



# **Intel® Omni-Path Fabric**

## **Setup Guide**

---

***Rev. 12.0***

***January 2020***



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.

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

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 © 2016–2020, Intel Corporation. All rights reserved.



## Revision History

Date	Revision	Description
January 2020	12.0	Updates to this document include: <ul style="list-style-type: none"> <li>Updated links to Intel® Omni-Path documentation set in <a href="#">Intel® Omni-Path Documentation Library</a>.</li> <li>Updated Download Software links in <a href="#">Before You Begin</a> and <a href="#">Programming and Verifying Option ROM EEPROM Device</a>.</li> </ul>
October 2019	11.0	Updates to this document include: <ul style="list-style-type: none"> <li>Updated Preface to include new <a href="#">Best Practices</a>.</li> <li>Added new section <a href="#">Customizing MPI</a> including <a href="#">Building MPI Libraries</a> and <a href="#">Building and Running MPI Applications</a>.</li> </ul>
December 2018	10.0	Updates to this document include: <ul style="list-style-type: none"> <li>Globally, changed OS RPMs Installation Prerequisites pointer from <i>Intel® Omni-Path Fabric Software Installation Guide</i> to <i>Intel® Omni-Path Fabric Software Release Notes</i>.</li> <li>Updated Installing Drivers to include installing using the IFS-included Linux* OS repository.</li> </ul>
September 2018	9.0	Updates to this document include: <ul style="list-style-type: none"> <li>Globally, performed minor edits to text, format, and structure for consistency and usability.</li> </ul>
April 2018	8.0	Updates to this document include: <ul style="list-style-type: none"> <li>Added How to Search the Intel Omni-Path Documentation Set to Preface.</li> <li>Updated Generate Cable Map Topology Files, Defining Type in the Topology Spreadsheet, and Creating the Topology File.</li> <li>Updated Programming and Verifying Option ROM EEPROM Device.</li> <li>Updated Running Bandwidth Test.</li> <li>Updated Running Latency Test.</li> <li>Updated Running MPI Deviation Test.</li> <li>Updated Running run_mpi_stress.</li> <li>Updated Analyzing Links.</li> </ul>
October 2017	7.0	Updates to this document include: <ul style="list-style-type: none"> <li>Updated Before You Begin, Complete Installation Prerequisites, Install Intel® Omni-Path Software, and Link Debug CLI Commands.</li> <li>Changed bios_images filepath in Programming and Verifying Option ROM EEPROM Device to /usr/share/opa/bios_images.</li> <li>The <i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i> has been merged into the <i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>. In this document, all references have been updated appropriately. See the Intel® Omni-Path Documentation Library for details.</li> </ul>
August 2017	6.0	Updates to this document include: <ul style="list-style-type: none"> <li>Updated Configuring BIOS Settings.</li> <li>Added NVIDIA* bullet in the Installing Drivers section of Install Intel® Omni-Path Software.</li> <li>Clerical changes.</li> </ul>
April 2017	5.0	Updates to this document include: <div style="text-align: right;"><i>continued...</i></div>



Date	Revision	Description
		<ul style="list-style-type: none"><li>• Changed the title of this document from <i>Intel® Omni-Path Fabric Staging Guide</i> to <i>Intel® Omni-Path Fabric Setup Guide</i>.</li><li>• Updated Decoding the Physical Configuration of an HFI section with TMM details.</li><li>• Added Programming and Verifying Option ROM EEPROM Device section.</li><li>• Globally, updated the following filepaths:<ul style="list-style-type: none"><li>— from <code>/etc/sysconfig/opafm.xml</code> to <code>/etc/opa-fm/opafm.xml</code></li><li>— from <code>/usr/lib/opa/src/</code> to <code>/usr/src/opa/</code></li><li>— from <code>/etc/sysconfig/</code> to <code>/etc/</code></li></ul></li><li>• Added Intel® Omni-Path Documentation Library to Preface.</li></ul>
December 2016	4.0	Updates to this document include: <ul style="list-style-type: none"><li>• Updated Take State Dump of a Switch procedure.</li><li>• Globally, updated the following filepaths:<ul style="list-style-type: none"><li>— from <code>/opt/opa</code> to <code>/usr/lib/opa</code></li><li>— from <code>/var/opt/opa</code> to <code>/var/usr/lib/opa</code></li><li>— from <code>/opt/opafm</code> to <code>/usr/lib/opa-fm</code></li><li>— from <code>/var/opt/opafm</code> to <code>/var/usr/lib/opa-fm</code></li></ul></li><li>• Added Cluster Configurator for Intel® Omni-Path Fabric to Preface.</li></ul>
October 2016	3.0	Document has been updated to add the following sections: <ul style="list-style-type: none"><li>• Performing Initial Fabric Verification</li><li>• Editing Hosts and Allhosts Files</li><li>• Defining Type in the Topology Spreadsheet</li><li>• Configure Intel Omni-Path Director Class Switch 100 Series</li><li>• Configure Host Setup</li><li>• Modifying the Fabric Manager Routing Algorithm</li></ul>
August 2016	2.0	Updates to this document include: <ul style="list-style-type: none"><li>• Updated Generate Cable Map Topology Files to include Intel® Omni-Path Director Class Switch 100 Series information.</li></ul>
May 2016	1.0	Initial release.




