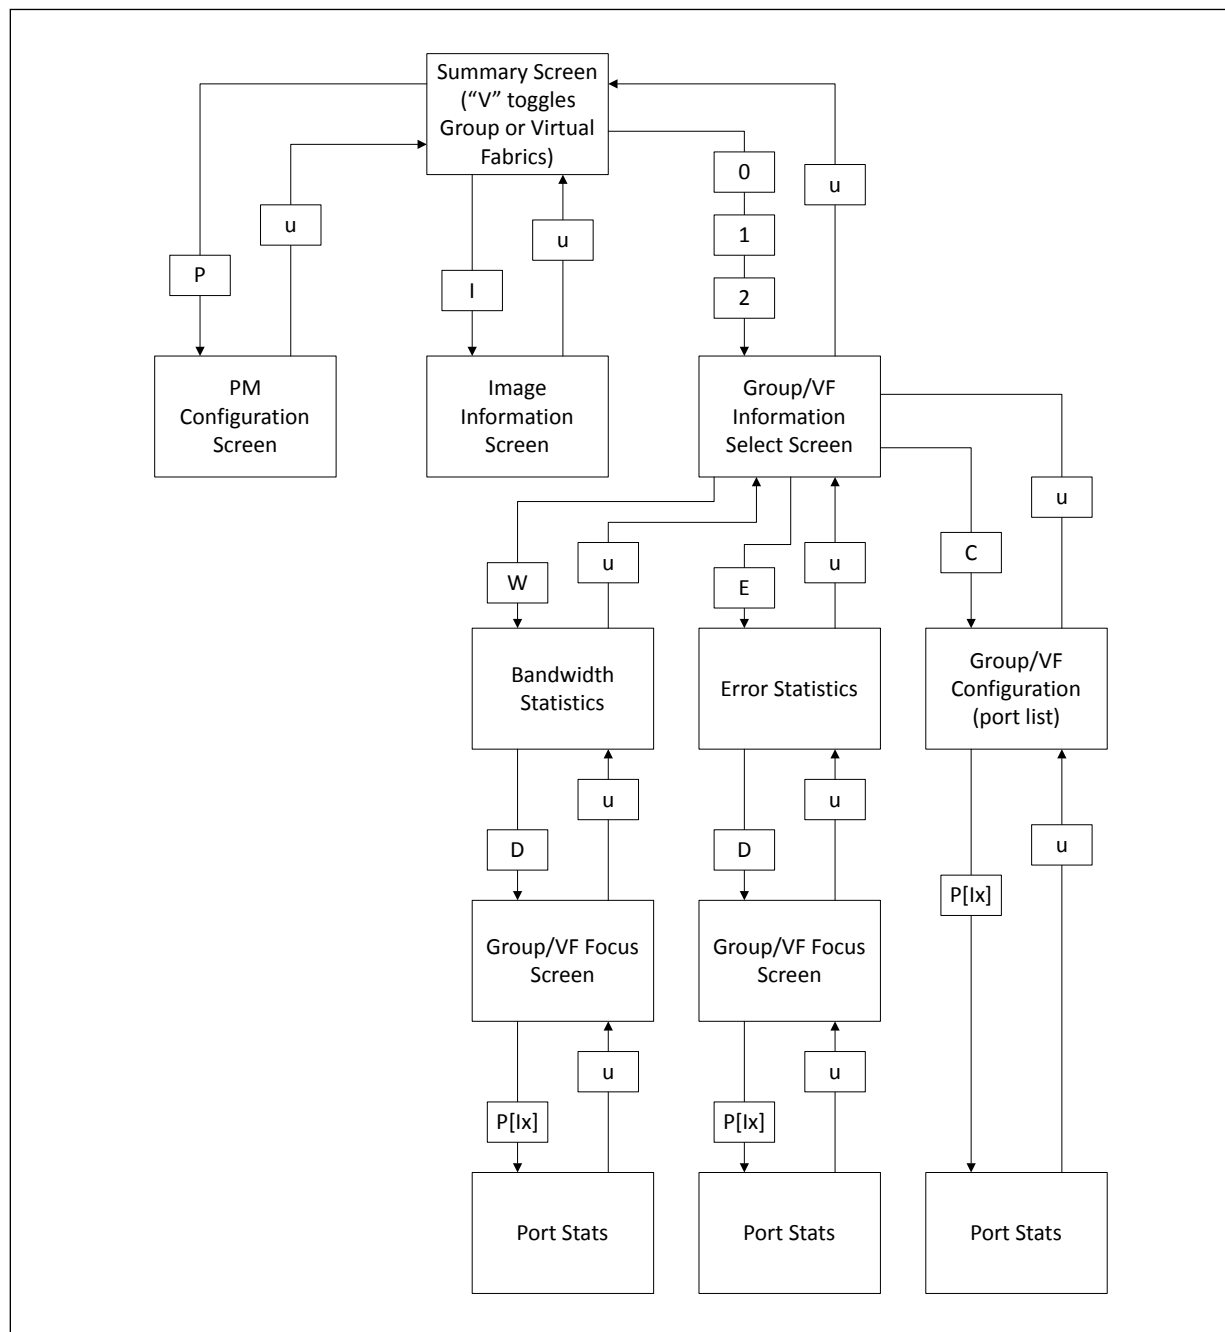
Section of Screen	Description
	<p><b>NOTE:</b> The interval can change when transitioning between images stored in memory and images stored on disk.</p> <ul style="list-style-type: none"> <li>Timestamp for the image being displayed in the format Day Month Date HR:MIN:SEC YYYY (example, Wed Sep 14 11:29:52 2016) If a Live image is not being displayed, the current time ('Now:') is also shown.</li> <li>Type of image <ul style="list-style-type: none"> <li>— Live</li> <li>— Hist (History)</li> <li>— Bkmk (Bookmark)</li> </ul> </li> </ul>
Screen-Specific Information	<p>Displays information and layout of the selected screen.</p> <p><b>NOTE:</b> Each screen is different and will be discussed in subsequent sections.</p>
Common Input Commands	<p>Displays the common input commands that appear on every screen and perform the same action.</p> <ul style="list-style-type: none"> <li>Q/q – Quit program</li> <li>u – Up to previous screen</li> <li>L – Select Live image</li> <li>r – Navigate reverse 1 sweep</li> <li>R – Navigate reverse 5 sweeps</li> <li>f – Navigate forward 1 sweep</li> <li>F – Navigate forward 5 sweeps</li> <li>b – Select (previously) bookmarked image</li> <li>B – Bookmark currently selected image</li> <li>U – Unbookmark image</li> <li>? – Help provides information about the screen contents and input commands.</li> </ul> <p>Commands are case insensitive except where specifically noted otherwise. The ENTER key must be pressed after multi-character commands and for Quit.</p>
Screen-Specific Input Commands	Displays the screen-specific commands.

## Navigating the Screens

The Fabric Performance Monitoring TUI allows you to access various screens in a hierarchal manner to examine the state of a fabric. Through the screen-specific commands, each screen will provide access to the next screen or back to the parent screen.

The Fabric Performance Monitoring TUI screen navigational hierarchy is shown below.

**Figure 6. Fabric Performance Monitoring TUI Navigation**



As an example, if you want to navigate from the Group Info Sel screen to the Group BW Stats screen, perform the following steps:

1. The Group Info Sel screen is shown below.

```

opatop: Img: 10s @ Thu Sep 22 15:44:47 2016, Live
Group Info Sel: HFIs
Int NumPorts: 2 Rate Min: 100g Max: 100g
  
```





```
Ext NumPorts: 0
  Group BW Summary (W)
  Group Err Summary (E)
  Group Config (C)

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | W E C:
```

The selections for the next level of screens are displayed as:

```
Group BW Summary (W)
Group Err Summary (E)
Group Config (C)
```

The menu options are shown in the screen-specific commands as:

```
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | W E C:
```

2. From the Group Info Sel screen, enter **W**.

The Group BW Stats screen is displayed.

```
opatop: Img: 10s @ Thu Sep 22 15:52:27 2016, Live
Group BW Stats: HFIs Criteria: Util-High Number: 10
Int: TotMBps AvgMBps MinMBps MaxMBps TotKPps AvgKPps MinKPps
MaxKPps
0 0 0 0 0 0 0
0 Buckt 0+% 10+% 20+% 30+% 40+% 50+% 60+% 70+% 80+% 90+%
2 0 0 0 0 0 0 0 0 0
Failed Int Ports: PMA: 0 Topo: 0

Int Congestion Max 0+ 25+ 50+ 75+ 100+
0 2 0 0 0 0

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | cC N0-n Detail:
```

3. Type **u** (lowercase) to return to the Group Info Sel screen.
4. Type **u** (lowercase) to return to the Summary screen.

**Important:** To switch between Port and Virtual Fabric Grouping screens, press **v** at the Summary screen and navigate through the hierarchy.

### 3.7 Configuration of IPoIB Name Mapping

The FastFabric tools support the concept of a management network and an IPoIB network. For some clusters, the management network will be a low-speed network such as 1 GB or 10 GB Ethernet. For other clusters, IPoIB may serve double duty as the host management network.

**Note:** When using IPoIB as the management network, the initial installation of Fabric software cannot be done using FastFabric.

The various FastFabric tools will translate from host names provided to and from IPoIB names as needed. This permits the given host names to be either management network or IPoIB network names.



- The default configuration file assumes that IPoIB host names are formed by adding a `-opa` suffix to the management network name.
- If a different suffix is desired, `FF_IPOIB_SUFFIX` can be changed.
- If IPoIB is also being used as the management network, `FF_IPOIB_SUFFIX` can be set to an empty string `""`.

The translation is driven by the following functions within `opafastfabric.conf`:

- `ff_host_basename` – Given a management network or IPoIB hostname, translate to management network name; should match hostname `-s`
- `ff_host_basename_to_ipoib` – Given a management network name, translate to IPoIB hostname

More complex mappings can be specified by implementing alternate algorithms for these functions.

**Note:** When managing a cluster where the IPoIB settings on the compute nodes are incompatible with the Fabric Management node, Intel recommends that you do not run IPoIB on the Fabric management nodes.

## 3.8 Configuration Files for FastFabric

The FastFabric configuration files allow you to configure and change the basic settings and variables for the fabric and each of its components. These files are pushed out across the network ensuring that each component is synchronized.

Configuration files are located under the `/etc/sysconfig/opa` directory.

Sample files are installed into `/usr/lib/opa/samples` with the suffix `-sample`. These files show the defaults of the given release.

**Note:** Do not edit the sample files.

Configuration files are self-documented as shown in the example snippet below.

```
#!/bin/bash
# [ICS VERSION STRING: @(#) ./fastfabric/samples/opafastfabric.conf-sample
10_3_0_0_51 [09/20/16 23:52]
# This is a bash sourced config file which defines variables used in
# fast fabric tools. Command line arguments will override these settings.
# Assignments should be scripted such that this file does not override
# exported environment settings, as shown in the defaults below

if [ "$CONFIG_DIR" = "" ]
then
    if [ -d /etc/sysconfig ]
    then
        CONFIG_DIR=/etc/sysconfig
    else
        CONFIG_DIR=/etc
    fi
    export CONFIG_DIR
fi

# Override default location for HOSTS_FILE
export HOSTS_FILE=${HOSTS_FILE:-$CONFIG_DIR/opa/hosts}
```



```
# Override default location for CHASSIS_FILE
export CHASSIS_FILE=${CHASSIS_FILE:-$CONFIG_DIR/opa/chassis}
```

You can find more information about the various configuration variables in the "Environment Variables" section for the applicable commands. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*

### 3.8.1 FastFabric Configuration File

The FastFabric configuration file allows you to view the default settings and modify the variables for most of the FastFabric command line options.

The file is located under `/etc/sysconfig/opa/opafastfabric.conf`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

**Note:** Command line arguments will override these settings.

#### Modifying the FastFabric Configuration File

1. To modify the configuration file, refer to the following FastFabric TUI procedures:
  - [Editing the Configuration Files for Chassis Setup](#) on page 46
  - [Editing the Configuration Files for Externally-Managed Switch Setup](#) on page 64
  - [Editing the Configuration Files for Host Setup](#) on page 76
  - [Editing the Configuration Files for Host Verification](#) on page 90
2. Adhere to the following requirements when editing the file:
  - The configuration file is a bash shell script that will be included by each tool. As such, the file should be implemented so that the environment variables defined prior to execution will not be altered.

The sample code below shows the bash syntax that allows only uninitialized variables to be overwritten by the configuration file:

```
var= "${var:-value}"
```

### 3.8.2 Ports List Configuration File

The Ports List configuration file allows you to specify the local HFI ports (i.e., subnets) that FastFabric will use in assorted commands for fabric access.

The file is located under `/etc/sysconfig/opa/ports`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

Alternate filenames may be specified in `opafastfabric.conf` using environment variables, or on the command line. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information.

#### Modifying the Ports List Configuration File

1. To modify the configuration file, refer to the following FastFabric TUI procedures:

- [Editing the Configuration Files for Chassis Setup](#) on page 46
  - [Editing the Configuration Files for Externally-Managed Switch Setup](#) on page 64
  - [Editing the Configuration Files for Host Verification](#) on page 90
2. Adhere to the following requirements when editing the file:
- Each line of the port list file may specify a single port, a comment, or another port list file to include.
  - Ports are specified as `hfi:port`. No spaces are permitted.

The first Host Fabric Interface Adapter is 1, and the first port is 1. The value 0 for Host Fabric Interface or port has special meaning. The allowed formats are shown in the example below.

```
# [ICS VERSION STRING: @(#) ./fastfabric/samples/ports-sample 10_3_0_0_51
[09/20/16 23:52]
# This file defines the local HFI ports to use to access the fabric(s)
#
# specify one line per HFI port of the form hfi:port such as:
#   0:0 = 1st active port in system
#   0:y = port y within system
#   x:0 = 1st active port on HFI x
#   x:y = HFI x, port y
# The first HFI in the system is 1. The first port on an HFI is 1.
0:0
```

- Files to be included may be specified using an `include` directive followed by a file name.  
In general, specified file names should be absolute path names. If relative path names are used, they will be searched for within the current directory, then `/etc/sysconfig/opa`.
- Comments may be placed on any line by using a `"#"` to precede the comment.  
On lines with a port or `include` directive, the `"#"` must be white-space separated from any preceding port or included file name.

### 3.8.3 Chassis List Configuration Files

The Chassis List configuration files allow you to specify the Intel chassis that FastFabric will operate against for many operations.

The `opagenchassis` command can be used to help locate chassis in the fabric and generate a chassis file.

The files are located under `/etc/sysconfig/opa/chassis` and `/etc/sysconfig/opa/esm_chassis`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

Alternate filenames may be specified in `opafastfabric.conf`, using environment variables or on the command line. Refer to *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information.



### Modifying the Chassis List Configuration Files

1. To modify the configuration files, refer to [Editing the Configuration Files for Chassis Setup](#) on page 46.
2. Adhere to the following requirements when editing the file:
  - Each line of the chassis list file may specify a single chassis, a comment, or another chassis list file to include.
  - Chassis are specified by the chassis management network IP address or by a resolvable TCP/IP name.  
*Note:* Typically, names are used for readability.
  - If Ethernet is being used for the management network, specify the name corresponding to the ethernet IP address of the chassis.
  - Files to be included may be specified using an `include` directive followed by a file name.  
 In general, specified file names should be absolute path names. If relative path names are used, they will be searched for within the current directory, then `/etc/sysconfig/opa` directory.
  - Comments may be placed on any line by using a "#" to precede the comment.  
 On lines with chassis or `include` directives, the # must be white-space separated from any preceding name, IP address, or included filename.

#### 3.8.3.1 Performing Operations Against a Selection of Slots Within a Chassis

Normally, operations are performed against the management card in the chassis. For operations such as `opacmdall`, the command is executed against the management interface for the given chassis. For more sophisticated operations, such as firmware update, a directory with firmware for each chassis card type can be supplied and all cards in the chassis will be updated with the appropriate firmware from that directory. However, in some cases it may be desirable to perform operations against a specific subset of cards within the chassis.

1. Augment the chassis IP address, a name within a chassis list, or a chassis file with a list of slot numbers on which to operate.

This is done in the form:

```
chassis:slot1,slot2,...
```

*Note:* There must be no spaces within the chassis name and/or slot list.

- This format is used by `opacmdall` and chassis firmware update.  
 It may be used anywhere a chassis name or IP address is valid, such as the `-H` option, the `CHASSIS` environment variable, or chassis list files.
- The slot number specified is ignored on some operations (such as `opapingall`).
- Only slots containing management cards may be specified with this format.
- For all Intel® Omni-Path Chassis 100 Series chassis, slot 0 is always an alias for the presently active management card for the chassis.



For the remainder of slot usages in the chassis, the `chassisQuery` command can be executed against a given chassis to identify which slots have management cards.

*Note:* For any operation, care should be taken that a given chassis is listed only once with all relevant slots as part of that single specification. This is important so that parallel operations do not cause conflicting concurrent operations against a given chassis.

### 3.8.4 Externally-Managed Switch List Configuration File

The Externally-Managed Switch List configuration file allows you to specify the externally-managed Intel switches that FastFabric will operate against for many operations.

The file is located under `/etc/sysconfig/opa/switches`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

Alternate file names may be specified in `opafastfabric.conf`, using environment variables or on the command line. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information.

#### Modifying the Externally-Managed Switch List Configuration File

1. To modify the configuration file, refer to [Editing the Configuration Files for Externally-Managed Switch Setup](#) on page 64.
2. Adhere to the following requirements when editing the file:
  - Each line of the switch list file may specify a single switch, a comment, or another switch list file to include.
  - Switches are specified in the comma-separated form:  
`guid,nodeDesc,distance` where
    - `guid` – Node GUID of the switch optionally followed by a colon and `hfi:port`
    - `nodeDesc` – Optional node description should be programmed into the switch by FastFabric.  
It is recommended to supply a unique nodeDesc for each switch to simplify management of the cluster.
    - `distance` – Optional relative distance of the switch from the FastFabric node.  
This is used by reboot operations to first operate on switches furthest from the FastFabric node. Nodes without a distance specified will be treated as furthest. Refer to [Defining the Distance Value](#) on page 39.
  - The GUID will be used to select the switch and on firmware update operations, the node description will be written to the switch such that other FastFabric tools (such as `opasaquery` and `opareport`) can provide a more easily readable name for the switch.  
The node description can also be updated as part of switch basic configuration.
  - The `hfi:port` may be used to specify which local port (subnet) to use to access the switch.



If this is omitted, all local ports specified will be checked for the switch and the first port found to be able to access the switch will be used to access it. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information about how to specify an `hfi:port` value.

- Files to be included may be specified using an `include` directive followed by a file name. File names specified should generally be absolute path names. If relative path names are used, they will be searched for within the current directory then `/etc/sysconfig/opa`.
- Comments may be placed on any line by using a `"#"` to precede the comment.

On lines with `chassis` or `include` directives, the `"#"` must be white-space separated from any preceding GUID, name, or included file name.

- Intel recommends that a unique node description be specified for each switch. This name should follow typical naming rules and use the characters a-z, A-Z, 0-9, and underscore. No spaces are allowed in the node description. Additionally, names should not start with a digit.
  - For externally-managed switches, the node GUID can be found on a label on the bottom of the switch.
  - Alternately the node GUIDs for switches in the fabric can be found using a command such as:

```
opasaquery -t sw -o nodeguid
```

**Note:** The preceding command will report all switch node GUIDs, including those of internally-managed chassis such as the Intel® Omni-Path Switch 100 Series switches. GUIDs for internally-managed chassis cannot be specified for use in the `switches` file.

## Defining the Distance Value

The `opagenswitches` command can be used to help locate externally-managed switches in the fabric and generate a `switches` file. The `opagenswitches` tool will by default provide the proper distance value relative to the FastFabric node from which it was run. This capability requires use of IBTA standard TraceRecord queries that are not supported by openSM, but can be supplied by the Intel® Omni-Path Fabric Suite Fabric Manager (FM). Alternatively the `opagenswitches -R` option can suppress generation of this field. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information.

In a typical pure fat tree topology, with externally-managed switches as edge switches and internally-managed switches as core switches, you can also manually specify proper distance by specifying `1` for the distance value of the switch next to the FastFabric node. Note that in such a topology, all other switches are an equal length from the FastFabric node, and a missing distance value will cause them to be treated as having a distance value that is larger than any other found in the file. Therefore, the other switches would be rebooted first and the FastFabric node's switch would be rebooted last.

FastFabric is topology-aware when updating externally-managed switch firmware or resetting the switches. Switches furthest from the FastFabric node are updated or reset first, and then each switch, working toward the FastFabric node. This way, switches that are rebooted are not in the path between the FastFabric node and others that are being rebooted.

The ordering is controlled by an optional `distance` field in the `switches` file or the `switches` provided on the command line. The `distance` field indicates the relative distance from the FastFabric node for each switch. Any `switches` file entries that do not specify a distance value are treated as having a value larger than any others in the file. The `switches` file contains any one of the following formats per line:

- `nodeguid`
- `nodeguid,,distance`
- `nodeguid:hfi:port`
- `nodeguid:hfi:port,,distance`
- `nodeguid,nodename`
- `nodeguid,nodename,distance`
- `nodeguid:hfi:port,nodename`
- `nodeguid:hfi:port,nodename,distance`

### 3.8.5 Hosts List Configuration Files

The Hosts List configuration files allow you to specify the hosts that FastFabric will operate against for many operations.

The files are located under `/etc/sysconfig/opa/hosts` and `/etc/sysconfig/opa/allhosts`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

Alternate filenames may be specified in `opafastfabric.conf`, using environment variables or on the command line. Refer to *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information.

#### Modifying the Hosts List Configuration Files

1. To modify the configuration file, refer to the following FastFabric TUI procedures:
  - [Editing the Configuration Files for Host Setup](#) on page 76
  - [Editing the Configuration Files for Host Verification](#) on page 90
2. Adhere to the following requirements when editing the file:
  - Each line of the host list file may specify a single host, a comment or another host list file to include.
  - Hosts are specified by IP address or by a resolvable TCP/IP hostname. Typically, hostnames are used for readability.





Also, some FastFabric tools will translate the supplied host names to IPoIB hostnames, in which case names are generally easier to translate than numeric IP addresses. Typically, management network host names are specified. However, if desired, IPoIB hostnames or IP addresses may be used. This can accelerate large file transfers and other operations.

- If Ethernet is being used for the management network, specify the hostname corresponding to the ethernet IP address.
- Files to be included may be specified using an `include` directive followed by a file name.

In general, specified file names should be absolute path names. If relative path names are used, they will be searched for within the current directory, then `/etc/sysconfig/opa` directory.

- Comments may be placed on any line by using a `#` to precede the comment.

On lines with hosts or include directives, the `#` must be white-space separated from any preceding host name, IP address, or included file name.

### 3.8.6 Port Statistics Thresholds Configuration File

The `opamon.conf` configuration file defines the thresholds for each port statistic. Error Counters are specified in absolute number of errors since last cleared. If the threshold for a given statistic is not defined or is set to 0 (disabled), the given statistic will not be checked. This file is used by the following commands:

- `opareport`

*Note:* When used by `opareport` or fabric health tools, the counts are absolute values and are applied against the counters as found in the system.

- `opafabricanalysis`
- `opalinkanalysis`
- `opaextractbadlinks`
- `opaextractstat`
- `opaextractstat2`
- `opaallanalysis`

The file is located under `/etc/sysconfig/opa/opamon.conf`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

### 3.8.7 Signal Integrity Thresholds Configuration File

The `opamon.si.conf` configuration file defines thresholds for port counter signal integrity. This file allows analysis for any non-zero error counters related to signal integrity (bad cables, etc.) and can be enabled by adding the `-c` option to many FastFabric tools including:

- `opareport`
- `opaextractbadlinks`
- `opaextractstat`
- `opaextractstat2,`



- opalinkanalysis
- opacabletest
- opafabricanalysis

The file is located under `/etc/sysconfig/opa/opamon.si.conf`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

### 3.8.8 Fabric Topology Input File

The Fabric Topology input file (`topology.0:0.xml`) allows you to specify the expected fabric topology and augmented fabric information (such as cable labels, types, lengths, SM details, node details, link details, etc.). If present, this file will be used by assorted FastFabric commands such as `opareports`, `opafabricanalysis`, and `opaallanalysis`.

The file is located under `/etc/sysconfig/opa/topology.0:0.xml`.

A sample file is provided, and matches the internal defaults of the FastFabric tools.

Alternate filenames may be specified in `opafastfabric.conf`, using environment variables or on the command line. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information.

#### Modifying the Fabric Topology Input File

Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information on how to create and modify a topology file describing the fabric.

An example of the topology input file (in XML format) is shown below:

```
<?xml version="1.0" encoding="utf-8" ?>
<Report date="day mmm dd hh:mm:ss yyyy" unixtime="1446650124" options="-o
topology" >
<Nodes>
  <FIs>
    <ConnectedFICount>2</ConnectedFICount>
    <Node id="0x00117501007067a2">
      <NodeGUID>0x00117501007067a2</NodeGUID>
      <NodeType>FI</NodeType>
      <NodeType_Int>1</NodeType_Int>
      <NodeDesc>mindy2 hfi-0</NodeDesc>
      <Port id="0x00117501007067a2:1">
        <PortNum>1</PortNum>
        <LID>0x0001</LID>
        <PortGUID>0x00117501007067a2</PortGUID>
        <LinkWidthActive>4</LinkWidthActive>
        <LinkWidthActive_Int>8</LinkWidthActive_Int>
        <LinkSpeedActive>25Gb</LinkSpeedActive>
        <LinkSpeedActive_Int>2</LinkSpeedActive_Int>
      </Port>
    </Node>
    <Node id="0x00117501007067e6">
      <NodeGUID>0x00117501007067e6</NodeGUID>
      <NodeType>FI</NodeType>
      <NodeType_Int>1</NodeType_Int>
      <NodeDesc>mindy2 hfi-0</NodeDesc>
      <Port id="0x00117501007067e6:1">
        <PortNum>1</PortNum>
```



```

        <LID>0x0002</LID>
        <PortGUID>0x00117501007067e6</PortGUID>
        <LinkWidthActive>4</LinkWidthActive>
        <LinkWidthActive_Int>8</LinkWidthActive_Int>
        <LinkSpeedActive>25Gb</LinkSpeedActive>
        <LinkSpeedActive_Int>2</LinkSpeedActive_Int>
    </Port>
</Node>
</FIs>
<Switches>
    <ConnectedSwitchCount>0</ConnectedSwitchCount>
</Switches>
<SMs>
    <ConnectedSMCount>1</ConnectedSMCount>
    <SM id="0x00117501007067a2:1">
        <SMState>Master</SMState>
        <SMState_Int>3</SMState_Int>
        <NodeGUID>0x00117501007067a2</NodeGUID>
        <NodeDesc>mindy2 hfi-0</NodeDesc>
        <PortNum>1</PortNum>
        <PortGUID>0x00117501007067a2</PortGUID>
        <NodeType>FI</NodeType>
        <NodeType_Int>1</NodeType_Int>
    </SM>
</SMs>
</Nodes>
<LinkSummary>
    <LinkCount>1</LinkCount>
    <Link id="0x00117501007067a2:1">
        <Rate>100g</Rate>
        <Rate_Int>16</Rate_Int>
        <Internal>0</Internal>
        <Port id="0x00117501007067a2:1">
            <NodeGUID>0x00117501007067a2</NodeGUID>
            <PortGUID>0x00117501007067a2</PortGUID>
            <PortNum>1</PortNum>
            <NodeType>FI</NodeType>
            <NodeType_Int>1</NodeType_Int>
            <NodeDesc>mindy2 hfi-0</NodeDesc>
        </Port>
        <Port id="0x00117501007067e6:1">
            <NodeGUID>0x00117501007067e6</NodeGUID>
            <PortGUID>0x00117501007067e6</PortGUID>
            <PortNum>1</PortNum>
            <NodeType>FI</NodeType>
            <NodeType_Int>1</NodeType_Int>
            <NodeDesc>mindy3 hfi-0</NodeDesc>
        </Port>
    </Link>
</LinkSummary>
</Report>

```



## 4.0 Managing the Chassis Configuration

The FastFabric OPA Chassis Setup/Admin menu allows you to set up and manage the Intel® Omni-Path Architecture internally-managed switches.

1. Log in to the server as root.
2. At the command prompt, enter **opafastfabric**.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

  1) Chassis Setup/Admin
  2) Externally Managed Switch Setup/Admin
  3) Host Setup
  4) Host Verification/Admin
  5) Fabric Monitoring

X) Exit
```

3. Type **1**.

The **FastFabric OPA Chassis Setup/Admin Menu** is displayed.


```
FastFabric OPA Chassis Setup/Admin Menu
Chassis File: /etc/sysconfig/opa/chassis

Setup:
0) Edit Config and Select/Edit Chassis File [ Skip ]
1) Verify Chassis via Ethernet Ping [ Skip ]
2) Update Chassis Firmware [ Skip ]
3) Set Up Chassis Basic Configuration [ Skip ]
4) Set Up Password-less ssh/scp [ Skip ]
5) Reboot Chassis [ Skip ]
6) Get Basic Chassis Configuration [ Skip ]
7) Configure Chassis Fabric Manager (FM) [ Skip ]
8) Update Chassis FM Security Files [ Skip ]
9) Get Chassis FM Security Files [ Skip ]
Admin:
a) Check OPA Fabric Status [ Skip ]
b) Control Chassis Fabric Manager (FM) [ Skip ]
c) Generate All Chassis Problem Report Info [ Skip ]
d) Run a Command on All Chassis [ Skip ]
Review:
e) View opachassisadmin Result Files [ Skip ]

P) Perform the Selected Actions N) Select None
X) Return to Previous Menu (or ESC)
```

4. Select one or more items by typing the alphanumeric character associated with the item to toggle the selection from **Skip** to **Perform**.
5. Type **P** to perform the operations.

*Note:* Each menu item will present you with prompts to complete the operation.

**Table 4. FastFabric OPA Chassis Setup/Admin Menu Descriptions**

Menu Item	Description
0) Edit the Configuration and Select/ Edit Chassis File	<p><b>(Switch)</b> Allows you to edit the following configuration files:</p> <ul style="list-style-type: none"> <li>• <code>/etc/sysconfig/opa/chassis</code> The chassis file lists the internally-managed Intel switching chassis.</li> <li>• <code>/etc/sysconfig/opa/ports</code> The ports file lists the local HFI ports (for example, subnets) to be used to access the fabric for analysis.</li> <li>• <code>/etc/sysconfig/opa/opafastfabric.conf</code> The opafastfabric.conf file lists the default settings for most of the FastFabric command line options</li> </ul> <p>Associated CLI Command: <code>opagenchassis</code></p>
1) Verify Chassis via Ethernet Ping	<p><b>(Switch)</b> Allows you to verify the existence of each selected chassis listed in the chassis file using a ping over the management network.</p> <p>Associated CLI Command: <code>opapingall -C -p -F</code></p>
2) Update Chassis Firmware	<p><b>(Switch)</b> Allows you to verify and update the chassis firmware version.</p> <p>Associated CLI Command: <code>opachassisadmin update</code></p>
3) Set Up Chassis Basic Configuration	<p><b>(Switch)</b> Prompts you for chassis configuration settings and then configures all the selected chassis accordingly.</p> <p>Associated CLI Command: <code>opachassisadmin configure</code></p>
4) Set Up Password-less ssh/scp	<p><b>(Switch)</b> Allows you to set up secure password-less SSH so that the Fabric Management Node can securely log into all the other chassis as admin through the management network without requiring a password.</p> <p>Associated CLI Command: <code>opasetup_ssh -p -S -C -F chassisfile</code></p>
5) Reboot Chassis	<p><b>(Switch)</b> Allows you to reboot each chassis listed in the <code>/etc/sysconfig/opa/chassis</code> file, ensuring that each chassis reboot is successful (as verified using ping over the management network).</p> <p>Associated CLI Command: <code>opachassisadmin -S -F chassisfile reboot</code></p>
6) Get Basic Chassis Configuration	<p><b>(Switch)</b> Allows you to retrieve basic information from the chassis, such as</p> <ul style="list-style-type: none"> <li>• Syslog</li> <li>• NTP configuration</li> <li>• Time zone information</li> <li>• Link Width</li> <li>• Link CRC Mode</li> <li>• Node description</li> </ul> <p>Associated CLI Command: <code>opachassisadmin -F /etc/sysconfig/opa/chassis getconfig</code></p>
7) Configure Chassis Fabric Manager (FM)	<p><b>(Switch)</b> Assists you in configuring the Intel® Omni-Path Fabric Suite Fabric Manager (FM) for any member of the Intel® Omni-Path Chassis 100 Series.</p> <p>Associated FM Command: <code>config_generate</code></p> <p>Associated CLI Command: <code>opachassisadmin fmconfig</code></p>
8) Update Chassis FM Security Files	<p><b>(Switch)</b> Allows you to verify and update the chassis security files.</p> <p>Associated CLI Command: <code>opachassisadmin fmsecurityfiles</code></p>
9) Get Chassis FM Security Files	<p><b>(Switch)</b> Allows you to retrieve the chassis FM security files from the chassis.</p> <p>Associated CLI Command: <code>opachassisadmin fmgetsecurityfiles</code></p>
a) Check OPA Fabric Status	<p><b>(Switch or All)</b> Allows you to check the state and error counts of all ports.</p>

*continued...*



Menu Item	Description
	Associated CLI Commands: opalinkanalysis, opareport -o errors, opareport -o slowlinks, opareport -o misconnlinks, opashowallports -C
b) Control Chassis Fabric Manager (FM)	<b>(Switch)</b> Assists you in controlling the FM for any Intel® Omni-Path Chassis 100 Series chassis. <b>NOTE:</b> This operation is skipped for other chassis models.
c) Generate All Chassis Problem Report Info	<b>(Switch)</b> Allows you to collect configuration and status information from all selected chassis and generates a single *.tgz file that can be sent to a support representative. Associated CLI Command: opacaptureall -C
d) Run a Command on All Chassis	<b>(Switch)</b> Allows you to execute a CLI command against all selected chassis. Associated CLI Command: opacmdall -C
e) View opachassisadmin Result Files	<b>(All)</b> Allows you to view the test.log and test.res files, which reflect the results from opachassisadmin operations (such as for updating Chassis Firmware or rebooting all chassis).

## 4.1 Editing the Configuration Files for Chassis Setup

**(Switch)** The **Edit Config and Select/Edit Chassis File** selection allows you to select and edit the chassis, ports, and FastFabric configuration files.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **0**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Chassis Admin: Edit Config and Select/Edit Chassis File
Using vi (to select a different editor, export EDITOR).
You will now have a chance to edit/review the FastFabric Config File:
/etc/sysconfig/opa/opafastfabric.conf
The values in this file will control the default operation of the
FastFabric Tools. With the exception of the host file to use,
the values you specify for defaults will be used for all FastFabric
Operations performed via this menu system
Beware existing environment variables will override the values in this file.

About to: vi /etc/sysconfig/opa/opafastfabric.conf
Hit any key to continue (or ESC to abort)...
```

3. Press any key to **open the opafastfabric.conf file** or **ESC** to abort the operation.

*Note:* To get to subsequent configuration files, you must access each file.

The configuration file opens.

4. Review the settings:
  - a. Review the `FF_CHASSIS_LOGIN_METHOD` and `FF_CHASSIS_ADMIN_PASSWORD`.



- FastFabric provides the opportunity to enter the chassis password interactively when needed. It is not necessary to place it within `opafastfabric.conf`. If the Intel chassis admin password is placed in `opafastfabric.conf`, change the `opafastfabric.conf` permissions to be 0x600 (root-only access).
  - All versions of Intel® Omni-Path Chassis 100 Series firmware permit SSH keys to be configured within the chassis for secure password-less login. There is no need to configure a `FF_CHASSIS_ADMIN_PASSWORD`, and `FF_CHASSIS_LOGIN_METHOD` can be set to `SSH` (the default).
- b. Select the location for the result files from FastFabric with the `FF_RESULT_DIR` parameter. The default is the directory from which a given session of FastFabric is invoked. Alternatively, it can be set to a directory relative to your home directory. For example:

```
export FF_RESULT_DIR=${FF_RESULT_DIR:-$HOME/
fastfabric_results}
```

Refer to [FastFabric Configuration File](#) on page 35 for more information.

5. After saving and closing the `opafastfabric.conf` file in the editor, you will be given the opportunity to edit the `ports` file.

```
You will now have a chance to edit/review the FastFabric PORTS_FILE:
/etc/sysconfig/opa/ports
Some of the FastFabric operations which follow will use this file to
specify the local HFI ports to use to access the fabric(s) to operate on
Beware existing environment variables will override the values in this file.

About to: vi /etc/sysconfig/opa/ports
Hit any key to continue (or ESC to abort)...
```

6. Press any key to **open the ports file** or ESC to abort the operation.  
The configuration file opens.
7. Review the file:
  - For typical single-subnet clusters, the default of "0:0" may be used. This uses the first active port on the Management Node to access the fabric.
  - For configuring a cluster with multiple subnets, refer to *Intel® Omni-Path Fabric Software Installation Guide*.

Refer to [Ports List Configuration File](#) on page 35 for more information.

For further details about the Port List File format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

8. After saving and closing the `ports` file in the editor, you will be given the opportunity to select the `chassis` file.

```
The FastFabric operations which follow will require a file
listing the chassis to operate on
Select Chassis File to Use/Edit [/etc/sysconfig/opa/chassis]:
```

9. Press **Enter** to edit the file.

```
About to: vi /etc/sysconfig/opa/chassis
Hit any key to continue (or ESC to abort)...
```



10. Press any key to **open the chassis file** or **ESC** to abort the operation.

The configuration file opens.

11. Create the file with a list of the chassis names (the TCP/IP Ethernet management port names assigned) or IP addresses.

*Note:* Intel recommends you use chassis names.

Enter one chassis name or IP address per line. For example:

```
Chassis1
Chassis2
```

*Note:* Do not list externally-managed switches in this file.

Refer to [Chassis List Configuration Files](#) on page 36 for more information.

For further details about the Chassis List File format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

12. After saving and closing the `chassis` file in the editor, you will be given the opportunity to review and change the configuration files again.

```
Selected Chassis File: /etc/sysconfig/opa/chassis
Do you want to edit/review/change the files? [y]:
```

13. Press **Enter** to review and edit the files or type `n` and press **Enter** to end the operation.

## 4.2 Verifying Chassis via Ethernet Ping

**(Switch)** The **Verify Chassis via Ethernet Ping** selection allows you to ping each selected chassis over the management network.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **1**.

The menu item changes from `[Skip]` to `[Perform]`.

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The status is displayed.

```
Performing Chassis Admin: Verify Chassis via Ethernet Ping
Executing: /usr/sbin/opapingall -C -p -F /etc/sysconfig/opa/chassis
10.228.208.245: is alive
Hit any key to continue (or ESC to abort)...
```

3. Press any key to continue or **ESC** to abort the operation.
4. If some chassis were not found, use the following list to assist in troubleshooting:
  - Is chassis powered on and booted?
  - Is chassis connected to management network?
  - Are chassis IP address and network settings consistent with DNS or `/etc/hosts`?
  - Is Management node connected to the management network?





- Are Management node IP address and network settings correct?
- Is management network itself up (including switches, routers, and others)?
- Is correct set of chassis listed in the chassis file? You may need to repeat the previous step to review and edit the file.

## 4.3 Updating the Chassis Firmware

**(Switch)** The **Update Chassis Firmware** selection allows you to verify and update the chassis firmware version as needed.

**PREREQUISITE:** Before updating the firmware, refer to the relevant switch release notes for any prerequisites:

- *Intel® Omni-Path Fabric Managed Switches Release Notes*
- *Intel® Omni-Path Fabric Externally-Managed Switches Release Notes*

1. From the FastFabric OPA Chassis Setup/Admin menu, type **2**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Chassis Admin: Update Chassis Firmware
Multiple Firmware files and/or Directories may be space separated
Shell wildcards may be used
For Directories all .dpkg or .spkg files in the directory tree will be used
Enter Files/Directories to use (or none):
```

3. Specify the directory where the relevant firmware files have been stored and press **Enter**.

This can be the mount point of the CD or the directory to which the files were copied in a previous step.

```
Would you like to run the firmware now? [n]:
```

4. Type **y** and press **Enter**.

FastFabric ensures that all chassis are running the firmware level provided, and installs and/or reboots each chassis as needed.

If any chassis fails to be updated, use the **View opachassisadmin Result Files** option to review the result files from the update. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more details.

## 4.4 Setting Up Chassis Basic Configuration

**(Switch)** The **Setup Chassis Basic Configuration** allows you to perform the typical chassis setup operations for all chassis.

*Note:* First-time installation instructions are found the *Intel® Omni-Path Fabric Software Installation Guide*.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **3**.



The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Would you like to be prompted for chassis' password? [n]:	Allows you to enter a password for the chassis.
Do you wish to adjust syslog configuration settings? [y]:	Allows you to modify the syslog configuration settings.
Enter IP address for syslog server (or none):	Allows you to enter the IP address to the syslog server.
Do you wish to configure the syslog TCP/UDP port number? [n]:	Allows you to enter the TCP/UDP port number.
Do you wish to configure the syslog facility? [n]:	Allows you to configure the syslog facility.
Do you wish to configure an NTP server? [y]:	Allows you to configure an NTP server.
Enter IP address for NTP server (or none):	Allows you to set up the NTP server IP address.
Do you wish to configure timezone and DST information? [y]:	Allows you to set timezone and DST information from local server or manually.
Do you want to use the local timezone information from the local server? [y]:	Allows you to synchronize local timezone information with a local server.
Do you wish to configure the chassis link width? [n]:	Allows you to configure chassis link width.
Do you wish to configure OPA Node Desc to match ethernet chassis name? [y]:	Allows you to set OPA Node Desc to match ethernet chassis name. If you select yes [y], a reboot of all chassis devices is required in order to activate changes to the chassis OPA Node Desc.
Do you wish to configure the Link CRC Mode? [n]:	Allows you to configure link CRC mode.

After executing the prompts, the following is displayed:

```
Executing configure Test Suite (configure) Tue Oct 04 15:43:00 EDT 2016 ...
Executing TEST SUITE configure CASE (configure.10.228.208.245.emb.configure)
Chassis 10.228.208.245 configure ...
TEST SUITE configure CASE (configure.10.228.208.245.emb.configure) Chassis
10.228.208.245 configure PASSED
TEST SUITE configure: 1 Cases; 1 PASSED; 0 FAILED
TEST SUITE configure PASSED
Done configure Test Suite Tue Oct 04 15:43:02 EDT 2016

Hit any key to continue (or ESC to abort)...
```

4. Press any key or ESC to end the operation



## 4.5 Setting Up Password-less ssh/scp

**(Switch)** The **Set up Password-Less SSH/SCP** selection sets up secure password-less SSH, such that the Management Node can securely log into all the chassis as admin through the management network, without requiring a password.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **4**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **␣** to begin the operation.

```
Would you like to override the default Chassis password? [n]:
```

3. Choose one of the following actions:

- Press **Enter** to accept the default Chassis password.

```
Default Chassis password will be used to perform the setup
Executing: /usr/sbin/opasetupssh -p -C -F /etc/sysconfig/opa/chassis
Configuring 10.228.208.245...
Successfully processed: 1
Hit any key to continue (or ESC to abort)...
```

- Type **y** and press **Enter** to configure a new password for all chassis.

```
Executing: /usr/sbin/opasetupssh -p -S -C -F /etc/sysconfig/opa/chassis
Password for admin on all chassis:
```

- Enter the new password and press **Enter**.

## 4.6 Rebooting the Chassis

**(Switch)** The **Reboot Chassis** selection allows you to reboot all the selected chassis and ensures that they reboot fully (as verified through ping over the management network). When the chassis come back up following the reboot, they are running with all the new configuration settings.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **5**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **␣** to begin the operation.

```
Performing Chassis Admin: Reboot Chassis
Would you like to be prompted for chassis' password? [n]:
```

3. Press **Enter** to accept the default.

The chassis reboots.



## 4.7 Getting Basic Chassis Configuration

**(Switch)** The **Get Basic Chassis Configuration** selection allows you to retrieve basic information from chassis such as syslog, NTP configuration, time zone, node description, and other information.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **6**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The status is displayed as shown in the example below.

```
Performing Chassis Admin: Get basic Chassis configuration
Executing: /usr/sbin/opachassisadmin -F /etc/sysconfig/opa/chassis getconfig
Executing getconfig Test Suite (getconfig) day mmm dd hh:mm:ss timezone
YYYY ...
Executing TEST SUITE getconfig CASE (getconfig.xx.xx.xx.getconfig) get
xx.xx.xx.xx ...
TEST SUITE getconfig CASE (getconfig.xx.xx.xx.getconfig) get xx.xx.xx.xx
xx.xx.xx.xx:
    Firmware Active       : xx.xx.xx.xx
    Firmware Primary      : xx.xx.xx.xx
    Syslog Configuration  : Syslog host set to: 0.0.0.0 port 514 facility 22
    NTP                   : Configured to use the local clock
    Time Zone             : Time zone offset has not been configured
    LinkWidth Support     : 4X
    Node Description      : Node_Name
    Link CRC Mode         : 48b_or_14b_or_16b
PASSED
TEST SUITE getconfig: 1 Cases; 1 PASSED
TEST SUITE getconfig PASSED
Done getconfig Test Suite day mmm dd hh:mm:ss timezone yyyy

Hit any key to continue (or ESC to abort)...
```

3. Press any key or ESC to end the operation.

## 4.8 Configuring Chassis Fabric Manager

**(Switch)** The **Configure Chassis Fabric Manager (FM)** selection allows you to configure the Fabric Manager for any Intel® Omni-Path Chassis 100 Series.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **7**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Chassis Admin: Configure Chassis Fabric Manager (FM)
Enter FM Config file to use (or none or generate):
```

3. Type **generate**.



This performs the `config_generate` operation to guide you through selecting FM configuration options. See the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* for more information about `config_generate`.

4. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Anticipated maximum fabric size [100]:	Allows you to set the size of the subnet. FM resources and buffering are scaled to match the anticipated maximum size of the fabric. The size is specified in terms of the number of HFIs in a single fabric. For Embedded Fabric Manager, its recommended to use a value of 100 or less.
LMC value to use (there will be 2^LMC LIDs per HFI) [0]:	Allows you to set LMC value. LMC is used to control the number of LIDs per HFI. Multiple LIDs can be used to permit multiple routes between endpoints. This permits selected applications (such as MPIs using Intel(R) PSM) to optimize performance and/or resiliency by using dispersive routing. Default 0 assigns 1 LID per HFI.
Should Adaptive Routing be enabled [n]:	Allows you to enable or disable adaptive routing. Adaptive routing permits Intel® Omni-Path Architecture switches to dynamically adjust routing based on traffic patterns and hence reduce congestion and improve overall cluster performance and efficiency.
Node Appearance Log Message Threshold [100]:	Allows you to set the number of node appearance messages per sweep. When nodes appear or disappear from the fabric, a message is logged. A Threshold can be configured to limit the number of such messages per sweep. This Threshold can help to avoid excessive messages when fabric changes occur.
Name for FM instance 0 (Switch Port 0) [fm0]:	Allows you to set a name for each FM.
IPoIB rate for this FM (25g recommended):	Allows you to set the IPoIB rate for the FM. The FM configures the rate and MTU used for IPoIB multicast. The rate selected must be no greater than the rate of the slowest link in the fabric(s). The MTU selected must be no greater than the MTU of the smallest MTU link in the fabric(s). When selecting the rate and MTU, HFIs which won't run IPoIB can be ignored. However all Switches must be operating with at least the rate and MTU selected. Values are: 1) 25g 2) 50g 3) 75g 4) 100g
IPoIB MTU for this FM (2048 recommended):	Allows you to set the IPoIB MTU for the FM. Values are: 1) 2048 2) 4096
Do you want to configure a preferred primary or secondary FM [n]:	Allows you to configure primary or secondary FM. The FM supports failover. The FM to be preferred as the primary can be selected per FM instance. If no preferred primary is selected, FMs will negotiate based on HFI GUIDs.
Will this FM be the preferred primary [y]:	Allows you to set the current FM to be the primary FM. Values are: y - Primary n - Secondary
Should Sticky Failover be enabled [n]:	Allows you to set up the FM to support sticky failover.
continued...	