## Contents

---

<b>Revision History.....</b>	<b>3</b>
<b>Preface.....</b>	<b>8</b>
Intended Audience.....	8
Intel® Omni-Path Documentation Library.....	8
How to Search the Intel® Omni-Path Documentation Set.....	10
Cluster Configurator for Intel® Omni-Path Fabric.....	11
Documentation Conventions.....	11
Best Practices.....	12
License Agreements.....	12
Technical Support.....	12
<b>1.0 Introduction.....</b>	<b>13</b>
1.1 Before You Begin.....	13
<b>2.0 Complete Installation Prerequisites.....</b>	<b>15</b>
2.1 Configuring BIOS Settings.....	15
2.2 Configuring OS Settings .....	15
<b>3.0 Resolve TCP/IP Host Names.....</b>	<b>17</b>
<b>4.0 Install Intel® Omni-Path Software.....</b>	<b>18</b>
4.1 Installing Drivers.....	18
4.2 Verifying HFI Speed and Bus Width Using lspci.....	19
4.3 Performing Initial Fabric Verification.....	19
4.4 Editing Hosts and Allhosts Files.....	20
<b>5.0 Generate Cable Map Topology Files.....</b>	<b>21</b>
5.1 Defining Type in the Topology Spreadsheet.....	21
5.2 Creating the Topology File.....	22
<b>6.0 Configure FastFabric.....</b>	<b>23</b>
6.1 Formatting for IPoIB Host Names.....	23
6.2 Specifying Test Areas for opaallanalysis.....	23
6.3 Modifying the Location of mpi_apps Directory.....	23
<b>7.0 Configure Managed Intel® Omni-Path Edge Switches.....</b>	<b>25</b>
<b>8.0 Configure Intel® Omni-Path Director Class Switch 100 Series.....</b>	<b>27</b>
<b>9.0 Configure Externally-Managed Intel® Omni-Path Edge Switches.....</b>	<b>28</b>
<b>10.0 Configure Host Setup.....</b>	<b>31</b>
<b>11.0 Verify Cable Map Topology.....</b>	<b>32</b>
<b>12.0 Verify Server and Fabric.....</b>	<b>33</b>
<b>13.0 Apply Best Known Methods for Site Installation.....</b>	<b>34</b>
13.1 Enabling Intel® Omni-Path Fabric Manager GUI.....	34
13.2 Review Server and Fabric Verification Test Results.....	35
13.3 Debugging Intel® Omni-Path Physical Link Issues.....	36