Prompt	Description
	The FM supports sticky failover. When enabled sticky failover will prevent a master FM from relinquishing control even if the preferred primary FM comes online. This can prevent situations where a bouncing preferred primary repeatedly takes over then fails.
Subnet Prefix upper bits for cluster [0xfe80000000000000]:	Allows you to set to the subnet prefix upper bits for the cluster.  Each fabric in a cluster must have a unique 64 bit subnet prefix. The subnet prefix must be consistently configured on all FMs which manage the given fabric (e.g., on the primary and secondaries).  To simplify input, you will be prompted for the upper bits for the cluster, then you will be prompted for the lower bits for each instance. The two values will be OR'ed together to form the subnet prefix for each fabric.
Subnet Prefix lower bits for FM instance 0 (fm0) (Switch Port 0) [0x0]:	Allows you to set to the subnet prefix lower bits for the cluster.
PM Sweep Interval in seconds [10]:	Allows you to set the PM sweep interval.  The Fabric Manager includes a Performance Manager (PM) which can monitor the data movement and error counters in all devices. The PM monitors the counters periodically and computes the delta for counters. If the PM Sweep Interval is set to 0, no automatic sweeps occur. The PM Sweep Interval must be > 0 when using tools such as Fabric Performance Monitoring (opatop).
PM Error Threshold Exceeded Log Message Limit [10]:	Allows you to limit the number of PM error messages per sweep.  When a port exceeds the threshold for Integrity, Security, or Routing errors, a message is logged. A Threshold can be configured to limit the number of such messages per sweep. This Threshold can help to avoid excessive messages.
How many concurrent clients are expected? [3]:	Allows you to set the number of concurrent PM clients to expect.  The PM can retain some recent history in memory. This history can then be viewed in tools such as Fabric Performance Monitoring (opatop). For each historical sweep, both the topology and performance data is retained. Each dataset is referred to as an "image". The values will be adjusted based on the number of concurrent PA clients expected.
How many images should be retained? [10]:	Allows you to set the number of images to be retained for history. Images include: Pm.TotalImages Pm.FreezeFrameImages

After executing the prompts, the following is displayed:

```
Generated ./opafm.xml
To activate this configuration, ./opafm.xml must be transfered to
the chassis and the FM must be restarted.
The fastfabric TUI provides an easy way to do this.
You have selected to use: ./opafm.xml
Syntax Checking ./opafm.xml...
Executing: /usr/lib/opa/fm_tools/config_check -s -c ./opafm.xml
Valid FM Config file: ./opafm.xml

After push, the FM may be started/restarted
Would you like to restart the FM? [n]:
```

5. Enter **y**.

This causes the FM to be started with the new configuration.

```
Would you like to run the FM on slave MMs? [n]:
```

## 6. Refer to the following If/Then table:

If	Then
Your fabric has a single chassis running the Fabric Manager. You can run the Fabric Manager on the slave management module (MM). This causes the Fabric Manager to be started in the applicable chassis.	Enter <b>y</b>
Your fabric has multiple chassis running the Fabric Manager. Intel recommends you run Fabric Manager on the master management module. This causes the Fabric Manager to be started only on the master management module in the applicable chassis.	Enter <b>n</b>

```
There will be a disruption as FMs are restarted
Doing the operation in parallel (on multiple chassis) will finish the fastest
Doing it serially may reduce disruption
Would you like to do the operation in parallel? [y]:
```

7. Press **Enter** to select the default option: **y**.

Intel recommends doing the operation in parallel.

```
You have selected to perform the push and FM restart in parallel
Would you like to enable FM start at boot? [n]:
```

8. Enter **y**.

This causes the Fabric Manager to be started on all applicable chassis each time those chassis boot.

```
Would you like to enable FM start on slave MMs at boot? [n]:
```

## 9. Refer to the following If/Then table:

If	Then
Your fabric has a single chassis running the Fabric Manager. You can run the Fabric Manager on the slave management module. This causes the Fabric Manager to be started in the applicable chassis.	Enter <b>y</b>
Your fabric has multiple chassis running the Fabric Manager. Intel recommends you run Fabric Manager on the master management module. This causes the Fabric Manager to only be started on the master management module in the applicable chassis.	Enter <b>n</b>

System prompts:

```
Would you like to be prompted for chassis' password? [n]:
```

10. Press **Enter** to select the default **n** option.

```
Are you sure you want to proceed? [n]:
```

11. Enter **y**.



This updates the Fabric Manager.

```
Hit any key to continue (or ESC to abort)...
```

12. Press any key or ESC to end the operation.

## 4.9 Updating the Chassis FM Security Files

**(Switch)** The **Update Chassis FM Security Files** selection allows you to verify and update the chassis security files.

**Note:** The FM security files are the private key, public key, and certificate files required by the FM, in order to support secure socket connection via OpenSSL. Refer to the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* for instructions on the administration tasks required to support these files.

1. From the FastFabric OPA Chassis Setup/Admin menu, type 8.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Chassis Admin: Update Chassis FM Security Files
Multiple FM Security files and/or Directories may be space separated
Shell wildcards may be used
For Directories all .pem files in the directory tree will be used
Enter Files/Directories to use (or none):
```

3. Enter the files/directories and press **Enter**.

4. For each subsequent prompt, provide the required information and press **Enter**:

Prompts guide you through the options:

- `push` - Ensures given security files are pushed to each chassis.
- `restart` - After push, restart FM on master, stop on slave.
- `restartall` - After push, restart FM on all MM.

Additional options prompted for:

- Selection of security files or directory containing pem files
- Parallel vs serial update
- Chassis password (default is to have password in `fastfabric.conf` or to use password-less SSH)

If any chassis fails to be updated, use the `View opachassisadmin results files` selection to review the result files from the update. Refer to [Viewing opachassisadmin Result Files](#) on page 60 for more information.

Refer to *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more details on the `opachassisadmin update` command.





## 4.10 Getting Chassis FM Security Files

**(Switch)** The **Get Chassis FM Security Files** selection allows you to run the `opachassisadmin fmgetsecurityfiles` command to retrieve the chassis FM security files from the chassis.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **9**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The status is displayed.

```
Performing Chassis Admin: Get Chassis FM Security Files
Executing: /usr/sbin/opachassisadmin -F /etc/sysconfig/opa/chassis
fmgetsecurityfiles
Executing get FM security files Test Suite (fmgetsecurityfiles) Tue Oct 04
16:27:55 EDT 2016 ...
Executing TEST SUITE get FM security files CASE (fmgetsecurityfiles.
10.228.208.245.fm_get_security_files) get 10.228.208.245 *.pem ./uploads/
10.228.208.245/ ...
TEST SUITE get FM security files CASE (fmgetsecurityfiles.
10.228.208.245.fm_get_security_files) get 10.228.208.245 *.pem ./uploads/
10.228.208.245/ PASSED
TEST SUITE get FM security files: 1 Cases; 1 PASSED; 0 FAILED
TEST SUITE get FM security files PASSED
Done get FM security files Test Suite Tue Oct 04 16:27:57 EDT 2016

Hit any key to continue (or ESC to abort)...
```

3. Press any key to complete this procedure.

## 4.11 Checking the OPA Fabric Status

The **Check OPA Fabric Status** selection allows you to analyze the links in a fabric.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **a**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Would you like to perform fabric error analysis? [y]:	Allows you to start the analysis.
Clear error counters after generating report? [n]:	Allows you to clear the error counters after generating the report.
Would you like to perform fabric link speed error analysis? [y]:	Allows you to analyze fabric link speed errors.
continued...	



Prompt	Description
Check for links configured to run slower than supported? [n]:	Allows you to check for Links running slower than expected.
Check for links connected with mismatched speed potential? [n]:	Allows you to check for links connected with mismatched speed.
Enter filename for results [/root/linkanalysis.res]:	Allows you to enter a filename for the results or use the default file.

After executing the prompts, the following is displayed:

```
Executing: /usr/sbin/opalinkanalysis errors slowlinks > /root/
linkanalysis.res 2>&l
About to: vi /root/linkanalysis.res
Hit any key to continue (or ESC to abort)...
```

4. Press any key to view the results file in the editor.

An example output is shown below.

```
Links running slower than expected Summary

Links running slower than expected:
1 of 1 Links Checked, 0 Errors found
-----
-
Links with errors >= threshold Summary

Configured Thresholds:
  LinkQualityIndicator           4
  UncorrectableErrors           1
  LinkDowned                     1
  RcvErrors                     1
  ExcessiveBufferOverruns       1
  FMConfigErrors                1
1 of 1 Links Checked, 0 Errors found
-----
-
```

## 4.12 Controlling Chassis Fabric Manager

**(Switch)** The **Control Chassis Fabric Manager (FM)** selection allows you to control the FM.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **b**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Would you like to restart the FM? [n]:	Allows you to restart the FM.
Would you like to make sure the FM is not running? [n]:	Allows you to ensure that the FM is not running.
continued...	



Prompt	Description
Would you like to make sure the FM is running? [n]:	Allows you to ensure that the FM is running.
Would you like to run FM on slave MMs? [n]:	Allows you to run FM on slave management modules.
Would you like to do the operation in parallel? [y]:	Allows you to perform operations in parallel (on multiple chassis). Doing the operation in parallel will finish the fastest.
Would you like to change FM boot state to enable FM start at boot? [n]:	Allows you to enable FM start on slave management modules at boot.
Would you like to change FM boot state to disable FM start at boot? [n]:	Allows you to disable FM start on slave management modules at boot.
Would you like to be prompted for chassis' password? [n]:	Allows you to be prompted for the chassis password.

After executing the prompts, the following is displayed:

```
Are you sure you want to proceed? [n]:
```

4. Select **y** to complete the operation.
5. When complete, press any key or **ESC** to end the operation.

## 4.13 Generating All Chassis Problem Report Information

**(Switch)** The **Generate All Chassis Problem Report Info** selection allows you to collect configuration and status information from all chassis and generate a single \*.tgz file that can be sent to an Intel support representative.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **c**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Chassis Admin: Generate All Chassis Problem Report Info
Would you like to be prompted for chassis' password? [n]:
```

3. Press **Enter** to select the default or **y** to prompt for password.

opacaptureall is initiated and results gathered in chassiscapture.all.tgz.

```
Executing: /usr/sbin/opacaptureall -C -p -F /etc/sysconfig/opa/chassis
Running capture on all chassis ...
admin@10.228.208.245: capture: Command execution PASSED (Login): ...
Combining captured files into ./uploads/chassiscapture.all.tgz ...
Done.
Hit any key to continue (or ESC to abort)...
```

4. When complete, press any key or **ESC** to end the operation.



## 4.14 Running a Command on All Chassis

**(Switch)** The **Run a Command on All Chassis** selection allows you to perform other operations on all chassis. Each time this is executed, a single chassis CLI command may be specified to be executed against all selected chassis. When using these commands, additional setup or verification of the chassis may be performed.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **d**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Chassis Admin: Run a Command on All Chassis
Would you like to be prompted for chassis' password? [n]:
```

3. Press **Enter** to select the default or **y** to prompt for password.

```
Enter Command to run on all chassis (or none):
```

4. Enter the CLI command to run and press **Enter**.

```
Run in parallel on all chassis? [y]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.

```
About to run: /usr/sbin/opacmdall -C -F /etc/sysconfig/opa/chassis
'opashowmc -v'
Are you sure you want to proceed? [n]:
```

6. Select **y** (yes) or **n** (no) and press **Enter**.

```
Executing: /usr/sbin/opacmdall -C -F /etc/sysconfig/opa/chassis 'opashowmc -v'
[admin@10.228.208.245]# opashowmc -v
admin@10.228.208.245: opashowmc -v: Command execution ...
Hit any key to continue (or ESC to abort)...
```

7. When complete, press any key or **ESC** to end the operation.

## 4.15 Viewing opachassisadmin Result Files

**(Switch)** The **View opachassisadmin Result Files** selection allows you to open and view the punchlist.csv, test.res, and test.log files.

1. From the FastFabric OPA Chassis Setup/Admin menu, type **e**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.



2. Type **P** to begin the operation.

```
Performing Chassis Admin: View opachassisadmin Result Files
Using vi (to select a different editor, export EDITOR).
About to: vi /root/punchlist.csv /root/test.res /root/test.log
Hit any key to continue (or ESC to abort)...
```

3. Press any key to view the opashassisadmin results files.
4. After reviewing and closing the log, you are prompted to remove the following files.

```
3 files to edit
Would you like to remove test.res test.log test_tmp* and save_tmp
in /root ? [n]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.
6. If you chose **y** in the step above, press any key or **ESC** to end the operation



## 5.0 Managing the Switch Configuration

The FastFabric OPA Switch Setup/Admin menu allows you to set up and manage the Intel® Omni-Path externally-managed Edge Switches.

1. Log in to the server as root.
2. At the command prompt, enter **opafastfabric**.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

  1) Chassis Setup/Admin
  2) Externally Managed Switch Setup/Admin
  3) Host Setup
  4) Host Verification/Admin
  5) Fabric Monitoring

X) Exit
```

3. Type **2**.


The **FastFabric OPA Switch Setup/Admin Menu** is displayed.

```
FastFabric OPA Switch Setup/Admin Menu
Externally Managed Switch File: /etc/sysconfig/opa/switches
Setup:
0) Edit Config and Select/Edit Switch File      [ Skip ]
1) Generate or Update Switch File               [ Skip ]
2) Test for Switch Presence                     [ Skip ]
3) Verify Switch Firmware                      [ Skip ]
4) Update Switch Firmware                      [ Skip ]
5) Set Up Switch Basic Configuration            [ Skip ]
6) Reboot Switch                              [ Skip ]
7) Report Switch Firmware & Hardware Info      [ Skip ]
8) Get Basic Switch Configuration              [ Skip ]
Admin:
9) Report Switch VPD Information                [ Skip ]
Review:
a) View opaswitchadmin Result Files            [ Skip ]

P) Perform the Selected Actions                 N) Select None
X) Return to Previous Menu (or ESC)
```

4. Select one or more items by typing the alphanumeric character associated with the item to toggle the selection from **Skip** to **Perform**.
5. Type **P** to perform the operations.

**Note:** Each menu item will present you with prompts to complete the operation.

**Table 5. FastFabric OPA Switch Setup/Admin Menu Descriptions**

Menu Item	Description
0) Edit Config and Select/Edit Switch File	<p><b>(Switch)</b> Allow you to edit the following configuration files:</p> <ul style="list-style-type: none"> <li>• <code>/etc/sysconfig/opa/ports</code> The <code>ports</code> file lists the local HFI ports (for example, subnets) to be used to access the fabric for analysis.</li> <li>• <code>/etc/sysconfig/opa/switches</code> The <code>switches</code> file lists the externally-managed Intel® Omni-Path Switch 100 Series switches.</li> <li>• <code>/etc/sysconfig/opa/opafastfabric.conf</code> The <code>opafastfabric.conf</code> file lists the default settings for most of the FastFabric command line options.</li> </ul> <p>Associated CLI Command: <code>opagenswitches</code></p>
1) Generate or Update Switch File	<p><b>(Switch)</b> Allows you to generate or update the <code>/etc/sysconfig/opa/switches</code> file based on the <code>/etc/sysconfig/opa/topology.%P.xml</code> files.</p>
2) Test for Switch Presence	<p><b>(Switch)</b> Allows you to test for the presence of the selected switches in the fabric.</p> <p>Associated CLI Command: <code>opaswitchadmin ping</code></p>
3) Verify Switch Firmware	<p><b>(Switch)</b> Allows you to verify the integrity of the present firmware in the switch.</p> <p>If this operation fails prior to any switch reboots or power-offs of the switch, perform <code>Update Switch Firmware</code> to correct the firmware in the switch.</p> <p>Associated CLI Command: <code>opaswitchadmin fwverify</code></p>
4) Update Switch Firmware	<p><b>(Switch)</b> Allow you to update the switch firmware version and set the switch node name.</p> <p>Associated CLI Command: <code>opaswitchadmin upgrade</code></p>
5) Set Up Switch Basic Configuration	<p><b>(Switch)</b> Prompts you for switch configuration settings and then configures all the selected Intel® Omni-Path Edge Switch 100 Series externally managed switches accordingly.</p> <p>Associated CLI Command: <code>opaswitchadmin configure</code></p>
6) Reboot Switch	<p><b>(Switch)</b> The <code>Reboot Switch</code> selection runs the <code>opaswitchadmin reboot</code> command to reboot all the switches listed in the <code>/etc/sysconfig/opa/switches</code> file that was created in a previous step.</p> <p>Associated CLI Command: <code>opaswitchadmin reboot</code></p>
7) Report Switch Firmware & Hardware Info	<p><b>(Switch)</b> Provides you with a summary of the present state for all the selected switches.</p> <p>Associated CLI Command: <code>opaswitchadmin info</code></p>
8) Get Basic Switch Configuration	<p><b>(Switch)</b> Allows you to retrieve basic information from an externally-managed switch, such as:</p> <ul style="list-style-type: none"> <li>• MTU</li> <li>• VL Cap</li> <li>• Credit Distribution</li> <li>• Link Width</li> <li>• Link Speed</li> <li>• Node description</li> </ul> <p>Associated CLI Command: <code>opaswitchadmin -S -L /etc/sysconfig/opa/switches getconfig</code></p>
9) Report Switch VPD Information	<p><b>(Switch)</b> Provides you with the Virtual Product Data (VPD) for all the selected switches. This information is useful for inventory and asset control as well as to provide details about the product to customer support.</p>

*continued...*



Menu Item	Description
	Associated CLI Command: <code>opaswitchadmin hwypd</code>
a) View <code>opaswitchadmin</code> Result Files	Allows you to view the <code>test.log</code> and <code>test.res</code> files that reflect the results from <code>opaswitchadmin</code> runs (such as those for updating switch firmware, or for rebooting all switches per menu items above).

## 5.1 Editing the Configuration Files for Externally-Managed Switch Setup

**(Switch)** The **Edit Config and Select/Edit Switch File** selection allows you to select and edit the switches, ports, and FastFabric configuration files.

**Note:** Intel® Omni-Path Fabric Suite FastFabric is topology-aware when updating externally-managed switch firmware or resetting the switches. The update or restart starts at the switches farthest from the FastFabric node and then works toward the FastFabric node. This way, switches that are rebooted are not in the path between the FastFabric node and others that are being updated or reset.

1. From the FastFabric OPA Switch Setup/Admin menu, type **0**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Switch Admin: Edit Config and Select/Edit Switch File
Using vi (to select a different editor, export EDITOR).
You will now have a chance to edit/review the FastFabric Config File:
/etc/sysconfig/opa/opafastfabric.conf
The values in this file will control the default operation of the
FastFabric Tools. With the exception of the host file to use,
the values you specify for defaults will be used for all FastFabric
Operations performed via this menu system
Beware existing environment variables will override the values in this file.

About to: vi /etc/sysconfig/opa/opafastfabric.conf
Hit any key to continue (or ESC to abort)...
```

3. Press any key to **open the opafastfabric.conf file** or ESC to abort the operation.

**Note:** To get to subsequent configuration files, you must access each file.

The configuration file opens.

4. Review the settings.

Refer to [FastFabric Configuration File](#) on page 35 for more information.

5. After saving and closing the `opafastfabric.conf` file in the editor, you will be given the opportunity to edit the ports file.

```
You will now have a chance to edit/review the FastFabric PORTS_FILE:
/etc/sysconfig/opa/ports
Some of the FastFabric operations which follow will use this file to
specify the local HFI ports to use to access the fabric(s) to operate on
Beware existing environment variables will override the values in this file.
```





```
About to: vi /etc/sysconfig/opa/ports
Hit any key to continue (or ESC to abort)...
```

6. Press any key to **open the ports file** or ESC to abort the operation.

The configuration file opens.

7. Review the file:

- For typical single-subnet clusters, the default of "0:0" may be used. This uses the first active port on the Management Node to access all externally managed switches. .
- For configuring a cluster with multiple subnets, refer to *Intel® Omni-Path Fabric Software Installation Guide*.

Refer to [Ports List Configuration File](#) on page 35 for more information.

For further details about the Port List File format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

8. After saving and closing the `ports` file in the editor, you will be given the opportunity to select the `switches` file.

```
The FastFabric operations which follow will require a file
listing the externally managed switches to operate on
Select Switch File to Use/Edit [/etc/sysconfig/opa/switches]:
```

9. Press **Enter** to edit the file.

```
About to: vi /etc/sysconfig/opa/switches
Hit any key to continue (or ESC to abort)...
```

10. Press any key to **open the switches file** or ESC to abort the operation.

The configuration file opens.

11. Create the file with a list of the switch node GUID and required switch names.

Enter one switch node GUID and required switch name per line. Do not use any spaces before or after the comma separating the switch node GUID and the name, as shown in this example:

```
0x00117500d9000138,edge1
0x00117500d9000139,edge2
```

**Note:** Do not list internally-managed chassis in this file.

- Tips:**
- The **Generate or Update Switch File** menu item or `opagenswitches` may be used to generate a list of the externally-managed switches presently in the fabric. For example, when using the `vi` editor, the command `:r ! opagenswitches` may be used to add the output from this command to the file.
  - If needed, an SA query can be used to get a list of all switches. This includes both internally- and externally-managed switches. Consequently, the output must be edited to leave only the Intel externally managed switches. An example SA query is:

```
opasaquery -t sw -o nodeguid
```



Refer to [Externally-Managed Switch List Configuration File](#) on page 38 for more information.

For further details about the (externally-managed) Switch List File format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

12. After saving and closing the `switches` file in the editor, you will be given the opportunity to review and change the configuration files again.

```
Selected Externally Managed Switch File: /etc/sysconfig/opa/switches
Do you want to edit/review/change the files? [y]:
```

13. Press **Enter** to review and edit the files or type `n` and press **Enter** to end the operation.

## 5.2 Generating or Updating Switch File

**(Switch)** The **Generate or Update Switch File** selection generates or updates the `switches` file. It can also update switch names in the `switches` file by comparing the actual fabric to topology xml data.

## 5.3 Testing for Switch Presence

**(Switch)** The **Test for Switch Presence** selection allows you to verify that each externally-managed switch specified in the `switches` file can be accessed by the Management Node through the Fabric Network.

1. From the FastFabric OPA Switch Setup/Admin menu, type `2`.

The menu item changes from `[Skip]` to `[Perform]`.

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The Test Suite report for switch ping is performed.

```
Performing Switch Admin: Test for Switch Presence
Executing: /usr/sbin/opaswitchadmin -L /etc/sysconfig/opa/switches ping
Executing report switch ping Test Suite (switchping) Thu Oct 06 10:02:00 EDT
2016 ...
Executing TEST SUITE report switch ping CASE (switchping.
10.228.208.247.i2c.extmgd.switchping) ping switch 10.228.208.247 ...
TEST SUITE report switch ping CASE (switchping.
10.228.208.247.i2c.extmgd.switchping) ping switch 10.228.208.247
...
TEST SUITE report switch ping: 1 Cases; 1 PASSED; 0 FAILED
TEST SUITE report switch ping PASSED
Done report switch ping Test Suite Thu Oct 06 10:04:01 EDT 2016

Hit any key to continue (or ESC to abort)...
```

3. Press any key or `ESC` to end the operation.
4. If some switches were not found, use the following list to assist in troubleshooting:
  - Is switch powered on and booted?
  - Is switch connected to Intel® Omni-Path Fabric?



- Is Subnet Manager running?
- Is Management Node's Port active?
- Is Management Node connected to the correct Intel® Omni-Path Fabric?
- Is FM Switch LED activated on the switch port to which the Fabric Management node is connected?

For more information, refer to the "FM Switch" section in the *Intel® Omni-Path Fabric Switches Hardware Installation Guide*.

- Is the correct set of switches listed in the switches file?

You may need to perform the [Generating or Updating Switch File](#) on page 66 operation to review and edit the file.

## 5.4 Verifying Switch Firmware

**(Switch)** The **Verify Switch Firmware** selection allows you to check that each externally-managed switch is operational and that its firmware is valid and accessible.

1. From the FastFabric OPA Switch Setup/Admin menu, type 3.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The TEST SUITE report switch fwverify is performed.

*Note:* The operation may take several minutes to complete.

```
Performing Switch Admin: Verify Switch Firmware
Executing: /usr/sbin/opa-switch-admin -L /etc/sysconfig/opa/switches fwverify
Executing report switch fwverify Test Suite (switchfwverify) Thu Oct 06
10:22:50 EDT 2016 ...
Executing TEST SUITE report switch fwverify CASE (switchfwverify.
10.228.208.247.i2c.extmgd.switchfwverify) retrieve switch 10.228.208.247 ...
TEST SUITE report switch fwverify CASE (switchfwverify.
10.228.208.247.i2c.extmgd.switchfwverify) retrieve switch 10.228.208.247
...
TEST SUITE report switch fwverify: 1 Cases; 1 PASSED; 0 FAILED
TEST SUITE report switch fwverify PASSED
Done report switch fwverify Test Suite Thu Oct 06 10:26:55 EDT 2016

Hit any key to continue (or ESC to abort)...
```

3. Press any key or **ESC** to end the operation.

## 5.5 Updating Switch Firmware

**(Switch)** The **Update Switch Firmware** selection allows you to update the switch firmware version and set the switch node name.

*Note:* Refer to the *Intel® Omni-Path Fabric Externally-Managed Switches Release Notes* to ensure that any prerequisites for the upgrade to the new firmware level have been met prior to performing the upgrade through FastFabric.

1. From the FastFabric OPA Switch Setup/Admin menu, type 4.



The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **p** to begin the operation.

```
Performing Switch Admin: Update Switch Firmware
Multiple Firmware files and/or Directories may be space separated
Shell wildcards may be used
For Directories all .emfw files in the directory tree will be used
Enter Files/Directories to use (or none):
```

3. Specify the directory where the relevant firmware files are located.

```
After upgrade, the switch may be optionally rebooted.
Would you like to reboot the switch after the update? [n]:
```

4. Type **y**.

```
The firmware on the switch will be checked, and if the running version is the
same
as the version being used for the update, the update operation will be
skipped.
Would you like to override this check, and force the update to occur? [n]:
```

5. Press **Enter** to select default (n).

**Note:** The fabric is not yet operational.

```
You have selected to update the switch firmware and reboot.
There will be a disruption as switch or switches are rebooted.
Doing the operation in parallel (on multiple switches) will finish the
fastest.
Doing it serially may reduce disruption.
Would you like to do the operation in parallel? [y]:
```

**Note:** Because the Intel® Omni-Path Fabric itself is used to update externally-managed switches, updating multiple switches with the reboot option may disrupt parallel update operations. If there are not any selected externally-managed switches in the path from the Management Node to any other externally-managed switch, parallel operations can be established. For example, if the Management Node is connected directly to a core switch and externally-managed switches are only at the edges.

To control the order of the rebooting of externally-managed switches by FastFabric, refer to the [distance](#) option for the switches file in [Externally-Managed Switch List Configuration File](#) on page 38.

6. Press **Enter** for yes, or type **n** for no and press **Enter**.

**Note:** Be aware that non-parallel operation for a fabric with many externally managed switches can take a significant amount of time.

FastFabric updates the firmware on all switches and sets the node names, as per the switches file. Each switch is then rebooted.

If any switch fails to be updated, use the **View opaswitchadmin result files** option to review the result files from the update. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more details.



## 5.6 Setting Up Switch Basic Configuration

**(Switch)** The **Setup Switch Basic Configuration** selection allows you to perform typical switch setup operations using a wizard to configure all switches.

The following steps walk you through the procedure:

1. From the FastFabric OPA Switch Setup/Admin menu, type 5.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Switch Admin: Set Up Switch Basic Configuration
Executing: /usr/sbin/opaswitchadmin -L /etc/sysconfig/opa/switches configure
Do you wish to configure the switch Link Width Options? [n]:
```

3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Do you wish to configure the switch Link Width Options? [n]:	<ul style="list-style-type: none"> <li>• Selecting <b>n</b> (no) causes the default switch Link Width Options to be used for all switches. If switches have previously been manually configured for different switch Link Width Options, this option keeps the previously configured switch Link Width Options. See the <i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i> for more information.</li> <li>• Selecting <b>y</b> (yes) prompts for setting the switch link width supported setting for all ports on all switches.</li> </ul> <p><b>Note:</b> This operation is only applicable to Intel® Omni-Path Edge Switch 100 Series switches.</p>
Do you wish to configure the switch Node Description as it is set in the switches file? [n]:	<ul style="list-style-type: none"> <li>• Selecting <b>n</b> (no) causes the default switch Node Description on each switch to be used. If the switches have previously been manually configured for a customized switch Node Description, this option keeps the previously configured switch Node Descriptions. See the <i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i> for more information.</li> <li>• Selecting <b>y</b> (yes) causes the Node Description on each switch to be updated as specified by the switches file.</li> </ul> <p><b>Note:</b> Only node descriptions on Intel® Omni-Path Edge Switch 100 Series switches can be changed in this step.</p>
Do you wish to configure the switch FM Enabled option? [n]:	<ul style="list-style-type: none"> <li>• Selecting <b>n</b> (no) causes all of the externally-managed switch ports to stay FM disabled.</li> <li>• Selecting <b>y</b> (yes) prompts for setting the switch FM-enabled capability for all ports on all switches.</li> </ul> <p>Setting it to enabled allows the FM to be connected to any port on any externally managed switch.</p> <p>If this is not desired, then select the default for the answer (disabled) and set the desired ports on the externally-managed switch to be FM-enabled using the FM switch.</p> <p>Refer to the "FM Switch" section in the <i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i> to set the port to FM enabled.</p>

**continued...**



Prompt	Description
	<i>Note:</i> This operation is only applicable to Intel® Omni-Path Edge Switch 100 Series switches.
Do you wish to configure the switch Link CRC Mode? [n]:	<ul style="list-style-type: none"><li>Selecting <b>n</b> (no) causes all of the externally managed switch ports Link CRC Mode to stay disabled.</li><li>Selecting <b>y</b> (yes) prompts for setting the link CRC Mode for all ports on all switches.</li></ul> <i>Refer to the Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide for more information.</i>

After executing the prompts, you will be notified whether the operation passed or failed.

```
Executing configure Test Suite (configure) Thu Oct 06 16:21:04 EDT 2016 ...
Executing TEST SUITE configure CASE (configure.
10.228.208.247.i2c.extmgd.switchconfigure) configure switch 10.228.208.247 ...
TEST SUITE configure: 1 Cases; 1 PASSED; 0 FAILED
TEST SUITE configure PASSED
Done configure Test Suite Tue Oct 04 15:43:02 EDT 2016

Hit any key to continue (or ESC to abort)...
```

4. Press any key or ESC to end the operation

## 5.7 Rebooting the Switch

**(Switch)** The **Reboot Switch** selection allows you to reboot all switches, ensuring that all the configuration changes become effective and are discovered by the Intel® Omni-Path Fabric Suite Fabric Manager.

1. From the FastFabric OPA Switch Setup/Admin menu, type **6**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

*Note:* Reboot begins immediately.

```
Performing Switch Admin: Reboot Switch
Executing: /usr/sbin/opaswitchadmin -L /etc/sysconfig/opa/switches reboot
Executing reboot Test Suite (reboot) Thu Oct 06 16:37:30 EDT 2016 ...
Executing TEST SUITE reboot CASE (reboot.10.228.208.247.i2c.extmgd.reset)
reset switch 10.228.208.247 ...
```

## 5.8 Reporting Switch Firmware and Hardware Information

**(Switch)** The **Report Switch Firmware & Hardware Info** selection allows you to review reports on the firmware and hardware versions for each switch, along with other information for all of the externally-managed switches. Review the results against the expected models and firmware versions.

1. From the FastFabric OPA Switch Setup/Admin menu, type **7**.

The menu item changes from [Skip] to [Perform].



**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Switch Admin: Report Switch firmware & hardware info
Executing: /usr/sbin/opaswitchadmin -L /etc/sysconfig/opa/switches info
Executing report switch info Test Suite (switchinfo) day mmm dd hh:mm:ss
timezone yyyy ...
Executing TEST SUITE report switch info CASE (switchinfo.
0x00117500ff513121,Node_Name.i2c.extmgd.switchinfo)
  retrieve switch 0x00117500ff513121,Node_Name ...
TEST SUITE report switch info CASE (switchinfo.
0x00117500ff513121,Node_Name.i2c.extmgd.switchinfo)
  retrieve switch 0x00117500ff513121,Node_Name
0x00117500ff513121,hds1swb8171:
  F/W ver:xx.xx.xx.xx H/W ver:XXX H/W pt num:NNNNNN-NNN
  Fan status:Normal/Normal/Normal/Normal/Normal/Normal PS1 Status:N/A PS2
Status:ONLINE
PASSED
TEST SUITE report switch info: 1 Cases; 1 PASSED
TEST SUITE report switch info PASSED
Done report switch info Test Suite day mmm dd hh:mm:ss timezone yyyy
```

If any Intel® Omni-Path Switch 100 Series switches were purposely skipped, this operation should be repeated for those switches. In this case, Intel recommends that you create a separate file with a name other than switches.

## 5.9 Getting Basic Switch Configuration

**(Switch)** The **Get Basic Switch Configuration** selection allows you to view the switch configuration report for all of the ports.

1. From the FastFabric OPA Switch Setup/Admin menu, type **8**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Switch Admin: Get basic Switch configuration
Executing: /usr/sbin/opaswitchadmin -L /etc/sysconfig/opa/switches getconfig
Executing report switch getconfig Test Suite (switchgetportconfig) day mmm dd
hh:mm:ss timezone yyyy ...
Executing TEST SUITE report switch getconfig CASE (switchgetportconfig.
0x00117500ff513121,hds1swb8171.i2c
.extmgd.switchgetportconfig) retrieve switch 0x00117500ff513121,Node_Name ...
TEST SUITE report switch getconfig CASE (switchgetportconfig.
0x00117500ff513121,Node_Name.i2c
.extmgd.switchgetportconfig) retrieve switch 0x00117500ff513121,Node_Name
Link Width : 1,2,3,4
Link Speed : 25Gb
FM Enabled : Yes
Link CRC Mode : 14-bit,16-bit,48-bit
vCU : 0
External Loopback Allowed : Yes
Node Description : Node_Name

PASSED
TEST SUITE report switch getconfig: 1 Cases; 1 PASSED
TEST SUITE report switch getconfig PASSED
Done report switch getconfig Test Suite day mmm dd hh:mm:ss timezone yyyy
```



The results show the number of cases, how many of the cases passed, and how many of the cases failed. It also gives an overall summary of configuration and passed or failed.

## 5.10 Reporting Switch VPD Information

**(Switch)** The **Report Switch VPD Information** selection allows you to view the vital product data (VPD) for all of the nodes listed in `/etc/sysconfig/opa/switches`.

1. From the FastFabric OPA Switch Setup/Admin menu, type **9**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Switch Admin: Report Switch VPD information
Executing: /usr/sbin/opaswitchadmin -L /etc/sysconfig/opa/switches hwpd
Executing report switch hwpd Test Suite (switchhwpd) day mmm dd hh:mm:ss
timezone yyyy ...
Executing TEST SUITE report switch hwpd CASE (switchhwpd.
0x00117500ff513121,Node_Name.i2c
.extmgd.switchhwpd) retrieve switch 0x00117500ff513121,Node_Name ...
TEST SUITE report switch hwpd CASE (switchhwpd.
0x00117500ff513121,Node_Name.i2c
.extmgd.switchhwpd) retrieve switch 0x00117500ff513121,Node_Name

0x00117500ff513121,hds1swb8171: H/W VPD serial number: USFU13150000D
0x00117500ff513121,hds1swb8171: H/W VPD part number : NNNNNN-NNN
0x00117500ff513121,hds1swb8171: H/W VPD model : 100SWE48QF2
0x00117500ff513121,hds1swb8171: H/W VPD h/w version : 004
0x00117500ff513121,hds1swb8171: H/W VPD manufacturer : Intel Corporation
0x00117500ff513121,hds1swb8171: H/W VPD prod desc : 100 OP Edge 48p Q7
forward 2PSU
0x00117500ff513121,hds1swb8171: H/W VPD mfg id : 001175
0x00117500ff513121,hds1swb8171: H/W VPD mfg date : m-dd-yyyy
0x00117500ff513121,hds1swb8171: H/W VPD mfg time : hh:mm
PASSED
TEST SUITE report switch hwpd: 1 Cases; 1 PASSED
TEST SUITE report switch hwpd PASSED
Done report switch hwpd Test Suite day mmm dd hh:mm:ss timezone yyyy
```

## 5.11 Viewing opaswitchadmin Result Files

**(Switch)** The **View opaswitchadmin Result Files** selection allows you to open and view the `punchlist.csv`, `test.res`, and `test.log` files.

1. From the FastFabric OPA Switch Setup/Admin menu, type **a**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.





2. Type **P** to begin the operation.

```
Performing Switch Admin: View opaswitchadmin Result Files
Using vi (to select a different editor, export EDITOR).
About to: vi /root/punchlist.csv /root/test.res /root/test.log
Hit any key to continue (or ESC to abort)...
```

3. Press any key to view opaswitchadmin results files.
4. After reviewing and closing the log, you are prompted to remove the following files.

```
3 files to edit
Would you like to remove test.res test.log test_tmp* and save_tmp
in /root ? [n]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.
6. If you chose **y** in the step above, press any key or **ESC** to end the operation



## 6.0 Managing the Host Configuration

The FastFabric OPA Host Setup menu allows you to set up and install the Fabric software on all the hosts.

To access up the FastFabric OPA Host Setup Menu, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opafastfabric**.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

3. Type **3**.


The FastFabric OPA Host Setup menu is displayed.

```
FastFabric OPA Host Setup Menu
Host File: /etc/sysconfig/opa/hosts
Setup:
0) Edit Config and Select/Edit Host File      [ Skip ]
1) Verify Hosts Pingable                     [ Skip ]
2) Set Up Password-Less SSH/SCP               [ Skip ]
3) Copy /etc/hosts to All Hosts               [ Skip ]
4) Show uname -a for All Hosts                [ Skip ]
5) Install/Upgrade OPA Software               [ Skip ]
6) Configure IPoIB IP Address                 [ Skip ]
7) Build Test Apps and Copy to Hosts          [ Skip ]
8) Reboot Hosts                              [ Skip ]
Admin:
9) Refresh SSH Known Hosts                   [ Skip ]
a) Rebuild MPI Library and Tools              [ Skip ]
b) Run a Command on All Hosts                 [ Skip ]
c) Copy a File to All Hosts                   [ Skip ]
Review:
d) View opahostadmin Result Files             [ Skip ]

P) Perform the Selected Actions               N) Select None
X) Return to Previous Menu (or ESC)
```

4. Select one or more items by typing the alphanumeric character associated with the item to toggle the selection from **Skip** to **Perform**.
5. Type **P** to perform the operations.

*Note:* Each menu item will present you with prompts to complete the operation.

**Table 6. FastFabric OPA Host Setup Menu Descriptions**

Menu Item	Description
0) Edit Config and Select/Edit Host File	<p>Allows you to edit the following configuration files:</p> <ul style="list-style-type: none"> <li>• <code>/etc/sysconfig/opa/hosts</code> The <code>hosts</code> file lists the names of the hosts in a cluster except the FastFabric toolset node.</li> <li>• <code>/etc/sysconfig/opa/opafastfabric.conf</code> The <code>opafastfabric.conf</code> file lists the default settings for most of the FastFabric command line options</li> </ul> <p><b>NOTE:</b> The <code>hosts</code> file selected and created using this menu should not list the FastFabric host itself.</p> <p>Associated CLI Command: <code>opahostadmin</code></p>
1) Verify Hosts Pingable	<p>Allows you to ping all the hosts listed through the Management Network.</p> <p>Associated CLI Command: <code>opapingall</code></p>
2) Set Up Password-Less SSH/SCP	<p><b>(Linux)</b> Allows you to set up secure password-less SSH such that the Fabric Management Node can securely log into all the other hosts as root through the management network without requiring a password.</p> <p>Associated CLI Command: <code>opasetupssh -p -S -i ""</code></p>
3) Copy <code>/etc/hosts</code> to All Hosts	<p><b>(Linux)</b> Allow you to copy the <code>/etc/hosts</code> file on this host to all the other selected hosts.</p> <p><b>NOTE:</b> This is not necessary when using a DNS server to resolve host names for the cluster.</p> <p>Associated CLI Command: <code>scpall /etc/hosts /etc/hosts</code></p>
4) Show <code>uname -a</code> for All Hosts	<p><b>(Linux)</b> Allows you to view the OS version on all the hosts.</p> <p>In typical clusters, all hosts are running the same OS and kernel version.</p> <p>Associated CLI Command: <code>opacmdall "uname -a"</code></p>
5) Install/Upgrade OPA Software	<p><b>(Host)</b> Allows you to install the Intel® OPA software on all the hosts.</p> <p>Associated CLI Commands: <code>opahostadmin load</code>, <code>opahostadmin update</code></p>
6) Configure IPoIB IP Address	<p><b>(Host)</b> Allow you to create the <code>ifcfg-ib0</code> files on each host.</p> <p>The file will be created with a statically-assigned IPv4 address.</p> <p>Associated CLI Command: <code>opahostadmin configipoib</code></p>
7) Build Test Apps and Copy to Hosts	<p><b>(Host)</b> Allows you to build the MPI sample benchmarks on the Fabric Management Node and copy the resulting object files to all the hosts.</p>
8) Reboot Hosts	<p><b>(Linux)</b> Allows you to reboot all the selected hosts and to ensure they reboot fully (as verified using ping over the management network). When the hosts come back up, they will be running the software installed.</p> <p>Associated CLI Command: <code>opahostadmin reboot</code></p>
9) Refresh SSH Known Hosts	<p><b>(Linux)</b> Allows you to refresh the ssh known hosts list on this server for the Management Network.</p> <p>This option needs to be executed after the configuration of IPoIB interfaces on any or all hosts.</p> <p>In addition, this option may be used to update security for this host to complete installation of the hosts or if hosts are installed, replaced, reinstalled, renamed, or repaired.</p> <p>Associated CLI Command: <code>opasetupssh -p -U -f /etc/sysconfig/opa/hosts</code></p>
a) Rebuild MPI Library and Tools	<p><b>(Host)</b> Allows you to rebuild the MPI Library and related tools (such as <code>mpirun</code>).</p>
b) Run a Command on All Hosts	<p><b>(Linux)</b> Allows you to run a command on all hosts.</p> <p><b>NOTE:</b> A Linux shell command (or sequence of commands separated by semicolons) may be specified to be executed against all selected hosts.</p>

*continued...*



Menu Item	Description
	Associated CLI Command: opacmdall
c) Copy a File to All Hosts	<b>(Linux)</b> Allow you to copy a file to all hosts. <b>NOTE:</b> A file on the local host may be specified to be copied to all selected hosts. Associated CLI Command: opascpall
d) View opahostadmin Result Files	Allows you to view the test.log and test.res files that reflect the results from opahostadmin runs (such as for installing software or rebooting all hosts per menu items above). Associated CLI Command: opahostadmin

## 6.1 Editing the Configuration Files for Host Setup

The **Edit Config and Select/Edit Host File** selection allows you to edit the hosts and FastFabric configuration files.

1. From the FastFabric OPA Host Setup menu, type 0.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Edit Config and Select/Edit Host File
Using vi (to select a different editor, export EDITOR).
You will now have a chance to edit/review the FastFabric Config File:
/etc/sysconfig/opa/opafastfabric.conf
The values in this file will control the default operation of the
FastFabric Tools. With the exception of the host file to use,
the values you specify for defaults will be used for all FastFabric
Operations performed via this menu system
Beware existing environment variables will override the values in this file.

About to: vi /etc/sysconfig/opa/opafastfabric.conf
Hit any key to continue (or ESC to abort)...
```

3. Press any key to **open the opafastfabric.conf file** or ESC to abort the operation.

*Note:* To get to subsequent configuration files, you must access each file.

The configuration file opens.

4. Review the settings.

Refer to [FastFabric Configuration File](#) on page 35 for more information.

5. After saving and closing the opafastfabric.conf file in the editor, you will be given the opportunity to edit the hosts file.

```
The FastFabric operations which follow will require a file
listing the hosts to operate on
You should select a file which OMITS this host
Select Host File to Use/Edit [/etc/sysconfig/opa/hosts]:
```

6. Press any key to **open the hosts file** or ESC to abort the operation.

The configuration file opens.



Refer to [Hosts List Configuration Files](#) on page 40 for more information.

For further details about the Host Lists file format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

7. Create the file with a list of the hosts names (the TCP/IP management network names), except the Management Node from which FastFabric is presently being run.

Enter one host's name per line. For example:

```
host1
host2
```

**Note:** Do not list the Management Node itself (the node where FastFabric is currently running).

If additional Management Nodes are to be used, they may be listed at this time, and FastFabric can aid in their initial installation and verification.

8. After saving and closing the `hosts` file in the editor, you will be given the opportunity to review and change the configuration files again.

```
Selected Host File: /etc/sysconfig/opa/hosts
Do you want to edit/review/change the files? [y]:
```

9. Press **Enter** to review and edit the files or type `n` and press **Enter** to end the operation.

## 6.2 Verifying Hosts are Pingable

**(All)** The **Verify Hosts Pingable** selection pings each selected host over the management network.

1. From the FastFabric OPA Host Setup menu, type **1**.

The menu item changes from `[Skip]` to `[Perform]`.

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Verify Hosts Pingable
Would you like to verify hosts are ssh-able? [n]:
```

3. Press **Enter** to select the default (`n`) or enter **y** and press **Enter**.

The status is displayed.

```
Executing: /usr/sbin/opafindgood -A -Q -R -f /etc/sysconfig/opa/hosts
1 hosts will be checked
1 hosts are pingable (alive)
1 hosts are alive (good)
0 hosts are bad (bad)
Bad hosts have been added to /root/punchlist.csv
Hit any key to continue (or ESC to abort)...
```



4. Press any key to continue or **ESC** to abort the operation.

```
Would you like to now use /etc/sysconfig/opa/good as Host File? [y]:
```

5. Press **Enter** to select the default (y) or enter **n** and press **Enter** to end the operation.
6. If some hosts were not found, use the following list to assist in troubleshooting:
  - Host powered on and booted?
  - Host connected to management network?
  - Host management network IP address and network settings consistent with DNS or `/etc/hosts`?
  - Management node connected to the management network?
  - Management node IP address and network settings correct?
  - Management network itself up (including switches, routers, and others)?
  - Correct set of hosts listed in the hosts file? You may need to repeat the previous step to review and edit the file.

## 6.3 Setting Up Password-Less SSH/SCP

**(Linux)** The **Setup Password-less ssh/scp** selection allows you to set up secure password-less SSH (root password) such that the Management Node can securely log in to all the other hosts as root through the management network without requiring a password.

*Note:* Password-less SSH is required by Intel® Omni-Path Fabric Suite FastFabric, MPI test applications, and most versions of MPI (including OpenMPI and MVAPICH2).

1. From the FastFabric OPA Host Setup menu, type **2**.

The menu item changes from `[Skip]` to `[Perform]`.