13.3.1 OPA Link Transition Flow.....	36
13.3.2 Verify the Fabric Manager is Running.....	37
13.3.3 Check the State of All Links in the System.....	37
13.3.4 Check the State of HFI Links from a Server.....	38
13.3.5 Link Width, Downgrades, and opafm.xml.....	38
13.3.6 How to Check Fabric Connectivity.....	39
13.3.7 Physical Links Stability Test Using opacabletest.....	39
13.3.8 How to Debug and Fix Physical Link Issues.....	40
13.3.9 Link Debug CLI Commands.....	41
13.4 Using opatop to View Bandwidth and Error Summary.....	42
13.5 Using the Beacon LED to Identify HFI and Switch Ports.....	42
13.6 Decoding the Physical Configuration of an HFI.....	43
13.7 Programming and Verifying Option ROM EEPROM Device.....	44
13.8 Verifying Fabric Manager Sweep.....	46
13.9 Verifying PM Sweep Duration.....	46
13.10 Checking Credit Loop Operation.....	47
13.11 Modifying the Fabric Manager Routing Algorithm.....	47
<b>14.0 Run Benchmark and Stress Tests.....</b>	<b>48</b>
14.1 Running Bandwidth Test.....	48
14.2 Running Latency Test.....	48
14.3 Running MPI Deviation Test.....	48
14.4 Running run_mpi_stress.....	48
<b>15.0 Take State Dump of a Switch.....</b>	<b>50</b>
<b>16.0 Apply Best Known Methods for OPA Commands.....</b>	<b>51</b>
16.1 Retrieving the Host Fabric Interface Temperature.....	51
16.2 Reading Error Counters.....	51
16.3 Clearing Error Counters.....	52
16.4 Loading and Unloading Intel® Omni-Path Host HFI Driver.....	52
16.5 Analyzing Links.....	52
16.6 Tracing a Route between Two Nodes.....	53
16.7 Analyzing All Fabric ISLs Routing Balance.....	53
16.8 Dumping Switch ASIC Forwarding Tables.....	53
16.9 Configuring Redundant Fabric Manager Priority.....	54
16.9.1 FM Priority from a Local or Remote Terminal.....	54
16.9.2 FM Elevated Priority.....	54
16.9.3 Configuration Consistency for Priority/Elevated Priority.....	54
16.9.4 FM states from the Management Node.....	55
<b>17.0 Perform Final Fabric Checks.....</b>	<b>56</b>
<b>18.0 Customizing MPI.....</b>	<b>57</b>
18.1 Building MPI Libraries.....	57
18.2 Building and Running MPI Applications.....	57



## Tables

1	HFI Temperature Output Definitions.....	51
2	Link Quality Values and Description.....	53



## 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 the next generation, High-Performance Computing (HPC) network solution that is designed to cost-effectively meet the growth, density, and reliability requirements of large-scale HPC clusters.

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

## Intended Audience

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

## Intel® Omni-Path Documentation Library


Intel® Omni-Path publications are available at the following URL, under *Latest Release Library*:

<https://www.intel.com/content/www/us/en/design/products-and-solutions/networking-and-io/fabric-products/omni-path/downloads.html>

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

Task	Document Title	Description
Using the Intel® OPA documentation set	<i>Intel® Omni-Path Fabric Quick Start Guide</i>	A roadmap to Intel's comprehensive library of publications describing all aspects of the product family. This document outlines the most basic steps for getting your Intel® Omni-Path Architecture (Intel® OPA) cluster installed and operational.
Setting up an Intel® OPA cluster	<i>Intel® Omni-Path Fabric Setup Guide</i>	Provides a high level overview of the steps required to stage a customer-based installation of the Intel® Omni-Path Fabric. Procedures and key reference documents, such as Intel® Omni-Path user guides and installation guides, are provided to clarify the process. Additional commands and best known methods are defined to facilitate the installation process and troubleshooting.
continued...		