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Set Up Password-Less SSH/SCP
Executing: /usr/sbin/opasetupssh -S -p -i '' -f /etc/sysconfig/opa/hosts
Password for root on all hosts:
```

3. Type the password for root on all hosts and press **Enter**.

## 6.4 Copying /etc/hosts to All Hosts

**(Linux)** The **Copy /etc/hosts to all hosts** selection allows you to copy the `/etc/hosts` file on this host to all the other selected hosts.

*Note:* If DNS is being used, this task is not required.



**Note:** Typically, `/etc/resolv.conf` is set up as part of OS installation for each host. However, if `/etc/resolv.conf` was not set up on all the hosts during OS installation, the **Copy a File to All Hosts** operation could be used at this time to copy `/etc/resolv.conf` from the Management Node to all the other nodes.

1. From the FastFabric OPA Host Setup menu, type **3**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Copy /etc/hosts to All Hosts
Executing: /usr/sbin/opascall -p -f /etc/sysconfig/opa/hosts /etc/hosts /etc/
hosts
scp -q /etc/hosts root@[phgppriv11]:/etc/hosts
Hit any key to continue (or ESC to abort)...
```

3. Press any key to continue or **ESC** and press **y** to cancel the operation.

## 6.5 Showing `uname -a` for All Hosts

**(Linux)** The **Show `uname -a` for All Hosts** selection allows you to show the OS version on all the hosts.

1. From the FastFabric OPA Host Setup menu, type **4**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Show uname -a for All Hosts
Executing: /usr/sbin/opacmdall -T 60 -f /etc/sysconfig/opa/hosts 'uname -a'
[root@phgppriv11]# uname -a
Linux phgppriv11.ph.intel.com 3.10.0-123.el7.x86_64 #1 SMP Mon May 5 11:16:57
EDT 2014 x86_64 x86_64 x86_64 GNU/Linux
Hit any key to continue (or ESC to abort)...
```

3. Press any key or **ESC** to end the operation.
4. Review the results to verify all the hosts have the expected OS version.
  - In typical clusters, all hosts are running the same OS and kernel version.
  - If any hosts are identified with an incorrect OS version, the OS on those hosts should be corrected at this time.

After the OS versions have been corrected, perform [Copying a File to All Hosts](#) on page 86.



## 6.6 Installing/Upgrading OPA Software

**(Host)** The **Install/Upgrade OPA Software** selection allows you to install or upgrade the Intel® Omni-Path Fabric Host Software on all the hosts. By default, it looks in the current directory for the `IntelOPA-[Basic|IFS].DISTRO.VERSION.tgz` file. If the file is not found in the current directory, the installer application prompts for a directory name where this file can be found.

**Note:** Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for performing first-time installations and upgrades.

1. From the FastFabric OPA Host Setup menu, type **5**.

The menu item changes from `[Skip]` to `[Perform]`.

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Enter Directory to get IntelOPA-Basic.RHEL7-x86_64.10.3.0.0.52.tgz from (or none):	Allows you to enter the directory to the software. If none, you will be prompted whether you want to proceed: <ul style="list-style-type: none"><li>• Select <b>y</b> to continue.</li><li>• Select <b>n</b> to abort.</li></ul>
Do you want to use ./IntelOPA-[Basic IFS].DISTRO.VERSION.tgz? [y]:	Allows you to select the tgz that is required for the installation or upgrade.
Would you like to do a fresh [i]ninstall, an [u]pgrade or [s]kip this step? [u]:	<ul style="list-style-type: none"><li>• Select <b>i</b> to install software.</li><li>• Select <b>u</b> to upgrade software.</li><li>• Select <b>s</b> to skip this step.</li></ul>
Are you sure you want to proceed? [n]:	

After executing the prompts, the following is displayed:

```
/usr/sbin/opahostadmin -f /etc/sysconfig/opa/hosts -d . load
Executing load Test Suite (load) Day Mth DD HH:MM:SS timezone yyyy ...
.
.
Hit any key to continue (or ESC to abort)...
```

**Note:** If any hosts fail to be installed, you will see results as shown in the following example:

```
TEST SUITE load: 1 Cases; 0 PASSED; 1 FAILED
TEST SUITE load FAILED
```

Use the [Viewing opahostadmin Result Files](#) on page 104 option to review the result files from the update. For more details, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

4. Press any key or **ESC** to end the operation.





## 6.7 Configuring IPoIB IP Address

**(Host)** The **Configure IPoIB IP Address** selection allows you to create the `ifcfg-ib0` files on each host. The file is created with a statically-assigned IPv4 address. The IPoIB IP address for each host is determined by the resolver (Linux\* host command). If not found through the resolver, the `/etc/hosts` file for the given host is checked.

1. From the FastFabric OPA Host Setup menu, type **6**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Configure IPoIB IP Address
Executing: /usr/sbin/opahostadmin -f /etc/sysconfig/opa/hosts configipoib
Executing configure IPoIB Test Suite (configipoib) Fri Oct 07 11:31:49 EDT
2016 ...
Executing TEST SUITE configure IPoIB CASE (configipoib.phgppriv11.config)
config ipoib on phgppriv11 ...
...
Done configure IPoIB Test Suite Fri Oct 07 11:31:51 EDT 2016

Hit any key to continue (or ESC to abort)...
```

3. Press any key or **ESC** to end the operation.

## 6.8 Building Test Applications and Copying to Hosts

**(Host)** The **Build Test Apps and Copy to Hosts** selection allows you to build the MPI and/or SHMEM sample applications on the Management Node and copy the resulting object files to all the hosts. This is in preparation for execution of MPI and/or SHMEM performance tests and benchmarks.

1. From the FastFabric OPA Host Setup menu, type **7**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Build Test Apps and Copy to Hosts
Do you want to build MPI Test Apps? [y]:
```

3. Press **Enter**.

The MPI Directory Selection TUI is displayed.

```
Host Setup: Build Test Apps and Copy to Hosts
MPI Directory Selection

Please Select MPI Directory:
0) /usr/mpi/gcc/mvapich2-2.1
1) /usr/mpi/gcc/mvapich2-2.1-hfi
2) /usr/mpi/gcc/openmpi-1.10.2
3) /usr/mpi/gcc/openmpi-1.10.2-hfi
4) /usr/mpi/intel/mvapich2-2.1-hfi
5) /usr/mpi/intel/openmpi-1.10.2-hfi
```



```
6) /usr/mpi/pgi/mvapich2-2.1-hfi
7) /usr/mpi/pgi/openmpi-1.10.0-hfi
8) Enter Other Directory

X) Return to Previous Menu (or ESC)
```

4. Select the target menu item or type X to return to the operation.

The next prompt is shown:

```
Do you want to build SHMEM Test Apps? [y]:
```

5. Press **Enter**.

The MPI Directory Selection for SHMEM Job Launch TUI is displayed.

```
Host Setup: Build Test Apps and Copy to Hosts
MPI Directory Selection for SHMEM Job Launch

Please Select MPI Directory:
0) /usr/mpi/gcc/openmpi-1.10.2
1) /usr/mpi/gcc/openmpi-1.10.2-hfi
2) /usr/mpi/intel/openmpi-1.10.2-hfi
3) /usr/mpi/pgi/openmpi-1.10.0-hfi
4) Enter Other Directory
5) Skip MPI Directory Selection for SHMEM Job Launch

X) Return to Previous Menu (or ESC)
```

6. Select the target menu item or type X to return to the operation.
7. Follow the prompts to complete the operation.

## 6.9 Rebooting Hosts

**(Linux)** The **Reboot Hosts** selection allows you to reboot all the selected hosts and ensure they fully reboot, as verified through ping over the management network.

1. From the FastFabric OPA Host Setup menu, type 8.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

Reboot begins immediately.

```
Performing Host Setup: Reboot Hosts
Executing: /usr/sbin/opahostadmin -f /etc/sysconfig/opa/hosts reboot
Executing reboot Test Suite (reboot) Fri Oct 07 11:58:48 EDT 2016 ...
Executing TEST SUITE reboot CASE (reboot.phgppriv11.reboot) phgppriv11
reboot ...
```

## 6.10 Refreshing SSH Known Hosts

**(Linux)** The **Refresh SSH Known Hosts** selection allows you to refresh the SSH known hosts list on this server for the Management Network. This may be used to update security for this host if hosts are replaced, reinstalled, renamed, or repaired.

1. From the FastFabric OPA Host Setup menu, type 9.



The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Refresh SSH Known Hosts
Executing: /usr/sbin/opasetupssh -p -U -f /etc/sysconfig/opa/hosts
Verifying localhost ssh...
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.
localhost: Connected
Warning: Permanently added 'phgppriv10,10.228.209.74' (ECDSA) to the list of
known hosts.
phgppriv10: Connected
ssh: Could not resolve hostname phgppriv10-opa: Name or service not known
Connecting to phgppriv11...
Warning: Permanently added 'phgppriv11,10.228.209.75' (ECDSA) to the list of
known hosts.
phgppriv11: Connected
ssh: Could not resolve hostname phgppriv11-opa: Name or service not known
setup_self_ssh 100% 5599 5.5KB/s 00:00
phgppriv11: Verifying localhost ssh...
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.
localhost: Connected
Warning: Permanently added 'phgppriv11,10.228.209.75' (ECDSA) to the list of
known hosts.
phgppriv11: Connected
ssh: Could not resolve hostname phgppriv11-opa: Name or service not known
phgppriv11: Configured localhost ssh
Successfully processed: 3
Hit any key to continue (or ESC to abort)...
```

3. Press any key or **ESC** to end the operation.

## 6.11 Rebuilding MPI Library and Tools

**(Host)** The **Rebuild MPI Library and Tools** allows you to rebuild the MPI Library and related tools (such as `mpirun`), and install the resulting rpms on all the hosts.

This operation is performed using the `do_build` tool supplied with the MPI Source. When rebuilding MPI, `do_build` prompts you for selection of which MPI to rebuild, and provides choices as to which available compiler to use. Refer to *Intel® Omni-Path Host Fabric Interface Installation Guide* and *Intel® Omni-Path Fabric Host Software User Guide* for more information.

1. From the FastFabric OPA Host Setup menu, type **a**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Rebuild MPI Library and Tools
Executing: cd //usr/lib/opa/src/MPI; ./do_build

OFED MVAPICH MPI Library/Tools rebuild
1) openmpi
2) mvapich2
Select MPI to Build:
```



3. Enter the menu item to rebuild and press **Enter**.
  - a. For openmpi, go to step 4.
  - b. For mvapich2, go to step 5.
4. Rebuild openmpi.

```
OFA OpenMPI MPI Library/Tools rebuild
1) gcc
Select Compiler:
```

- a. Enter the menu item and press **Enter**.

```
Build for Omnipath HFI PSM [y]:
```

- b. Press **Enter** to continue or **n** and **Enter** to abort.

```
Executing: cd /usr/lib/opa/src/MPI && /usr/sbin/opascpall -p -f /etc/
sysconfig/opa/hosts /var/tmp
Usage: opascpall [-p] [-r] [-f hostfile] source_file ... dest_file
       opascpall -t [-p] [-f hostfile] [source_dir [dest_dir]]
       or
       opascpall --help
--help - produce full help text
-p - perform copy in parallel on all hosts
-r - recursive copy of directories
-t - optimized recursive copy of directories using tar
    if dest_dir omitted, defaults to current directory name
    if source_dir and dest_dir omitted, both default to current
directory
-f hostfile - file with hosts in cluster, default is /etc/
sysconfig/opa/hosts
source_file - list of source files to copy
source_dir - source directory to copy, if omitted . is used
dest_file - destination for copy.
             If more than 1 source file, this must be a directory
dest_dir - destination for copy. If omitted current directory name is
used
example:
opascpall MPI-PMB /root/MPI-PMB
opascpall -t -p /usr/lib/opa/src/mpi_apps /usr/lib/opa/src/mpi_apps
opascpall a b c /root/tools/
user@ syntax cannot be used in filenames specified
To copy from hosts in the cluster to this host, use opauploadall
Hit any key to continue (or ESC to abort)...
```

- c. Press any key to continue.

```
Executing: /usr/sbin/opacmdall -p -f /etc/sysconfig/opa/hosts 'cd /var/
tmp; rpm -U --force ; rm -f '
[root@phgppriv11]# cd /var/tmp; rpm -U --force ; rm -f
...
Hit any key to continue (or ESC to abort)...
```

- d. Press any key or ESC to end operation.

5. Rebuild mvapich2.

```
OFA MVAPICH2 MPI Library/Tools rebuild
1) gcc
Select Compiler:
```



- a. Enter the menu item and press **Enter**.

```
1) ofa
2) opa-psm
Select MVAPICH2 Implementation (opa-psm recommended):
```

- b. Enter the menu item and press **Enter**.

```
Executing: cd /usr/lib/opa/src/MPI && /usr/sbin/opascpall -p -f /etc/
sysconfig/opa/hosts /var/tmp
Usage: opascpall [-p] [-r] [-f hostfile] source_file ... dest_file
       opascpall -t [-p] [-f hostfile] [source_dir [dest_dir]]
       or
       opascpall --help
--help - produce full help text
-p - perform copy in parallel on all hosts
-r - recursive copy of directories
-t - optimized recursive copy of directories using tar
    if dest_dir omitted, defaults to current directory name
    if source_dir and dest_dir omitted, both default to current
    directory
-f hostfile - file with hosts in cluster, default is /etc/
sysconfig/opa/hosts
source_file - list of source files to copy
source_dir - source directory to copy, if omitted . is used
dest_file - destination for copy.
             If more than 1 source file, this must be a directory
dest_dir - destination for copy. If omitted current directory name is
used
example:
opascpall MPI-PMB /root/MPI-PMB
opascpall -t -p /usr/lib/opa/src/mpi_apps /usr/lib/opa/src/mpi_apps
opascpall a b c /root/tools/
user@ syntax cannot be used in filenames specified
To copy from hosts in the cluster to this host, use opauploadall
Hit any key to continue (or ESC to abort)...
```

- c. Press any key to continue.

```
Executing: /usr/sbin/opacmdall -p -f /etc/sysconfig/opa/hosts 'cd /var/
tmp; rpm -U --force ; rm -f '
[root@phgppriv11]# cd /var/tmp; rpm -U --force ; rm -f
...
Hit any key to continue (or ESC to abort)...
```

- d. Press any key or **ESC** to end operation.

## 6.12 Running a Command on All Hosts

**(Linux)** The **Run a Command on All Hosts** selection allows you to perform other operations on all hosts. Each time this is executed, a Linux\* shell command may be specified to be executed against all selected hosts. You can also specify a sequence of commands separated by semicolons.

1. From the FastFabric OPA Host Setup menu, type **b**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.



2. Type **P** to begin the operation.

```
Performing Host Setup: Run a Command on All Hosts
Enter Command to run on all hosts (or none):
```

3. Enter a Linux command and press **Enter**.

```
Timelimit in minutes (0=unlimited): [1]:
```

4. Specify a time limit and press **Enter**.

```
Run in parallel on all hosts? [y]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.

```
About to run: /usr/sbin/opacmdall -T 60 -f /etc/sysconfig/opa/hosts 'xxxx'
Are you sure you want to proceed? [n]:
```

6. Type **y** and press **Enter** to proceed with the operation.  
The operation is completed.

## 6.13 Copying a File to All Hosts

**(Linux)** The **Copy a File to All Hosts** selection allows you to run the `opascpall` command. A file on the local host may be specified to be copied to all selected hosts.

1. From the FastFabric OPA Host Setup menu, type **c**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: Copy a File to All Hosts
Enter File to copy to all hosts (or none):
```

3. Enter the name of the file to copy and press **Enter**.

```
Are you sure you want to proceed? [n]:
```

4. Type **y** and press **Enter** to continue.

```
Executing: /usr/sbin/opascpall -p -f /etc/sysconfig/opa/hosts /root/xxx /root/
xxx
scp -q /root/xxx root@[phgppriv11]:/root/xxx
...
Hit any key to continue (or ESC to abort)...
```

5. Press any key or **ESC** to end the operation.



## 6.14 Viewing opahostadmin Result Files

**(All)** The **View opahostadmin Result File** selection allows you to display the `test.log` and `test.res` files that contain the results from prior `opahostadmin` runs, such as installing Fabric software or rebooting all hosts. You are also given the option to remove these files after viewing them.

If prior files are not removed, subsequent runs of `opachassisadmin`, `opahostadmin`, or `opaswitchadmin` from within the current directory continue to append to these files.

1. From the FastFabric OPA Host Setup menu, type **d**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Setup: View opahostadmin Result Files
Using vi (to select a different editor, export EDITOR).
About to: vi /root/test.res /root/test.log
Hit any key to continue (or ESC to abort)...
```

3. Press any key to view the `opahostadmin` results files.
4. After reviewing and closing the log, you are prompted to remove the following files.

```
Would you like to remove test.res test.log test_tmp* and save_tmp
in /root ? [n]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.
6. If you chose **y** in the step above, press any key or **ESC** to end the operation



## 7.0 Verifying the Host

The FastFabric OPA Host Verification/Admin Menu allows you to verify hosts and the fabric, as well as manage all the hosts.

To access up the FastFabric OPA Host Setup Menu, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opafastfabric**.

The Intel FastFabric OPA Tools menu is displayed.

```
Intel FastFabric OPA Tools
Version: X.X.X.X.X

1) Chassis Setup/Admin
2) Externally Managed Switch Setup/Admin
3) Host Setup
4) Host Verification/Admin
5) Fabric Monitoring

X) Exit
```

3. Type **4**.

The FastFabric OPA Host Verification/Admin Menu is displayed.

```
FastFabric OPA Host Verification/Admin Menu
Host File: /etc/sysconfig/opa/allhosts
Validation:
0) Edit Config and Select/Edit Host File [ Skip ]
1) Summary of Fabric Components [ Skip ]
2) Verify Hosts Are Pingable, SSHable, and Active [ Skip ]
3) Perform Single Host Verification [ Skip ]
4) Verify OPA Fabric Status and Topology [ Skip ]
5) Verify Hosts See Each Other [ Skip ]
6) Verify Hosts Ping via IPoIB [ Skip ]
7) Refresh SSH Known Hosts [ Skip ]
8) Check MPI Performance [ Skip ]
9) Check Overall Fabric Health [ Skip ]
a) Start or Stop Bit Error Rate Cable Test [ Skip ]
Admin:
b) Generate All Hosts Problem Report Info [ Skip ]
c) Run a Command on All Hosts [ Skip ]
Review:
d) View opahostadmin Result Files [ Skip ]

P) Perform the Selected Actions N) Select None
X) Return to Previous Menu (or ESC)
```

4. Select one or more items by typing the alphanumeric character associated with the item to toggle the selection from Skip to Perform.
5. Type **P** to perform the operations.

*Note:* Each menu item will present you with prompts to complete the operation.



**Table 7. FastFabric OPA Host Verification/Admin Menu Descriptions**

Menu Item	Description
0) Edit Config and Select/Edit Host File	<p>Allows you to edit the following configuration files:</p> <ul style="list-style-type: none"> <li>• <code>/etc/sysconfig/opa/allhosts</code> The <code>allhosts</code> file lists of all hosts including the FastFabric toolset node.</li> <li>• <code>/etc/sysconfig/opa/ports</code> The <code>ports</code> file lists the local HFI ports (for example, subnets) to be used to access the fabric for analysis.</li> <li>• <code>/etc/sysconfig/opa/opafastfabric.conf</code> The <code>opafastfabric.conf</code> file lists the default settings for most of the FastFabric command line options.</li> </ul> <p>Associated CLI Command: <code>opahostadmin</code></p>
1) Summary of Fabric Components	<p>Allows you to view a brief summary of the components in the fabric including the number of components, how many switch chips, HFIs, and links. It also indicates whether any degraded or omitted (quarantined or out of policy) links were found.</p> <p>Associated CLI Command: <code>opafabricinfo</code></p>
2) Verify Hosts Are Pingable, SSHable, and Active	<p>Allows you to ping all the hosts listed through the Management Network.</p> <p>Associated CLI Command: <code>opapingall</code></p>
3) Perform Single Host Verification	<p>Allows you to perform verification on all nodes in the selected host file including configuration, performance, and stability using a variety of tools and checks including single node HPL .</p> <p>For additional information on the verification that is performed, refer to the <code>/usr/lib/opa/samples/hostverify.sh</code> file.</p> <p>Associated CLI Command: <code>opapingall</code></p>
4) Verify OPA Fabric Status and Topology	<p><b>(Host or All)</b> Allows you to review the fabric state and error counts of all ports.</p> <p>Associated CLI Commands: <code>opashowallports</code>, <code>opareport</code></p>
5) Verify Hosts See Each Other	<p><b>(Host)</b> Allows you to verify that each host can see all the others through queries to the Subnet Administrator.</p> <p>Associated CLI Command: <code>opahostadmin sacache</code></p>
6) Verify Hosts Ping via IPoIB	<p><b>(Host)</b> Allows you to verify that IPoIB is properly configured and running on all the hosts. This is accomplished through the Fabric Management node pinging each host using IPoIB.</p> <p>Associated CLI Command: <code>opahostadmin ipoibping</code></p>
7) Refresh SSH Known Hosts	<p><b>(Linux)</b> Allows you to refresh the ssh known hosts list on this server for the IPoIB and Management Networks. This option may be used to update security for this host to complete installation of the hosts or if hosts are replaced, reinstalled, renamed, or repaired.</p> <p>Associated CLI Command: <code>opasetupssh -p -U</code></p>
8) Check MPI Performance	<p><b>(Host)</b> Allows you to perform a quick check of PCI and MPI performance using end-to-end latency and bandwidth tests.</p>
9) Check Overall Fabric Health	<p><b>(Host)</b> Allows you to check the overall fabric health.</p> <p>Associated CLI Command: <code>opaallanalysis</code></p>
a) Start or Stop Bit Error Rate Cable Test	<p><b>(Host)</b> Allows you to start or stop the Cable Bit Error Rate stress tests for HFI-to-switch links and/or ISLs.</p> <p>Associated CLI Command: <code>opacabletest</code></p>
b) Generate All Hosts Problem Report Info	<p><b>(Host)</b> Allows you to collect configuration and status information from all hosts and generates a single <code>*.tgz</code> file, which can be sent to a support representative.</p>

*continued...*



Menu Item	Description
	Associated CLI Command: opacaptureall
c) Run a Command on All Hosts	<b>(Linux)</b> Allows you to execute a command on all hosts. Associated CLI Command: opacmdall
d) View opahostadmin Result Files	Allows you to view the test.log and test.res files that reflect the results from opahostadmin runs (such as those for installing software or rebooting all hosts per menu items above). Associated CLI Command: opahostadmin

## 7.1 Editing the Configuration Files for Host Verification

The **Edit Config and Select/Edit Host File** section allows you to select and edit the hosts, ports, and FastFabric configuration files.

1. From the FastFabric OPA Host Verification/Admin menu, type 0.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Admin: Edit Config and Select/Edit Host File
Using vi (to select a different editor, export EDITOR).
You will now have a chance to edit/review the FastFabric Config File:
/etc/sysconfig/opa/opafastfabric.conf
The values in this file will control the default operation of the
FastFabric Tools. With the exception of the host file to use,
the values you specify for defaults will be used for all FastFabric
Operations performed via this menu system
Beware existing environment variables will override the values in this file.

About to: vi /etc/sysconfig/opa/opafastfabric.conf
Hit any key to continue (or ESC to abort)...
```

3. Press any key to **open the opafastfabric.conf file** or ESC to abort the operation.

*Note:* To get to subsequent configuration files, you must access each file.

The configuration file opens.

4. Review the settings.

Especially review the following:

- FF\_TOPOLOGY\_FILE
- FF\_IPOIB\_SUFFIX
- FF\_DEVIATION\_ARGS
- ff\_host\_basename\_to\_ipoib
- ff\_host\_basename

Refer to [FastFabric Configuration File](#) on page 35 for more information.



**Note:** Intel recommends that a FastFabric topology file is created as `/etc/sysconfig/opa/topology.0:0.xml` to describe the intended topology of the fabric. The file can also augment assorted fabric reports with customer-specific information, such as cable labels and additional details about nodes, SMs, links, ports, and cables. Refer to [Fabric Topology Input File](#) on page 42 and the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more information about topology verification files.

5. After saving and closing the `opafastfabric.conf` file in the editor, you will be given the opportunity to edit the `ports` file.

```
You will now have a chance to edit/review the FastFabric PORTS_FILE:
/etc/sysconfig/opa/ports
Some of the FastFabric operations which follow will use this file to
specify the local HFI ports to use to access the fabric(s) to operate on
Beware existing environment variables will override the values in this file.

About to: vi /etc/sysconfig/opa/ports
Hit any key to continue (or ESC to abort)...
```

6. Press any key to **open the ports file** or ESC to abort the operation.

The configuration file opens.

- a. Review the file:

- For typical single-subnet clusters, the default of "0:0" may be used. This uses the first active port on the Management Node to access all externally managed switches. .
- For configuring a cluster with multiple subnets, refer to *Intel® Omni-Path Fabric Software Installation Guide*.

Refer to [Ports List Configuration File](#) on page 35 for more information.

For further details about the Port List File format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

7. After saving and closing the `ports` file in the editor, you will be given the opportunity to edit the `allhosts` file.

```
The FastFabric operations which follow will require a file
listing the hosts to operate on
You should select a file which INCLUDES this host
Select Host File to Use/Edit [/etc/sysconfig/opa/allhosts]:
```

8. Select the host file to edit or leave blank for the default and press **Enter**.

```
About to: vi /etc/sysconfig/opa/allhosts
Hit any key to continue (or ESC to abort)...
```

9. Press any key to **open the allhosts file** or ESC to abort the operation.

The configuration file opens.

Refer to [Hosts List Configuration Files](#) on page 40 for more information.

For further details about the Host Lists file format, refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*.

10. Create the file with the Management Node's host name (the TCP/IP management network name, for example `mgmthost`) and include the hosts file previously created.



Enter one host's name per line. For example:

```
mgmthost
include /etc/sysconfig/opa/hosts
```

11. After saving and closing the `hosts` file in the editor, you will be given the opportunity to review and change the configuration files again.

```
Selected Host File: /etc/sysconfig/opa/allhosts
Do you want to edit/review/change the files? [y]:
```

12. Press **Enter** to review and edit the files again or type `n` and press **Enter** to end the operation.

## 7.2 Viewing a Summary of Fabric Components

The **Summary of Fabric Components** selection allows you to generate a brief summary of the counts of components in the fabric, including how many switch chips, hosts, and links are in the fabric. The summary also indicates whether any degraded or omitted links were found which can indicate a poorly seated or bad cable, incorrect fabric configuration, or security issues.

1. From the FastFabric OPA Host Verification/Admin menu, type **1**.

The menu item changes from `[Skip]` to `[Perform]`.

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The summary is generated.

```
Performing Host Admin: Summary of Fabric Components
Executing: /usr/sbin/opafabricinfo
Fabric 0:0 Information:
SM: phcppriv10 hfil_0 Guid: 0x0011750101575300 State: Master
Number of HFIs: 2
Number of Switches: 0
Number of Links: 1
Number of HFI Links: 1          (Internal: 0   External: 1)
Number of ISLs: 0              (Internal: 0   External: 0)
Number of Degraded Links: 0     (HFI Links: 0   ISLs: 0)
Number of Omitted Links: 0      (HFI Links: 0   ISLs: 0)
-----
-
Hit any key to continue (or ESC to abort)...
```

3. Press any key or **ESC** to end the operation.

## 7.3 Verifying Hosts Pingable, SSHable, and Active

The **Verify Hosts Pingable, SSHable, and Active** selection allows you to verify each host and provides a concise summary of the bad hosts found.

Interactive prompts allow you to select ping, SSH, and port active verification. After completion of this test, you have the option of using the resulting good hosts file for the remainder of the operations within this TUI session.

1. From the FastFabric OPA Host Verification/Admin menu, type **2**.



The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Would you like to verify hosts are ssh-able? [y]:	Allows you to see which hosts are ssh-able.
Would you like to verify host OPA ports are active? [y]:	Allows you to view which ports are active.
Would you like to verify host OPA ports are not quarantined? [y]:	Allows you to view which ports are not quarantined.

After executing the prompts, the results are displayed.

```
Executing: /usr/sbin/opafindgood -f /etc/sysconfig/opa/allhosts
2 hosts will be checked
2 hosts are pingable (alive)
2 hosts are ssh'able (running)
opasaquery: Failed to open port hfi 0:0: Resource temporarily unavailable
0 total hosts have FIs active on one or more fabrics (active)
opareport: No Active ports found in System
Parse error at line 1: no element found
Parse error at line 1: Fatal error parsing file 'stdin'
opaxmlextract: XML Parse error
0 hosts are alive, running, active (good)
2 hosts are bad (bad)
Bad hosts have been added to /root/punchlist.csv
Hit any key to continue (or ESC to abort)...
```

4. Press any key to continue.

```
Would you like to now use /etc/sysconfig/opa/good as Host File? [y]:
```

5. Press **Enter** to use the host file or type **n** and press **Enter** to end the operation

## 7.4 Performing Single Host Verification

The **Perform Single Host Verification** selection allows you to perform a single host test on all hosts.

### Notes:

- Prior to using this selection, you must have a copy of the `hostverify.sh` in the directory pointed to by `FF_HOSTVERIFY_DIR`.
- If the file does not exist in that directory, copy the sample file `/usr/lib/opa/samples/hostverify.sh` to the directory pointed to by `FF_HOSTVERIFY_DIR`. When placed in the editor to review `hostverify.sh`, review the settings near the top and the list of TESTS selected, edit and save as needed.
- Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for information on `opaverifyhosts`.

1. From the FastFabric OPA Host Verification/Admin menu, type **3**.



The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Would you like to edit /root/hostverify.sh and copy to hosts? [y]:	Allows you to edit the hostverify.sh file. The next prompt will appear after you close the file.
Would you like to specify tests to run? [n]:	Allows you to run specific tests.
Enter filename for upload destination file [hostverify.res]:	Allows you to enter a file name for the results file or use the default hostverify.res.
Timelimit in minutes: [1]:	Allows you to set the time limit for the tests.
View Load on hosts prior to verification? [y]:	Allows you to view the load on the hosts before verification begins.

After executing the prompts, the average loads per host are displayed.

```
Executing: /usr/sbin/opacheckload -f /etc/sysconfig/opa/allhosts
loadavg          host
0.66 0.41 0.20 2/880 193597      phgppriv10
0.00 0.01 0.05 1/813 4001       phgppriv11
Hit any key to continue (or ESC to abort)...
```

4. Press any key to start the tests.

```
Executing: /usr/sbin/opaverifyhosts -k -c -u hostverify.res -T 60 -f /etc/
sysconfig/opa/allhosts
Killing hostverify and xhpl on hosts...
[root@phgppriv11]# pkill -9 -f 'host[v]erify.sh'; pkill -9 '[x]hpl'; echo -n
[root@phgppriv10]# pkill -9 -f 'host[v]erify.sh'; pkill -9 '[x]hpl'; echo -n
2 hosts will be verified
SCPing /root/hostverify.sh to /root/hostverify.sh ...
scp -q /root/hostverify.sh root@[phgppriv10]:/root/hostverify.sh
scp -q /root/hostverify.sh root@[phgppriv11]:/root/hostverify.sh
Running /root/hostverify.sh -d /root ...
phgppriv11: FAIL ipoib: ib0 is in 'datagram' mode - should be in 'connected'
mode
phgppriv10: FAIL ipoib: ib0 is in 'datagram' mode - should be in 'connected'
mode
Killing hostverify and xhpl on hosts...
[root@phgppriv11]# pkill -9 -f 'host[v]erify.sh'; pkill -9 '[x]hpl'; echo -n
[root@phgppriv10]# pkill -9 -f 'host[v]erify.sh'; pkill -9 '[x]hpl'; echo -n
Uploading /root/hostverify.res to ./uploads/hostverify.res ...
scp -q root@[phgppriv10]:/root/hostverify.res ./uploads/phgppriv10/
hostverify.res
scp -q root@[phgppriv11]:/root/hostverify.res ./uploads/phgppriv11/
hostverify.res
About to: vi /root/verifyhosts.res
Hit any key to continue (or ESC to abort)...
```

5. Press any key to view the results file.

The results of the test are shown in the editor.

6. Close the results file to end the operation.



## 7.5 Verifying OPA Fabric Status and Topology

The **Verify OPA Fabric Status and Topology** selection allows you to run various checks on the fabric and topology.

1. From the FastFabric OPA Host Verification/Admin menu, type **4**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:

Prompt	Description
Would you like to perform fabric error analysis? [y]:	Allows you to perform fabric error analysis.
Clear error counters after generating report? [n]:	Allows you to clear the error counters after the report is generated.
Would you like to perform fabric link speed error analysis? [y]:	Allows you to perform link speed error analysis.
Check for links configured to run slower than supported? [n]:	Allow you to look for links that are configured to run slower than supported.
Check for links connected with mismatched speed potential? [n]:	Allows you to look for connected links with mismatch speed potential.
Would you like to verify fabric topology? [y]:	Allows you to verify the fabric topology. <b>NOTE:</b> The fabric deployment can be verified against the planned topology. Typically the planned topology will have been converted to an XML topology file using <code>opaxlattopology</code> , <code>opaxlattopology_cust</code> or a customized variation. If this step has been done and a topology file has been placed in the location specified by the <code>FF_TOPOLOGY_FILE</code> in <code>opafastfabric.conf</code> file, then a topology verification can be performed. Refer to the FastFabric CLI reference guide for more info.
Verify all aspects of topology (links, nodes, SMs)? [y]:	Allows you to verify all links, nodes and SMs in the topology.
Include unexpected devices in punchlist? [y]:	Allows you to include unexpected devices in the punchlist.
Enter filename for results [/root/linkanalysis.res]:	Allows you to enter a file name for the result file or accept the default <code>linkanalysis.res</code> .

After executing the prompts, the average loads per host are displayed.

```
Executing: /usr/sbin/opacheckload -f /etc/sysconfig/opa/allhosts
loadavg          host
0.66 0.41 0.20 2/880 193597 phgppriv10
0.00 0.01 0.05 1/813 4001   phgppriv11
Hit any key to continue (or ESC to abort)...
```

4. Press any key or **ESC** to end the operation.

## 7.6 Verifying Hosts See Each Other

**(Host)** The **Verify Hosts See Each Other** selection allows you to confirm that each host can see all the others through queries to the Subnet Administrator. This ensures all nodes are connected to the same fabric and can properly access the Subnet Administrator.

1. From the FastFabric OPA Host Verification/Admin menu, type 5.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Admin: Verify Hosts See Each Other
Executing: /usr/sbin/opahostadmin -f /etc/sysconfig/opa/allhosts sacache
Executing sacache Test Suite (sacache) Fri Oct 07 16:06:08 EDT 2016 ...
Executing TEST SUITE sacache CASE (sacache.phgppriv10.dsc) phgppriv10 can see
phgppriv10 phgppriv11 ...
Executing TEST SUITE sacache CASE (sacache.phgppriv11.dsc) phgppriv11 can see
phgppriv10 phgppriv11 ...
...
Done sacache Test Suite Fri Oct 07 16:08:10 EDT 2016

Hit any key to continue (or ESC to abort)...
```

3. Press any key or ESC to end the operation.

## 7.7 Verifying Hosts Ping via IPoIB

**(Host)** The **Verify Hosts Ping via IPoIB** selection allows you to confirm that IPoIB is properly configured and running on all the hosts. This is accomplished through the Management Node pinging each host through IPoIB.

*Note:* This operation requires that IPoIB be enabled on the Management Node as well as on each host selected for verification. Also, the management host must have IPoIB configured.

1. From the FastFabric OPA Host Verification/Admin menu, type 6.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The status is displayed.

```
Performing Host Admin: Verify Hosts Ping via IPoIB
Executing: /usr/sbin/opahostadmin -f /etc/sysconfig/opa/allhosts ipoibping
Executing ipoib ping Test Suite (ipoibping) Fri Oct 07 16:10:53 EDT 2016 ...
Executing TEST SUITE ipoib ping CASE (ipoibping.localhost_ping1) simple ping
from localhost ...
TEST SUITE ipoib ping CASE (ipoibping.localhost_ping1) simple ping from
localhost ...
...
TEST CASE simple ping from localhost: 2 Items; 2 PASSED; 0 FAILED
TEST SUITE ipoib ping CASE (ipoibping.localhost_ping1) simple ping from
localhost PASSED
```





```
TEST SUITE ipoib ping: 1 Cases; 1 PASSED; 0 FAILED
TEST SUITE ipoib ping PASSED
Done ipoib ping Test Suite Fri Oct 07 16:10:54 EDT 2016

Hit any key to continue (or ESC to abort)...
```

3. Press any key or **ESC** to end the operation.

## 7.8 Refreshing SSH Known Hosts

**(Linux)** The **Refresh SSH Known Hosts** selection allows you to refresh the SSH `known_hosts` file on the Management Node to include the IPoIB hostnames of all the hosts.

**Note:** This operation requires that IPoIB be enabled on the Management Node as well as on each host selected for verification.

1. From the FastFabric OPA Host Verification/Admin menu, type 7.

The menu item changes from `[Skip]` to `[Perform]`.

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The status is displayed.

```
Performing Host Admin: Refresh SSH Known Hosts
Executing: /usr/sbin/opasetupssh -p -U -f /etc/sysconfig/opa/allhosts
Verifying localhost ssh...
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.
localhost: Connected
Warning: Permanently added 'phgppriv10,10.228.209.74' (ECDSA) to the list of
known hosts.
phgppriv10: Connected
ssh: Could not resolve hostname phgppriv10-opa: Name or service not known
Connecting to phgppriv10...
Connecting to phgppriv11...
Warning: Permanently added 'phgppriv10,10.228.209.74' (ECDSA) to the list of
known hosts.
Warning: Permanently added 'phgppriv11,10.228.209.75' (ECDSA) to the list of
known hosts.
phgppriv11: Connected
phgppriv10: Connected
ssh: Could not resolve hostname phgppriv11-opa: Name or service not known
ssh: Could not resolve hostname phgppriv10-opa: Name or service not known
setup_self_ssh          100% 5599      5.5KB/s   00:00
setup_self_ssh          100% 5599      5.5KB/s   00:00
phgppriv11: Verifying localhost ssh...
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.
phgppriv10: Verifying localhost ssh...
Warning: Permanently added 'localhost' (ECDSA) to the list of known hosts.
localhost: Connected
localhost: Connected
Warning: Permanently added 'phgppriv11,10.228.209.75' (ECDSA) to the list of
known hosts.
Warning: Permanently added 'phgppriv10,10.228.209.74' (ECDSA) to the list of
known hosts.
phgppriv11: Connected
phgppriv10: Connected
ssh: Could not resolve hostname phgppriv11-opa: Name or service not known
phgppriv11: Configured localhost ssh
ssh: Could not resolve hostname phgppriv10-opa: Name or service not known
```



```
phgppriv10: Configured localhost ssh
Successfully processed: 2
Hit any key to continue (or ESC to abort)...
```

3. Press any key or ESC to end the operation.

## 7.9 Checking MPI Performance

**(Host)** The **MPI Performance** selection allows you to perform a quick check of PCIe and MPI performance through end-to-end latency and bandwidth tests.

*Note:* This test identifies nodes whose performance is not consistent with others in the fabric. It is not intended as a benchmark of fabric latency and bandwidth. This test purposely uses techniques to reduce test runtime.

1. From the FastFabric OPA Host Verification/Admin menu, type **8**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

The status is displayed.

```
Performing Host Admin: Check MPI Performance
Test Latency and Bandwidth deviation between all hosts? [y]:
```

3. Press **Enter** to select default (y).

```
View Load on hosts prior to test? [y]:
```

4. Press **Enter** to select default (y).

```
Performing Host Admin: Check MPI Performance
Test Latency and Bandwidth deviation between all hosts? [y]:
View Load on hosts prior to test? [y]:
Executing: /usr/sbin/opacheckload -f /etc/sysconfig/opa/allhosts
loadavg          host
0.00 0.01 0.05 1/857 234917      phgppriv10
0.00 0.01 0.05 1/814 4642       phgppriv11
Hit any key to continue (or ESC to abort)...
```

5. Press any key to continue.

```
Executing: /usr/sbin/opahostadmin -f /etc/sysconfig/opa/allhosts
mpiperfdeviation
Executing mpi lat/bw deviation Test Suite (mpiperfdeviation) Fri Oct 07
16:40:18 EDT 2016 ...

TEST SUITE FAILURE: FAILURE during test suite: mpi_apps not built
TEST SUITE mpi lat/bw deviation TESTS ABORTED
TEST SUITE mpi lat/bw deviation: 0 Cases; 0 PASSED
TEST SUITE mpi lat/bw deviation FAILED
Done mpi lat/bw deviation Test Suite Fri Oct 07 16:40:18 EDT 2016

Hit any key to continue (or ESC to abort)...
```



The results display the pair-wise analysis of latency and bandwidth for the selected hosts and report pairs outside an acceptable tolerance range. By default, performance is compared relative to other hosts in the fabric. It is assumed that all hosts selected for a given run have comparable fabric performance. Failing hosts are clearly indicated.

Intel recommends that you review the `FF_DEVIATION_ARGS` parameter in `opafastfabric.conf` and adjust it as appropriate for the cluster. The default can accommodate a wide range of cluster designs.

The results are also written to the `test.res` file, which may be viewed through the **View opahostadmin result files** option. Refer to the *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide* for more details.

6. Press any key or `ESC` to end the operation.

### Additional Details

If any hosts fail, carefully examine the failing hosts to verify the HFI models, PCIe slot used, BIOS settings, and any motherboard or BIOS settings related to devices on PCIe buses or slot speeds. Also verify the HFI and any riser cards are properly seated.

The bandwidth that is reported should also be checked against the PCIe speeds in the Performance Impact table below. If all pairs are not in the expected performance range, carefully examine all hosts to verify the HFI models, PCIe slot used, BIOS settings and any motherboard or BIOS settings related to devices on PCIe buses or slot speeds. Also verify the HFI and any riser cards are properly seated.

**Table 8. Performance Impact**

PCIe Speed	Fabric Speed	Typical Bandwidth
PCIe 8GT/s x16 (Gen3)	100 Gbps	12.0 - 12.4 GBps
PCIe 8GT/s x8 (Gen3)	100 Gbps	6.4 - 6.8 GBps
PCIe 5GT/s x16 (Gen2)	100 Gbps	6.4 - 6.8 GBps
PCIe 5GT/s x8 (Gen2)	100 Gbps	3.2 - 3.4 GBps
<i>Note:</i> 1 GBps = 1,000,000,000 bytes/second		

## 7.10 Checking Overall Fabric Health

The **Check Overall Fabric Health** selection allows you to baseline the present fabric configuration for use in future fabric health checks. Perform this check after configuring any additional Management Nodes and establishing a healthy fabric via successful execution of all the other tests. If desired, a baseline of an incomplete or unhealthy fabric may be taken for future comparison after making additions or corrections to the fabric.

Refer to Configure and Initialize Health Check Tools in the *Intel® Omni-Path Fabric Software Installation Guide* for more information.

1. From the FastFabric OPA Host Verification/Admin menu, type **9**.

The menu item changes from `[Skip]` to `[Perform]`.

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.



2. Type **P** to begin the operation.

The status is displayed.

```
Performing Host Admin: Check Overall Fabric Health
Baseline present configuration? [n]:
```

3. Press **Enter** (n) to analyze the configuration without baselining it.

```
Executing: /usr/sbin/opaallanalysis
opafabricanalysis: Port 0:0 Error: Previous baseline run required
opafabricanalysis: Possible fabric errors or changes found
opachassisanalysis: Warning: showAllConfig command failed for 1 or more
chassis. See /var/usr/lib/opa/analysis/latest/chassis.showAllConfig
opachassisanalysis: Error: Chassis error. See /var/usr/lib/opa/analysis/
latest/chassis.hwCheck
opachassisanalysis: Possible Chassis errors or changes found
opaallanalysis: Possible errors or changes found
Hit any key to continue (or ESC to abort)...
```

4. Type **y** and press **Enter** to baseline the configuration.

The configuration is baselined.

```
Executing: /usr/sbin/opaallanalysis -b
opafabricanalysis: Port 0:0 Error: Unable to access fabric.
See /var/usr/lib/opa/analysis/latest/fabric.0:0.snapshot.stderr
opafabricanalysis: Possible fabric errors or changes found
opachassisanalysis: Warning: showAllConfig command failed for 1 or more
chassis. See /var/usr/lib/opa/analysis/latest/chassis.showAllConfig
opachassisanalysis: Baselined
opaallanalysis: Possible errors or changes found
Hit any key to continue (or ESC to abort)...
```

5. Press any key or ESC to end the operation.

## 7.11 Starting or Stopping Bit Error Rate Cable Test

The **Start or Stop Bit Error Rate Cable Test** selection allows you to perform host and/or ISL cable testing. The test allows for starting and stopping an extended Bit Error Rate test. The system prompts to clear hardware counters.

**Note:** Clearing of hardware counters (-A option) is optional and may affect the PM and other tools. See "PM Running Counters to Support opareport" section in the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* for more information.

Intel recommends that you run this test for 20-60 minutes for a thorough test. While the test is running, monitor the fabric for signal integrity or stability errors using opatop, opareport, and/or the Fabric Manager GUI. Once the desired test time has elapsed, return to this item in the menu and stop the test.

1. From the FastFabric OPA Host Verification/Admin menu, type **a**.

The menu item changes from [Skip] to [Perform].

**Note:** More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.
3. For each prompt, provide the required information and press **Enter**:



Prompt	Description
Stop or cleanup any already running Cable Test? [y]:	Allows you to stop and clean up any cable tests in process.
Stop HFI-Switch Cable Test? [y]:	Allows you to stop HFI-Switch cable test.
Stop ISL Cable Test? [y]:	Allows you to stop ISL cable test.
Start Cable Test? [y]:	Allows you to start a new cable test.
Clear error counters? [y]:	Allows you to clear the error counters.
Force Clear of hardware error counters too? [y]:	Allows you to clear hardware counters.
Start HFI-Switch Cable Test? [y]:	Allows you to start a new HFI-Switch cable test.
Start ISL Cable Test? [y]:	Allows you to start a new ISL cable test.

After executing the prompts, the following is displayed.

```
About to run: /usr/sbin/opacabletest -A
Hit any key to continue (or ESC to abort)...
```

- Press any key to execute the cabletest or ESC to end the operation.

## 7.12 Generating All Hosts Problem Report Information

**(Host)** The **Generate all Hosts Problem Report Info** selection allows you to collect configuration and status information from all hosts and generate a single \*.tgz file that can be sent to an Intel support representative.

- From the FastFabric OPA Host Verification/Admin menu, type **b**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

- Type **p** to begin the operation.

The status is displayed.

```
Performing Host Admin: Generate All Hosts Problem Report Info
Capture detail level (1-Normal 2-Fabric 3-Fabric+FDB 4-Analysis): [4]:
```

The detail levels are cumulative and shown below:

Detail Level	Description
1-Normal	Obtains local information from each host.
2-Fabric	In addition to "Normal", obtains basic fabric information by queries to the SM and fabric error analysis using iba_report.
continued...	



Detail Level	Description
3-Fabric+FDB	In addition to "Fabric", obtains all the switch forwarding tables and OPA multicast membership lists from the SM.
4-Analysis	In addition to "Fabric+FDB", obtains all <code>all_analysis</code> results. If <code>all_analysis</code> has not yet been run, it is run as part of the capture.
<b>Notes:</b> <ul style="list-style-type: none"><li>• Detail levels 2-4 can be used when fabric operational problems occur. If the problem appears to be node-specific, detail level 1 should be sufficient.</li><li>• Detail levels 2-4 require an operational Fabric Manager. Typically, your support representative requests a given detail level. If a given detail level takes excessively long or fails to be gathered, try a lower detail level.</li><li>• For detail levels 2-4, the additional information is only gathered on the node running the <code>captureall</code> command. The information is gathered for every fabric specified in the <code>/etc/sysconfig/opa/ports</code> file.</li></ul>	

3. Type the menu item for the level of details required for the report and press **Enter**.

`opacaptureall` is initiated and results gathered in a `hostcapture.all.tgz`.

A sample of a "Normal" analysis is shown below.

```
Executing: /usr/sbin/opacaptureall -p -D 1 -f /etc/sysconfig/opa/allhosts
Running capture on all non-local hosts ...
[root@phkpstl042]# rm -f ~root/hostcapture.tgz; opacapture ~root/
hostcapture.tgz
Getting software and firmware version information ...
Capturing FM binaries and debuginfo if available
Getting TMM information...
Obtaining OS configuration ...
Obtaining dmesg logs ...
Obtaining present process and module list ...
Obtaining PCI device list ...
ls: cannot access /dev/ipath*: No such file or directory
Obtaining environment variables ...
Obtaining network interfaces ...
Obtaining DMI information ...
Obtaining Shared Memory information ...
Obtaining OmniPath information ...
Obtaining MPI configuration ...
Copying configuration and statistics for OPA drivers from /proc ...
Obtaining additional CPU info...
Obtaining HFI statistics ...
Copying kernel debug information from /sys/kernel/debug/hfil...
Copying configuration and statistics for ib_ drivers from /sys ...
  Getting statedump for hfil_0 ...
Copying configuration and statistics for OPA from /sys/module ...
Gathering Host FM Information ...
Creating dump directory...
Getting systemd information...
Getting FM rpm version...
Copying FM configuration...
Copying FM core dumps...
Skipping SM 0: Not Running
Skipping SM 1: Not Running
Skipping SM 2: Not Running
Skipping SM 3: Not Running
Skipping SM 4: Not Running
Skipping SM 5: Not Running
Skipping SM 6: Not Running
Skipping SM 7: Not Running
Packaging capture file...
Saved FM capture as smdump-10Oct16095350.tgz
Gathering Distributed SA data...
Creating tar file /root/hostcapture.tgz ...
Done.
```



```

Please include /root/hostcapture.tgz with any problem reports to Customer
Support
Uploading capture from each host ...
Running capture on local host ...
scp root@[phpkstl042]:hostcapture.tgz ./uploads/phpkstl042/.
Getting software and firmware version information ...
hostcapture.tgz 100% 17MB 17.2MB/s 00:00
Capturing FM binaries and debuginfo if available
Getting TMM information...
Obtaining OS configuration ...
Obtaining dmesg logs ...
Obtaining present process and module list ...
Obtaining PCI device list ...
ls: cannot access /dev/ipath*: No such file or directory
Obtaining environment variables ...
Obtaining network interfaces ...
Obtaining DMI information ...
Obtaining Shared Memory information ...
Obtaining OmniPath information ...
Obtaining MPI configuration ...
Copying configuration and statistics for OPA drivers from /proc ...
Obtaining additional CPU info...
Obtaining HFI statistics ...
Copying kernel debug information from /sys/kernel/debug/hfil...
Copying configuration and statistics for ib_ drivers from /sys ...
Getting statedump for hfil_0 ...
Copying configuration and statistics for OPA from /sys/module ...
Gathering Host FM Information ...
Creating dump directory...
Getting systemd information...
Getting FM rpm version...
Copying FM configuration...
Copying FM core dumps...
Getting SM 0 counters...
Getting PM 0 counters...
Getting SM 0 run-time core file
Skipping SM 1: Not Running
Skipping SM 2: Not Running
Skipping SM 3: Not Running
Skipping SM 4: Not Running
Skipping SM 5: Not Running
Skipping SM 6: Not Running
Skipping SM 7: Not Running
Packaging capture file...
Saved FM capture as smdump-100Oct16095357.tgz
Gathering Distributed SA data...
Creating tar file /root/./uploads/phpkstl041/hostcapture.tgz ...
Done.