Task	Document Title	Description
Installing hardware	<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>	Describes the hardware installation and initial configuration tasks for the Intel® Omni-Path Switches 100 Series. This includes: Intel® Omni-Path Edge Switches 100 Series, 24 and 48-port configurable Edge switches, and Intel® Omni-Path Director Class Switches 100 Series.
	<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>	Contains instructions for installing the HFI in an Intel® OPA cluster.
Installing host software Installing HFI firmware Installing switch firmware (externally-managed switches)	<i>Intel® Omni-Path Fabric Software Installation Guide</i>	Describes using a Text-based User Interface (TUI) to guide you through the installation process. You have the option of using command line interface (CLI) commands to perform the installation or install using the Linux* distribution software.
Managing a switch using Chassis Viewer GUI Installing switch firmware (managed switches)	<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	Describes the graphical user interface (GUI) of the Intel® Omni-Path Fabric Chassis Viewer GUI. This document provides task-oriented procedures for configuring and managing the Intel® Omni-Path Switch family. Help: GUI embedded help files
Managing a switch using the CLI Installing switch firmware (managed switches)	<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	Describes the command line interface (CLI) task information for the Intel® Omni-Path Switch family. Help: -help for each CLI
Managing a fabric using FastFabric	<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>	Provides instructions for using the set of fabric management tools designed to simplify and optimize common fabric management tasks. The management tools consist of Text-based User Interface (TUI) menus and command line interface (CLI) commands. Help: -help and man pages for each CLI. Also, all host CLI commands can be accessed as console help in the Fabric Manager GUI.
Managing a fabric using Fabric Manager	<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>	The Fabric Manager uses a well defined management protocol to communicate with management agents in every Intel® Omni-Path Host Fabric Interface (HFI) and switch. Through these interfaces the Fabric Manager is able to discover, configure, and monitor the fabric.
	<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>	Provides an intuitive, scalable dashboard and set of analysis tools for graphically monitoring fabric status and configuration. This document is a user-friendly alternative to traditional command-line tools for day-to-day monitoring of fabric health. Help: Fabric Manager GUI embedded help files
Configuring and administering Intel® HFI and IPoIB driver Running MPI applications on Intel® OPA	<i>Intel® Omni-Path Fabric Host Software User Guide</i>	Describes how to set up and administer the Host Fabric Interface (HFI) after the software has been installed. The audience for this document includes cluster administrators and Message-Passing Interface (MPI) application programmers.
Writing and running middleware that uses Intel® OPA	<i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>	Provides a reference for programmers working with the Intel® PSM2 Application Programming Interface (API). The Performance Scaled Messaging 2 API (PSM2 API) is a low-level user-level communications interface.
<b>continued...</b>		



Task	Document Title	Description
Optimizing system performance	<i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>	Describes BIOS settings and parameters that have been shown to ensure best performance, or make performance more consistent, on Intel® Omni-Path Architecture. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.
Designing an IP or LNet router on Intel® OPA	<i>Intel® Omni-Path IP and LNet Router Design Guide</i>	Describes how to install, configure, and administer an IPoIB router solution (Linux* IP or LNet) for inter-operating between Intel® Omni-Path and a legacy InfiniBand* fabric.
Building Containers for Intel® OPA fabrics	<i>Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note</i>	Provides basic information for building and running Docker* and Singularity* containers on Linux*-based computer platforms that incorporate Intel® Omni-Path networking technology.
Writing management applications that interface with Intel® OPA	<i>Intel® Omni-Path Management API Programmer's Guide</i>	Contains a reference for programmers working with the Intel® Omni-Path Architecture Management (Intel OPAMGT) Application Programming Interface (API). The Intel OPAMGT API is a C-API permitting in-band and out-of-band queries of the FM's Subnet Administrator and Performance Administrator.
Using NVMe* over Fabrics on Intel® OPA	<i>Configuring Non-Volatile Memory Express* (NVMe*) over Fabrics on Intel® Omni-Path Architecture Application Note</i>	Describes how to implement a simple Intel® Omni-Path Architecture-based point-to-point configuration with one target and one host server.
Learning about new release features, open issues, and resolved issues for a particular release	<i>Intel® Omni-Path Fabric Software Release Notes</i>	
	<i>Intel® Omni-Path Fabric Manager GUI Release Notes</i>	
	<i>Intel® Omni-Path Fabric Switches Release Notes</i> (includes managed and externally-managed switches)	
	<i>Intel® Omni-Path Fabric Unified Extensible Firmware Interface (UEFI) Release Notes</i>	
	<i>Intel® Omni-Path Fabric Thermal Management Microchip (TMM) Release Notes</i>	
	<i>Intel® Omni-Path Fabric Firmware Tools Release Notes</i>	

## How to Search the Intel® Omni-Path Documentation Set

Many PDF readers, such as Adobe\* Reader and Foxit\* Reader, allow you to search across multiple PDFs in a folder.