Please include /root/./uploads/phpkstl041/hostcapture.tgz with any problem
reports to Customer Support
Combining captured files into ./uploads/hostcapture.all.tgz ...
Done.
Hit any key to continue (or ESC to abort)...

```

4. Press any key or ESC to end the operation.

## 7.13 Running a Command on All Hosts

**(Linux)** The **Run a command on all hosts** selection allows you to perform other operations on all hosts. Each time this is executed, a Linux\* shell command may be specified to be executed against all selected hosts. You can also specify a sequence of commands separated by semicolons.

1. From the FastFabric OPA Host Verification/Admin menu, type **c**.



The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Admin: Run a Command on All Hosts
Enter Command to run on all hosts (or none):
```

3. Enter a Linux command and press **Enter**.

```
Timelimit in minutes (0=unlimited): [1]:
```

4. Specify a time limit and press **Enter**.

```
Run in parallel on all hosts? [y]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.

```
About to run: /usr/sbin/opacmdall -T 60 -f /etc/sysconfig/opa/hosts 'xxxx'
Are you sure you want to proceed? [n]:
```

6. Type **y** and press **Enter** to proceed with the operation.

The operation is completed.

## 7.14 Viewing opahostadmin Result Files

The **View opahostadmin result files** allows you to display the `test.log` and `test.res` files that contain the results from prior `opahostadmin` runs, such as installing fabric software or rebooting all hosts. You are also given the option to remove these files after viewing them.

If prior files are not removed, subsequent runs of `opachassisadmin`, `opahostadmin`, or `opaswitchadmin` from within the current directory continue to append to these files.

1. From the FastFabric OPA Host Verification/Admin menu, type **d**.

The menu item changes from [Skip] to [Perform].

*Note:* More than one menu item may be selected. The operations will perform individually and in sequence with the menu.

2. Type **P** to begin the operation.

```
Performing Host Admin: View opahostadmin Result Files
Using vi (to select a different editor, export EDITOR).
About to: vi /root/punchlist.csv /root/verifyhosts.res /root/test.res /root/
test.log
Hit any key to continue (or ESC to abort)...
```

3. Press any key to view the `opahostadmin` results files.





4. After reviewing and closing the log, you are prompted to remove the following files.

```
4 files to edit
Would you like to remove verifyhosts.res test.res test.log test_tmp* and
save_tmp
in /root ? [n]:
```

5. Select **y** (yes) or **n** (no) and press **Enter**.
6. If you chose **y** in the step above, press any key or **ESC** to end the operation



## 8.0 Monitoring Fabric Performance

Both the FastFabric OPA Fabric Monitoring menu and `opatop` CLI allow you to start up the Fabric Performance Monitoring TUI so that you can monitor the performance of the fabric.

The Fabric Performance Monitor TUI displays performance, congestion, and error information about a fabric. Fabric information is divided into two main starting points for analyzing fabric traffic:

- **Performance** (bandwidth utilization): Can identify over-utilized areas (bottle necks) and under-utilized areas (potentially mis-configured).
- **Errors Statistics**: Can identify problems in fabric hardware or configuration, as well as congestion and other performance situations.

### 8.1 Viewing the Fabric Performance Monitoring Summary Screen

The top-level Summary screen shows the basic fabric configuration information as well as performance and error information. This is the initial screen you see when you start up the TUI.

After looking at the Summary screen you can decide which area of the fabric (performance or error) and which port group or virtual fabric most warrants investigation, and can then drill down into that area.

To view the Fabric Performance Monitoring Summary screen, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter `opatop`.

The Summary screen is displayed.

```
opatop: Img: 10s @ Wed Sep 14 11:29:52 2016, Live
Summary: SW:      0 Ports: SW:      0 HFI:      2      Link:      1
          SM:      1 Node Fail:      0 Skip:      0 Port Fail:      0 Skip:      0
          AvgMBps  MinMBps  MaxMBps  AvgKPps  MinKPps  MaxKPps
0 All      Int      0        0        0        0        0        0
   Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
1 HFIs      Int      0        0        0        0        0        0
   Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
2 SWs      No ports in group

Master-SM: LID: 0x0001 Port: 1 Priority: 0 State: Master
           Name: phcppriv10 hfi1_0
           PortGUID: 0x0011750101575300
Secondary-SM: none
```



```
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

- To change to the Virtual Fabrics (VF) Summary screen, type **v**.

The VF Summary screen is shown as in the example below.

```
opatop: Img: 10s @ Thu Sep 22 15:20:07 2016, Live
Summary: SW:      0 Ports: SW:      0 HFI:      2      Link:      1
          SM:      1 Node Fail:      0 Skip:      0 Port Fail:      0 Skip:      0
          AvgMBps  MinMBps  MaxMBps  AvgKPPs  MinKPPs  MaxKPPs
0 Admin      Int      0      0      0      0      0      0
  Integ:min  Congst:min  SmaCong:min  Bubble:min  Secure:min  Routing:min

Master-SM: LID: 0x0001 Port: 1  Priority: 0  State: Master
           Name: phcppriv10 hfi1_0
           PortGUID: 0x0011750101575300
Secondary-SM: none

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

### Summary Screen Field Descriptions

The table below describes the Summary screen field descriptions.

**Table 9. Summary Screen Field Descriptions**

Field	Description
Fabric Configuration Information	<p>Fabric configuration information includes</p> <ul style="list-style-type: none"> <li>Numbers of links</li> <li>Numbers of switches (SW)</li> <li>Numbers of SMs</li> <li>Numbers of ports</li> <li>Master SM details</li> <li>Secondary SM details (if present)</li> </ul>
Performance and Error Statistics for Each Port Group	<p>Fabric performance and error statistics are presented based on port groupings and virtual fabrics grouping:</p> <p>For Port Groups:</p> <ul style="list-style-type: none"> <li>All <ul style="list-style-type: none"> <li>In the All group, all ports are Internal because, by definition, the neighbor port must be in the All group.</li> </ul> </li> <li>HFI <ul style="list-style-type: none"> <li>In the HFIs groups, all neighbor ports are outside the group, so statistics are contained in the Send and Receive subgroups.</li> </ul> </li> <li>SWs <ul style="list-style-type: none"> <li>In the SWs group, neighbor ports are either outside the group (HFI) or inside the group (another switch), so statistics are contained in all three subgroups. A special case for a switch port is the special switch port 0, which is always considered internal to the SWs group.</li> </ul> </li> </ul> <p>For Virtual Fabrics Group:</p> <ul style="list-style-type: none"> <li>Admin</li> <li>Default</li> </ul>

*continued...*



Field	Description
	<p>These groups provide a natural subdivision of the ports in a fabric for analysis. For more information about Groups and the operation of the PM, refer to the <i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>.</p> <p>For each group, the following statistics are reported:</p> <ul style="list-style-type: none"><li>• Average MBps (megabytes per second)</li><li>• Minimum MBps</li><li>• Maximum MBps</li><li>• Average KPs (kilopackets per second)</li><li>• Minimum KPs</li><li>• Maximum KPs</li><li>• Status indicator</li></ul>
Performance Statistics	<p>Performance statistics for each port group are further divided into up to three subgroups based on whether a port's neighbor port is in its group:</p> <ul style="list-style-type: none"><li>• Internal If a port's neighbor port is in its group, all performance statistics are contained in the Internal subgroup.</li><li>• Send If a port's neighbor is not in its group, statistics for data leaving the port (group) are contained in the Send subgroup</li><li>• Receive If a port's neighbor is not in its group, statistics for data entering the port are contained in the Receive subgroup.</li></ul>
Error Categories	<p>The error categories are:</p> <ul style="list-style-type: none"><li>• Integ – Integrity</li><li>• Congst – Congestion</li><li>• Bubble – Idles due to congestion</li><li>• SmaCong – SMA Congestion</li><li>• Secure – Security</li><li>• Routing – Routing</li></ul> <p>Error categories are each based on one or more port error counters. Each error category's status indicator is shown at one of five values/colors based on the error value as compared to a threshold value:</p> <ul style="list-style-type: none"><li>• Minimum – green</li><li>• Low – blue</li><li>• Moderate – cyan</li><li>• Warning – yellow</li><li>• OVER – red</li></ul>

## 8.2 Viewing the PM Configuration

The PM Configuration screen displays information as provided by the PM.

*Notes:*

- The PM Configuration screen is the same for VF and non-VF.
- The PM Configuration screen has no screen-specific input commands.

To view PM Configuration, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opatop**.  
The Summary screen is displayed.
3. Type **p**.



The PM Configuration screen is displayed as shown in the example below.

```

opatop: Img: 10s @ Thu Sep 22 15:23:17 2016, Live
PM Config:
  Sweep Interval: 10 sec  PM Flags(0x33):
    ProcessHFICntrs=On ProcessVLCntrs=On ClrDataCntrs=Off Clr64bitErrCntrs=Off
    Clr32bitErrCntrs=On Clr8bitErrCntrs=On
  Max Clients: 3
  Total Images: 10  Freeze Images: 5  Freeze Lease: 60 seconds
  Err Thresholds: Integrity: 100  Congestion: 100
                   SmaCongest: 100  Bubble: 100
                   Security: 10  Routing: 100
  Integrity Wts: Link Qual: 40  Uncorrectable: 100
                  Link Downed: 25  Rcv Errors: 100
                  Excs Bfr Ovrn: 100  FM Config Err: 100
                  Link Err Reco: 100  Loc Link Integ: 0
                  Lnk Wdth Dngd: 100
  Congest Wts: Cong Discards: 100  Rcv FECN: 5
                  Rcv BECN: 1  Mark FECN: 25
                  Xmit Time Cong 25  Xmit Wait: 10
  PM Memory Size: 169 MB (169295080 bytes)
  PMA MADs: MaxAttempts: 3  MinRespTimeout: 35  RespTimeout: 250
  Sweep: MaxParallelNodes: 10  PmaBatchSize: 2  ErrorClear: 7

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |

```

4. Type **u** (lowercase) to return to the Summary Screen.

### PM Configuration Screen Field Descriptions

For more information on field descriptions, refer to the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*.

The table below describes the PM Configuration screen field descriptions.

**Table 10. PM Configuration Field Descriptions**

Field	Description
Sweep Interval	The time over the image data is relevant. Default is 10 seconds <b>NOTE:</b> Normally, the opatop interval should be set to a value ≥ Sweep Interval.
PM Flags	Shows whether PM Flags are On or Off for: <ul style="list-style-type: none"> <li>• ProcessHFICntrs</li> <li>• ProcessVLCntrs</li> <li>• ClrDataCntrs</li> <li>• Clr64bitErrCntrs</li> <li>• Clr32bitErrCntrs</li> <li>• Clr8bitErrCntrs</li> </ul>
Max Clients	Maximum clients
Total Images	<ul style="list-style-type: none"> <li>• Freeze Images</li> <li>• Freeze Lease time</li> </ul>
Err Thresholds	Error thresholds <ul style="list-style-type: none"> <li>• Integrity - Integrity</li> <li>• Congestion - Congestion</li> <li>• Bubble - Idles due to congestion</li> <li>• SmaCongest - SMA Congestion</li> <li>• Security - Security</li> <li>• Routing - Routing</li> </ul>
<i>continued...</i>	



Field	Description
Integrity Wts	Integrity weights <ul style="list-style-type: none"><li>• Link Qual</li><li>• Uncorrectable</li><li>• Link Downed</li><li>• Rcv Errors</li><li>• Excs Bfr Ovrn</li><li>• FM Config Err</li><li>• Link Err Reco</li><li>• Loc Link Integ</li><li>• Lnk Wdth Dngd</li></ul>
Congest Wts	Congestion weights <ul style="list-style-type: none"><li>• Cong Discards</li><li>• Rcv FECN</li><li>• Rcv BECN</li><li>• Mark FECN</li><li>• Xmit Time Cong</li><li>• Xmit Wait</li></ul>
PM Memory Size	Size of the PM memory footprint in MB and bytes
PMA MADs	PMA MADs retry/timeout <ul style="list-style-type: none"><li>• MaxAttempts</li><li>• MinRespTimeout</li><li>• RespTimeout</li></ul>
Sweep	Sweep information <ul style="list-style-type: none"><li>• MaxParallelNodes</li><li>• PmaBatchSize</li><li>• ErrorClear</li></ul>

## 8.3 Viewing Image Information

The Image Information screen show the image information as provided by the PM.

- Notes:**
- The Image Information screen is the same for VF and non-VF.
  - The PM Configuration screen has no screen-specific input commands.

To view Image Information, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opatop**.

The Summary screen is displayed.

3. Type **I**.

The Image Info screen is displayed as shown in the example below.

```
opatop: Img: IIs @ Day Month Date HR:MIN:SEC YYYY, Live
Image Inopatop: Img: 10s @ Thu Sep 22 16:51:58 2016, Live
Image Info:
Sweep Start: Thu Sep 22 16:51:58 2016
Sweep Duration: 0.001 Seconds
Image Interval: 10 Seconds

Num SW-Ports:      0  HFI-Ports:      2
Num SWS:           0  Num Links:      1  Num SMs:      1
```



```

Num Fail Nodes:      0  Ports:      0  Unexpected Clear Ports: 0
Num Skip Nodes:      0  Ports:      0

    Master-SM: LID: 0x0001 Port: 1  Priority: 0  State: Master
                Name: phcppriv10 hfil_0
                PortGUID: 0x0011750101575300
    Secondary-SM: none

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |

```

4. Type **u** (lowercase) to return to the Summary Screen.

### Image Information Screen Field Descriptions

The table below describes the Image Information screen field descriptions.

**Table 11. Image Information Field Descriptions**

Field	Description
Sweep Start	Timestamp for the start of the sweep
Sweep Duration	Length of time for the sweep
Image Interval	The time over the image data is relevant. Default is 10 seconds
Num [Ports]	Number of ports in each group: <ul style="list-style-type: none"> <li>• SW-Ports</li> <li>• HFI-Ports</li> </ul>
Num SWs	Number of switches
Node Information	Node information including: <ul style="list-style-type: none"> <li>• Failed nodes</li> <li>• Skipped nodes</li> </ul>
Port Information	Port information including: <ul style="list-style-type: none"> <li>• Failed ports</li> <li>• Skipped ports</li> <li>• Unexpected clear ports</li> </ul>
SM Information	Master and secondary SM details <ul style="list-style-type: none"> <li>• LID</li> <li>• Port</li> <li>• Priority</li> <li>• State</li> <li>• Name</li> <li>• PortGUID</li> </ul>

## 8.4 Viewing Bandwidth Statistics

For each valid performance data subgroup, the Bandwidth Statistics screen displays the total, average, minimum, and maximum MBps and Kpps. For each subgroup, ten performance 'buckets' count the number of ports whose 'MBps compared to link rate' value corresponds to that bucket. This provides an indication of how the data rate of the group compares to its potential.

To view bandwidth statistics, perform the following steps:



1. Log in to the server as root.
2. At the command prompt, enter **opatop**.  
The Summary screen is displayed.
3. Determine which set of statistics you want to view:
  - To view Group information, continue to the next step.
  - To view VF information, type **v**.
4. Type the number for the specific group statistics that you want to view:

For Port Group:

- 0 – All
- 1 – HFIs
- 2 – SWs

For VF Group:

- 0 – Default
- 1 – Admin

The Info Select screen is displayed as shown in the example below.

```
opatop: Img: 10s @ Fri Sep 23 09:44:49 2016, Live
Group Info Sel: HFIs
Int NumPorts: 2   Rate Min: 100g   Max: 100g
Ext NumPorts: 0
  Group BW Summary (W)
  Group Err Summary (E)
  Group Config (C)

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | W E C:
```

5. Type **w**.

The Bandwidth (BW) Stats screen is displayed as shown in the example below.

```
opatop: Img: 10s @ Fri Sep 23 09:46:09 2016, Live
Group BW Stats: HFIs   Criteria: Util-High   Number: 10
Int: TotMBps AvgMBps MinMBps MaxMBps TotKPps AvgKPps MinKPps
MaxKPps
0      0      0      0      0      0      0
  Buckt 0+%   10+%   20+%   30+%   40+%   50+%   60+%   70+%   80+%   90+%
        2      0      0      0      0      0      0      0      0      0
  Failed Int Ports: PMA:    0   Topo:    0

Int Congestion      Max      0+%   25+%   50+%   75+%   100+%
                   0      2      0      0      0      0

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | cC N0-n Detail:
```





6. To set the BW stats Criteria for the focus query, type **c** (lowercase) to scroll forward or **C** (uppercase) to scroll in reverse to select one of the following choices:
  - **Util-High** – Bandwidth Utilization (highest first)
  - **UtilPkt-Hi** – Packet Utilization (highest first)
  - **Util-Low** – Bandwidth Utilization (lowest first)
7. To change the Number of entries in the BW stats list, type **N** and enter the target number of entries; then press **Enter**.
8. Type **D** to initiate the group focus query and access the detailed Group Focus screen (refer to [Viewing Focus Information](#) on page 119.)
9. Type **u** (lowercase) for each screen you've accessed until you are back to the screen you want.

### Bandwidth Statistics Screen Field Descriptions

The table below describes the bandwidth screen field descriptions.

**Table 12. Bandwidth Statistics Field Descriptions**

Field	Description
Group Name	<p>Name of the group examined</p> <p>For Port Groups:</p> <ul style="list-style-type: none"> <li>• <b>All</b> In the All group, all ports are Internal because, by definition, the neighbor port must be in the All group.</li> <li>• <b>HFIs</b> In the HFIs groups, all neighbor ports are outside the group, so statistics are contained in the Send and Receive subgroups.</li> <li>• <b>SWs</b> In the SWs group, neighbor ports are either outside the group (HFI) or inside the group (another switch), so statistics are contained in all three subgroups. A special case for a switch port is the special switch port 0, which is always considered internal to the SWs group.</li> </ul> <p>For Virtual Fabrics Group:</p> <ul style="list-style-type: none"> <li>• <b>Admin</b></li> <li>• <b>Default</b></li> </ul>
Criteria	<p>Focus criterion for Group Focus screen:</p> <ul style="list-style-type: none"> <li>• <b>Util-High</b> – Bandwidth Utilization (highest first)</li> <li>• <b>UtilPkt-Hi</b> – Packet Utilization (highest first)</li> <li>• <b>Util-Low</b> – Bandwidth Utilization (lowest first)</li> </ul>
Number	Number of ports for a group focus query
Performance Data Subgroup	<p>Performance statistics for each port group are further divided into up to three subgroups based on whether a port's neighbor port is in its group:</p> <ul style="list-style-type: none"> <li>• <b>Internal</b> If a port's neighbor port is in its group, all performance statistics are contained in the Internal subgroup.</li> <li>• <b>Send</b> If a port's neighbor is not in its group, statistics for data leaving the port (group) are contained in the Send subgroup</li> <li>• <b>Receive</b> If a port's neighbor is not in its group, statistics for data entering the port are contained in the Receive subgroup.</li> </ul>
Statistics	For each group, the following statistics are reported:

*continued...*



Field	Description
	<ul style="list-style-type: none"><li>• Average MBps</li><li>• Minimum MBps</li><li>• Maximum MBps</li><li>• Average KPps</li><li>• Minimum KPps</li><li>• Maximum KPps</li><li>• Status indicator</li></ul>
Performance Buckets	Count the number of ports whose 'MBps compared to link rate' value corresponds to that bucket. This provides an indication of how the data rate of the group compares to its potential. Ten buckets from 0+% to 90+%, in 10% increments
Failed Ports	Failed Ports per subgroup: <ul style="list-style-type: none"><li>• PMA PMA failures are port counter query failures during the PM Sweep.</li><li>• Topo Topology errors are failures caused by encountering missing neighbor information in the topology.</li></ul>
Congestion buckets	Provides context (from the Errors Screen) <ul style="list-style-type: none"><li>• Max</li><li>• 0+%</li><li>• 25+%</li><li>• 50+%</li><li>• 75+%</li><li>• 100+%</li></ul>

## 8.5 Viewing Error Statistics

The Error Statistics screen displays error statistics for a port group.

To view error statistics, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opatop**.  
The Summary screen is displayed.
3. Determine which set of statistics you want to view:
  - To view Group information, continue to the next step.
  - To view VF information, type **V**.
4. Type the number for the specific group statistics that you want to view:  
For Port Group:
  - 0 – All
  - 1 – HFIs
  - 2 – SWsFor VF Group:
  - 0 – Default
  - 1 – Admin



The Info Select screen is displayed as shown in the example below.

```
opatop: Img: 10s @ Fri Sep 23 09:44:49 2016, Live
Group Info Sel: HFIs
Int NumPorts: 2   Rate Min: 100g   Max: 100g
Ext NumPorts: 0
  Group BW Summary (W)
  Group Err Summary (E)
  Group Config (C)

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | W E C:
```

5. Type **E**.

The Error (Err) Stats screen is displayed as shown in the example below.

```
opatop: Img: 10s @ Fri Sep 23 11:55:09 2016, Live
Group Err Stats: HFIs Criteria: Integ Number: 10
Int      Max      0+%      25+%      50+%      75+%      100+%
Integrity      0        2        0        0        0        0
Congestion     0        2        0        0        0        0
SmaCongest     0        2        0        0        0        0
Bubble         0        2        0        0        0        0
Security       0        2        0        0        0        0
Routing        0        2        0        0        0        0
Utilization:   0.0% Discards: 0.0%
```

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | cC N0-n Detail:

6. To set the error stats Criteria for the focus query, type **c** (lowercase) to scroll forward or **C** (uppercase) to scroll in reverse to select one of the following choices:
  - Integrity errors (highest first)
  - Congestion errors (highest first)
  - SmaCongestion errors (highest first)
  - Bubble errors (highest first)
  - Security errors (highest first)
  - Routing routing (highest first)
7. To change the Number of entries in the Err Stats list, type **N** and enter the target number of entries; then press **Enter**.
8. Type **D** to initiate the group focus query and access the detailed Group Focus screen (refer to [Viewing Focus Information](#) on page 119.)
9. Type **u** (lowercase) for each screen you've accessed until you are back to the screen you want.

### Bandwidth Statistics Screen Field Descriptions

The table below describes the bandwidth screen field descriptions.



**Table 13. Bandwidth Statistics Field Descriptions**

Field	Description
Group Name	<p>Name of the group examined</p> <p>For Port Groups:</p> <ul style="list-style-type: none"> <li>All In the All group, all ports are Internal because, by definition, the neighbor port must be in the All group. All ports are Internal</li> <li>HFI In the HFIs groups, all neighbor ports are outside the group, so statistics are contained in the Send and Receive subgroups. All ports are External</li> <li>SWs In the SWs group, neighbor ports are either outside the group (HFI) or inside the group (another switch), so statistics are contained in all three subgroups. A special case for a switch port is the special switch port 0, which is always considered internal to the SWs group. Ports are Internal and External.</li> </ul> <p>For Virtual Fabrics Group:</p> <ul style="list-style-type: none"> <li>Admin</li> <li>Default</li> </ul>
Criteria (Error Categories)	<p>Focus criteria/error categories:</p> <ul style="list-style-type: none"> <li>Integrity <ul style="list-style-type: none"> <li>Link Quality Indicator</li> <li>Link Width Downgrade</li> <li>Local Link Integrity Errors</li> <li>Port Receive Erros</li> <li>Excessive Buffer Overrun Errors (neighbor port)</li> <li>Link Error Recovery</li> <li>Link Downed</li> <li>Uncorrectable Errors</li> <li>FM Config Errors</li> </ul> </li> <li>Congestion <ul style="list-style-type: none"> <li>Port Transmit Wait</li> <li>Switch Port Congestion</li> <li>Port Receive FECN (neighbor port)</li> <li>Port Receive BECN (only from FIs)</li> <li>Port Transmit Time Congestion</li> <li>Port Mark FECN</li> </ul> </li> <li>SmaCongestion The counters included in the SMA Congestion category are the VL 15 counters equivalent to the port counters in the Congestion category.</li> <li>Bubble <ul style="list-style-type: none"> <li>Port Transmit Wasted Bandwidth</li> <li>Port Transmit Wait Data</li> <li>Port Receive Bubble (neighbor port)</li> </ul> </li> <li>Security <ul style="list-style-type: none"> <li>Port Receive Constraint Errors (neighbor port)</li> <li>Port Transmit Constraint Errors</li> </ul> </li> <li>Routing <ul style="list-style-type: none"> <li>Port Receive Switch Relay Errors</li> </ul> </li> </ul> <p>The integrity and congestion error values are calculated by using a weighted sum. The weights for each and the threshold value for each error category can be seen in the PM Configuration screen (<a href="#">PM Configuration Screen Field Descriptions</a> on page 109). For more details about how the values for each error category is composed, refer to the <i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>.</p>
continued...	



Field	Description
Number	Number of entries for a group focus query
Performance Data Subgroup	<p>Performance statistics for each port group are further divided into up to three subgroups based on whether a port's neighbor port is in its group:</p> <ul style="list-style-type: none"> <li>• Internal If a port's neighbor port is in its group, all performance statistics are contained in the Internal subgroup.</li> <li>• Send If a port's neighbor is not in its group, statistics for data leaving the port (group) are contained in the Send subgroup</li> <li>• Receive If a port's neighbor is not in its group, statistics for data entering the port are contained in the Receive subgroup.</li> </ul>
Int or Ext	<p>Location of the port in relation to the group.</p> <ul style="list-style-type: none"> <li>• Int – The port's neighbor port is in its group (internal).</li> <li>• Ext – The port's neighbor port is not in its group (external).</li> </ul>
Error buckets	<p>For each error subgroup, the error buckets count the number of ports whose "error compared to error threshold" value corresponds to that bucket. This provides an indication of how error rates compare to their thresholds.</p> <ul style="list-style-type: none"> <li>• Max</li> <li>• 0+%</li> <li>• 25+%</li> <li>• 50+%</li> <li>• 75+%</li> <li>• 100+%</li> </ul>
Utilization	Percent of error utilization; aids congestion analysis.
Discards	Percent of errors discarded; aids congestion analysis.

## 8.6 Viewing Configuration Information

The Configuration screen displays a list of the ports in a group, including the LID, port number, port GUID, and NodeDesc for each.

To view configuration information, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opatop**.  
The Summary screen is displayed.
3. Determine which set of statistics you want to view:
  - To view Group information, continue to the next step.
  - To view VF information, type **V**.
4. Type the number for the specific group statistics that you want to view:  
For Port Group:
  - 0 – All
  - 1 – HFIs
  - 2 – SWs
 For VF Group:



- 0 – Default
- 1 – Admin

The Info Select screen is displayed as shown in the example below.

```
opatop: Img: 10s @ Fri Sep 23 09:44:49 2016, Live
Group Info Sel: HFIs
Int NumPorts: 2 Rate Min: 100g Max: 100g
Ext NumPorts: 0
  Group BW Summary (W)
  Group Err Summary (E)
  Group Config (C)

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | W E C:
```

5. Type **c**.

The Config screen is displayed as shown in the example below.

```
opatop: Img: 10s @ Fri Sep 23 12:07:29 2016, Live
Group Config: HFIs NumPorts: 2
  Ix LIDx Port Node GUID 0x NodeDesc
    0 0001 1 0011750101575300 phcpprivl0 hfil_0
    1 0002 1 001175010157E443 phcpprivl1 hfil_0

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | sS P0-n:
```

6. Type **s** (lowercase) to scroll forward or **S** (uppercase) to scroll backward through multiple screens of a long port list.
7. Type **P** and enter the target Ix number; then press **Enter** to view the Port Stats screen for the specified Ix (refer to [Viewing Port Statistics](#) on page 122).
8. Type **u** (lowercase) for each screen you've accessed until you are back to the screen you want.

### Configuration Information Screen Field Descriptions

The table below describes the Configuration screen field descriptions.

**Table 14. Configuration Information Field Descriptions**

Field	Description
Group Name	Name of the group examined For Port Groups: <ul style="list-style-type: none"><li>• All In the All group, all ports are Internal because, by definition, the neighbor port must be in the All group.</li><li>• HFIs In the HFIs groups, all neighbor ports are outside the group, so statistics are contained in the Send and Receive subgroups.</li><li>• SWs</li></ul>
continued...	



Field	Description
	<p>In the SWs group, neighbor ports are either outside the group (HFI) or inside the group (another switch), so statistics are contained in all three subgroups. A special case for a switch port is the special switch port 0, which is always considered internal to the SWs group.</p> <p>For Virtual Fabrics Group:</p> <ul style="list-style-type: none"> <li>Admin</li> <li>Default</li> </ul>
NumPorts	Number of ports returned in the group configuration query
Ix	An index value that is used to select a port to view in the Port Stats screen.
LIDx	LID information
Port	Port Index
Node GUID 0x	Global Unique Identifier (GUID) for the Node
NodeDesc	Description of the node

## 8.7 Viewing Focus Information

The Focus information screen displays a list of the ports within a group, including the LID, port number, focus criterion, port GUID and NodeDesc of each. If the port has a neighbor port, the same information is displayed for the neighbor.

**Note:** The Focus information screen is the same for VF and non-VF.

To view focus information, perform the following steps:

- Log in to the server as root.
- At the command prompt, enter **opatop**.  
The Summary screen is displayed.
- Determine which set of statistics you want to view:
  - To view Group information, continue to the next step.
  - To view VF information, type **V**.
- Type the number for the specific group statistics that you want to view:  
For Port Group:
  - 0 – All
  - 1 – HFIs
  - 2 – SWs
 For VF Group:
  - 0 – Default
  - 1 – Admin
 The Info Select screen is displayed.
- Determine the Information Select menu to access:
  - To view the Focus information screen for BW Summary, type **W**.
  - To view the Focus information screen for Err Summary, type **E**.



6. Determine the Criteria for the focus query:

- To set the BW stats Criteria for the focus query, type **c** (lowercase) to scroll forward or **C** (uppercase) to scroll in reverse to select one of the following choices:
  - Util-High – Bandwidth Utilization (highest first)
  - UtilPkt-Hi – Packet Utilization (highest first)
  - Util-Low – Bandwidth Utilization (lowest first)
- To set the error stats Criteria for the focus query, type **c** (lowercase) to scroll forward or **C** (uppercase) to scroll in reverse to select one of the following choices:
  - Integrity errors (highest first)
  - Congestion errors (highest first)
  - SmaCongestion errors (highest first)
  - Bubble errors (highest first)
  - Security errors (highest first)
  - Routing routing (highest first)

7. Type **D**.

The Focus information screen is displayed as shown in the example below.

```
opato: Img: 10s @ Fri Sep 23 13:03:09 2016, Live
Group Focus: HFIs GrpNumPorts: 2 NumPorts: 1 Number: 10
Ix Util-High LIDx Port Node GUID 0x NodeDesc
  0      0.0 0001  1 0011750101575300 phcppriv10 hfil_0
<->      0.0 0002  1 001175010157E443 phcppriv11 hfil_0
```

```
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | sS cC N0-n P0-n:
```

8. To change the criteria after accessing this screen, type **c** (lowercase) to scroll forward or **C** (uppercase) to scroll in reverse to select one of the following choices:

- Util-High – Bandwidth Utilization (highest first)
- UtilPkt-Hi – Packet Utilization (highest first)
- Util-Low – Bandwidth Utilization (lowest first)
- Integrity errors (highest first)
- Congestion errors (highest first)
- SmaCongestion errors (highest first)
- Bubble errors (highest first)
- Security errors (highest first)
- Routing routing (highest first)

9. To change the Number of entries in the focus list, type **N** and enter the target number of entries; then press **Enter**.





10. Type **s** (lowercase) to scroll forward or **S** (uppercase) to scroll backward through multiple screens of a long port list.
11. Type **P** and enter the target Ix number; then press **Enter** to view the detailed Port Stats screen (refer to [Viewing Port Statistics](#) on page 122).
12. Type **u** (lowercase) for each screen you've accessed until you are back to the screen you want.

### Focus Information Screen Field Descriptions

The table below describes the Focus screen field descriptions.

**Table 15. Focus Information Field Descriptions**

Field	Description
Group Name	<p>Name of the group examined</p> <p>For Port Groups:</p> <ul style="list-style-type: none"> <li>All In the All group, all ports are Internal because, by definition, the neighbor port must be in the All group.</li> <li>HFIIs In the HFIIs groups, all neighbor ports are outside the group, so statistics are contained in the Send and Receive subgroups.</li> <li>SWs In the SWs group, neighbor ports are either outside the group (HFI) or inside the group (another switch), so statistics are contained in all three subgroups. A special case for a switch port is the special switch port 0, which is always considered internal to the SWs group.</li> </ul> <p>For Virtual Fabrics Group:</p> <ul style="list-style-type: none"> <li>Admin</li> <li>Default</li> </ul>
GrpNumPorts	Number of ports selected, as determined by the combination of group, criteria, and requested ports
NumPorts	Number of ports returned in the group configuration query
Number	Number of ports for a group focus query
Ix	An index value that is used to select a port to view in the Port Stats screen.
Criteria	<p>Limits the focus to specific port statistics</p> <p>For BW stats:</p> <ul style="list-style-type: none"> <li>Util-High – Bandwidth Utilization (highest first)</li> <li>UtilPkt-Hi – Packet Utilization (highest first)</li> <li>Util-Low – Bandwidth Utilization (lowest first)</li> </ul> <p>For Err stats:</p> <ul style="list-style-type: none"> <li>Integrity errors (highest first)</li> <li>Congestion errors (highest first)</li> <li>SmaCongestion errors (highest first)</li> <li>Bubble errors (highest first)</li> <li>Security errors (highest first)</li> <li>Routing routing (highest first)</li> </ul>
LIDx	LID information
Port	Port Index

*continued...*



Field	Description
	<b>NOTE:</b> A symbol may be present on the first character of each line related to a port. This symbol is used to indicate a non-ideal condition was observed when calculating the relevant port's data. The possible conditions are, the PM was told to ignore this port ('~'), the PM failed to query this port ('!'), and the PM topology does not know this port's identity ('?').
Node GUID 0x	Global Unique Identifier (GUID) for the Node
NodeDesc	Description of the node


## 8.8 Viewing Port Statistics

The Port Statistics screen displays a specific port and LID's performance and error counters.

*Note:* The Port Statistics screen is the same for VF and non-VF.

To view port statistics, perform the following steps:

1. Log in to the server as root.
2. At the command prompt, enter **opatop**.  
The Summary screen is displayed.
3. Determine which set of statistics you want to view:
  - To view Group information, continue to the next step.
  - To view VF information, type **V**.
4. Type the number for the specific group statistics that you want to view:  
For Port Group:
  - 0 – All
  - 1 – HFIs
  - 2 – SWsFor VF Group:
  - 0 – Default
  - 1 – AdminThe Info Select screen is displayed.
5. Determine the Information Select menu to access:
  - To view the Port Stats screen for BW Summary, type **W**.
  - To view the Port Stats screen for Err Summary, type **E**.
  - To view the Port Stats screen for Configuration information, type **C**.  
If you are accessing the Port Stats screen from the Configuration information screen, skip to Step 7.
6. Determine the Criteria for the focus query as described in [Viewing Bandwidth Statistics](#) on page 111 or [Viewing Error Statistics](#) on page 114.
7. Type **D** to access the Focus information screen.



To make changes to the Focus information prior to accessing the Port Stats screen, refer to [Viewing Focus Information](#) on page 119.

8. Type **P** and enter the target Ix number; then press **Enter** to view the detailed Port Stats screen.

The Port Stats screen is displayed.

**Note:** Neighbor port and link information is available only when access through the Focus Information screen. It is not available through the Configuration information screen.

```

opatop: Img: 10s @ Fri Sep 23 14:07:40 2016, Live
Port Stats: HFIs LID: 0x2 PortNum: 1 Rate: 100g MTU: 4096
NodeDesc: phcppriv11 hfil_0 NodeGUID: 0x001175010157E443
Neighbor: phcppriv10 hfil_0 LID: 0x1 PortNum: 1
Xmit: Data: 0 MB ( 63 Flits) Pkts: 1
Rcv: Data: 0 MB ( 10 Flits) Pkts: 1
Multicast: Xmit Pkts: 0 Rcv Pkts: 0
Integrity: | Congestion:
Link Quality: 5 | Cong Discards: 0
Uncorrectable: 0 | Rcv FECN*: 0
Link Downed: 0 | Rcv BECN: 0
Lanes Down: 0 | Mark FECN: 0
Rcv Errors: 0 | Xmit Time Cong: 0
Excs Bfr Ovrn*: 0 | Xmit Wait: 0
FM Conf Err: 0 | Routing and Others:
Lnk Err Recov: 0 | Rcv Sw Relay: 0
Loc Lnk Integ: 0 | Xmit Discards: 0
Security: | Bubble:
Xmit Constrain: 0 | Xmit Wasted BW: 0
Rcv Constrain*: 0 | Xmit Wait Data: 0
SmaCongestion (VL15): | Rcv Bubble*: 0
Cong Discards: 0
Xmit Wait: 0
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | Neighbor |

```

9. Type **N** to switch between statistics for the port and its neighbor port.
10. Type **u** (lowercase) for each screen you've accessed until you are back to the screen you want.

### Port Statistics Screen Field Descriptions

The table below describes the Port Statistics screen field descriptions.

**Table 16. Port Statistics Field Descriptions**

Field	Description
Group Name	<p>Name of the group examined</p> <p>For Port Groups:</p> <ul style="list-style-type: none"> <li>All           <p>In the All group, all ports are Internal because, by definition, the neighbor port must be in the All group.</p> </li> <li>HFIs           <p>In the HFIs groups, all neighbor ports are outside the group, so statistics are contained in the Send and Receive subgroups.</p> </li> <li>SWs           <p>In the SWs group, neighbor ports are either outside the group (HFI) or inside the group (another switch), so statistics are contained in all three subgroups. A special case for a switch port is the special switch port 0, which is always considered internal to the SWs group.</p> </li> </ul> <p>For Virtual Fabrics Group:</p>

*continued...*



Field	Description
	<ul style="list-style-type: none"> <li>Admin Default</li> </ul>
LIDx	LID information for the node
PortNum	Port number of the node
Rate	link rate
MTU	MTU, if available
NodeDesc	Description of the node
NodeGUID	Global Unique Identifier (GUID) for the Node
Neighbor	Description of the neighboring node
Xmit Data	Size of the data transmitted in MB and Flits and the number of packets
Recv Data	Size of the data received in MB and Flits and the number of packets
Multicast: Xmit Pkts	Number of multicast packets transmitted
Multicast: Recv Pkts	Number of multicast packets received
Error Counters	<ul style="list-style-type: none"> <li>Integrity               <ul style="list-style-type: none"> <li>Link Quality</li> <li>Uncorrectable</li> <li>Link Downed</li> <li>Lanes Down</li> <li>Receive Errors</li> <li>Excessive Buffer Overrun*</li> <li>FM Config Errors</li> <li>Link Error Recovery</li> <li>Local Link Integrity</li> </ul> </li> <li>Security               <ul style="list-style-type: none"> <li>Transmit Constraint</li> <li>Receive Constraint*</li> </ul> </li> <li>SmaCongestion               <p>The counters included in the SMA Congestion category are the VL 15 counters equivalent to the port counters in the Congestion category.</p> <ul style="list-style-type: none"> <li>Cong Discards</li> <li>Xmit Wait</li> </ul> </li> <li>Congestion               <ul style="list-style-type: none"> <li>Cong Discards</li> <li>Receive FECN*</li> <li>Receive BECN*</li> <li>Mark FECN</li> <li>Transmit Time Congestion</li> <li>Transmit Wait</li> </ul> </li> <li>Routing and Others               <ul style="list-style-type: none"> <li>Receive Sw Relay</li> <li>Transmit Discards</li> </ul> </li> <li>Bubble               <ul style="list-style-type: none"> <li>Transmit Wasted Bandwidth</li> <li>Transmit Wait Data</li> <li>Receive Bubble*</li> </ul> </li> </ul> <p>A trailing asterisk (*) on the counter name indicates the count will be used in computing Error Category information for the neighbor port.</p>



## 8.9 Navigating PM Sweeps

The Fabric Performance Monitoring TUI allows you to access statistics from sequential PM sweeps (the PM keeps a history of previous sweep images) and queries the PM at a user-specified interval (10 seconds by default). Sweeps are accessed from the short term history database being recorded by the PM. This allows access to statistics from up to 24 hours in the past.

When the Fabric Performance Monitoring TUI queries for statistics for the most recent PM sweep, it is in "Live" mode. In Live mode, the data will change, at the `opato` interval rate, as `opato` queries new PM sweeps. At each screen (summary or detail), the data being displayed is refreshed for the current PM sweep.

A PM sweep can be in "frozen" mode. The data in a frozen sweep will not change, allowing the statistics to be examined in summary and detail screens.

The Fabric Performance Monitoring TUI allows you to navigate the focus to another sweep within the history of sweeps maintained by the PM. For the duration of focus on such a sweep, it will remain frozen. You cannot navigate to any other screen while in "Historic" mode. Navigation can be performed for the screen in focus backward or forward, 1 or 5 sweeps at a time.

To navigate the historical PM sweeps, perform the following steps:

1. Navigate to the screen that you want to analyze historically.
2. Type `r` (lowercase) to go back one sweep at a time.

The date stamp below shows the time of the freeze (highlighted in bold) and the current on-going time (highlighted in italics).

```
opato: Img: 10s @ Fri Sep 23 17:32:32 2016, Hist Now: Fri Sep 23 17:33:08 2016
Port Stats: HFIs LID: 0x1 PortNum: 1 Rate: 100g MTU: 4096
NodeDesc: phcpriv10 hfil_0 NodeGUID: 0x0011750101575300
Neighbor: phcpriv11 hfil_0 LID: 0x2 PortNum: 1
Xmit: Data: 0 MB ( 10 Flits) Pkts: 1
Recv: Data: 0 MB ( 63 Flits) Pkts: 1
Multicast: Xmit Pkts: 0 Recv Pkts: 0
Integrity: | Congestion:
Link Quality: 5 | Cong Discards: 0
Uncorrectable: 0 | Rcv FECN*: 0
Link Downed: 0 | Rcv BECN: 0
Lanes Down: 0 | Mark FECN: 0
Rcv Errors: 0 | Xmit Time Cong: 0
Excs Bfr Ovrn*: 0 | Xmit Wait: 0
FM Conf Err: 0 | Routing and Others:
Lnk Err Recov: 0 | Rcv Sw Relay: 0
Loc Lnk Integ: 0 | Xmit Discards: 0
Security: | Bubble:
Xmit Constrain: 0 | Xmit Wasted BW: 0
Rcv Constrain*: 0 | Xmit Wait Data: 0
SmaCongestion (VL15): | Rcv Bubble*: 0
Cong Discards: 0
Xmit Wait: 0
Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help | Neighbor |
```

3. Type `R` (uppercase) to go back five sweeps at a time.
4. Type `f` (lowercase) to move ahead one sweep at a time.
5. Type `F` (uppercase) to move ahead five sweeps at a time.
6. Type `L` to return to the Live data.



## 8.10 Bookmarking a Sweep

The Fabric Performance Monitoring TUI allows you to bookmark a sweep to review the information. For the duration of the Bookmark, all information is frozen. You can navigate through the various screens to review the frozen information. The sweep will remain frozen until you explicitly "Unbookmark" it.

### Adding a Bookmark

**Note:** opatop allows only one sweep at a time to be bookmarked.

To bookmark a PM sweep, perform the following steps:

1. Navigate to the screen you want to capture and analyze.
2. Type **B** (uppercase) to bookmark the screen.

In the Image Identification line (line 1), the Live image changes to Bkmk (bookmark) as highlighted in bold in the example screen below.

```
opatop: Img: 10s @ Fri Sep 23 16:44:42 2016, Bkmk Now: Fri Sep 23 16:44:53 2016
Summary: SW:      0 Ports: SW:      0 HFI:      2      Link:      1
        SM:      1 Node Fail:      0 Skip:      0 Port Fail:      0 Skip:      0
              AvgMBps    MinMBps    MaxMBps    AvgKPps    MinKPps    MaxKPps
0 All      Int      0      0      0      0      0      0
  Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
1 HFIs      Int      0      0      0      0      0      0
  Integ:min Congst:min SmaCong:min Bubble:min Secure:min Routing:min
2 SWs      No ports in group

      Master-SM: LID: 0x0001 Port: 1 Priority: 0 State: Master
                Name: phcppriv10 hfil_0
                PortGUID: 0x0011750101575300
      Secondary-SM: none

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:
```

The bookmark will remain until you explicitly remove it.

3. Type **L** to return to the Live data.
4. Type **b** (lowercase) to return to the bookmarked image.

### Removing a Bookmark

To remove a bookmark from a PM sweep, perform the following steps:

1. Type **b** (lowercase) to return to the bookmarked image.
2. Type **U** (uppercase).

In the Image Identification line (line 1), the Bkmk image changes back to Live (bookmark) as highlighted in bold in the example screen below.

```
opatop: Img: 10s @ Fri Sep 23 16:49:52 2016, Live
Summary: SW:      0 Ports: SW:      0 HFI:      2      Link:      1
        SM:      1 Node Fail:      0 Skip:      0 Port Fail:      0 Skip:      0
```



```

                                AvgMBps   MinMBps   MaxMBps   AvgKPps   MinKPps   MaxKPps
0 All          Int          0           0           0           0           0           0
   Integ:min Congst:min   SmaCong:min   Bubble:min   Secure:min   Routing:min
1 HFIs        Int          0           0           0           0           0           0
   Integ:min Congst:min   SmaCong:min   Bubble:min   Secure:min   Routing:min
2 SWs          No ports in group

Master-SM: LID: 0x0001 Port: 1   Priority: 0   State: Master
           Name: phcppriv10 hfi1_0
           PortGUID: 0x0011750101575300
Secondary-SM: none

Quit up Live/rRev/fFwd/bookmrked Bookmrk Unbookmrk ?help |
sS Pmcfg Imginfo View 0-n:

```

## 8.11 Using the opatop Command Line Options

While `opatop` starts the Fabric Performance Monitoring TUI, you can use the command line options as shown below:

### Syntax

```
opatop [-v] [-q] [-h hfi] [-p port] [-i seconds]
```

### Options

<code>--help</code>	Produces full help text.
<code>-v/--verbose level</code>	Specifies the verbose output level. Value is additive and includes: <ul style="list-style-type: none"> <li>1 Screen</li> <li>4 STDERR opatop</li> <li>16 STDERR PaClient</li> </ul>
<code>-q/--quiet</code>	Disables progress reports.
<code>-h/--hfi hfi</code>	Specifies the HFI, numbered 1..n. Using 0 specifies that the <code>-p port</code> port is a system-wide port number. (Default is 0.)
<code>-p/--port port</code>	Specifies the port, numbered 1..n. Using 0 specifies the first active port. (Default is 0.)
<code>-i/--interval seconds</code>	Interval in <i>seconds</i> at which PA queries are performed to refresh to the latest PA image. Default = 10 seconds.



**-h and -p options permit a variety of selections:**

- h 0 First active port in system (default).
- h 0 -p 0 First active port in system.
- h x First active port on HFI x.
- h x -p 0 First active port on HFI x.
- h 0 -p y Port y within system (no matter which ports are active).
- h x -p y HFI x, port y.