Follow these steps:

1. Download and unzip all the Intel® Omni-Path PDFs into a single folder.
2. Open your PDF reader and use **CTRL-SHIFT-F** to open the Advanced Search window.
3. Select **All PDF documents in...**
4. Select **Browse for Location** in the dropdown menu and navigate to the folder containing the PDFs.
5. Enter the string you are looking for and click **Search**.

Use advanced features to further refine your search criteria. Refer to your PDF reader Help for details.



## Cluster Configurator for Intel® Omni-Path Fabric

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

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

## Documentation Conventions

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

- *Note:* provides additional information.
- **Caution:** indicates the presence of a hazard that has the potential of causing damage to data or equipment.
- **Warning:** indicates the presence of a hazard that has the potential of causing personal injury.
- Text in [blue](#) font indicates a hyperlink (jump) to a figure, table, or section in this guide. Links to websites are also shown in blue. For example:  
See [License Agreements](#) on page 12 for more information.  
For more information, visit [www.intel.com](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.



## Best Practices

- Intel recommends that users update to the latest versions of Intel® Omni-Path firmware and software to obtain the most recent functional and security updates.
- To improve security, the administrator should log out users and disable multi-user logins prior to performing provisioning and similar tasks.

## License Agreements

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

## Technical Support

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



## 1.0 Introduction

---

This document provides a high-level overview of the steps required to set up an Intel® Omni-Path Fabric. Procedures and key reference documents, such as Intel® Omni-Path user guides and installation guides are provided to clarify the process. Additional commands and BKMs are defined to facilitate the installation process and troubleshooting.

Intel recommends that you use the Intel® Omni-Path FastFabric (FF) Text-based User Interface (TUI) as the initial tool suite for installation, configuration, and validation of the fabric. This tool includes a set of automated features that are specifically used for standalone host, Ethernet\*, and Intel® Omni-Path Fabric connectivity validation.

This document includes recommendations for processes and procedures that complement the FF tools to reduce the time required to install and configure the fabric.

You should check applicable release notes and technical advisories for key information that could influence installation steps outlined in this document.

This document assumes the following:

- Operating System (OS) Software is a release-supported OS. See the *Intel® Omni-Path Fabric Software Release Notes* for the complete list of supported OSes.
- Single Management Node (with Fabric Manager running) configured with the Intel® Omni-Path Fabric Suite Software, also known as IntelOPA-IFS.
- Intel® Omni-Path Fabric Manager enabled on management nodes.
- Compute Nodes configured with the Intel® Omni-Path Fabric Host Software, also known as IntelOPA-Basic.
- Password-less access enabled for all hosts and switches.

---

### NOTE

Before you run top500 HPL (High Performance Linpack) runs or customer acceptance tests, Intel recommends that you follow all steps outlined in this guide.

---

## 1.1 Before You Begin

- Familiarize yourself with the available documentation by referring to [Intel® Omni-Path Documentation Library](#) on page 8 in this document.
- Printable checklists that contain an overview of the required steps for installing the Intel® OPA software are located in the *Intel® Omni-Path Fabric Software Installation Guide*, Software Installation Checklists section.
- Download software from: <https://www.intel.com/content/www/us/en/design/products-and-solutions/networking-and-io/fabric-products/omni-path/downloads.html>, Latest Release Library



- Fabric Host Software: IFS package for management nodes.
- Fabric Host Software: Basic package for compute nodes.
- Switch Firmware \*.emfw: for externally-managed switches.
- Switch Firmware \*.spkg: for managed switches.
- HFI Platform Firmware: for UEFI, TMM, and Firmware Tools (optional)
- Intel® Omni-Path Fabric Suite Fabric Manager GUI software.
- Make sure you have access to OS packages and some extended packages that are prerequisites for installing the Intel Fabric Suite software, as described in the *Intel® Omni-Path Fabric Software Release Notes*, OS RPMs Installation Prerequisites.
- Intel recommends that you develop a cable map topology file to validate the physical installation of the cables and ensure no accidental hot-spots are created through incorrect cabling. It is also a beneficial extension to any datacenter and rack layout diagrams developed for an installation. Refer to the instructions described in [Defining Type in the Topology Spreadsheet](#) on page 21 in this document.



## 2.0 Complete Installation Prerequisites

---

The recommended fabric installation prerequisites are defined in the *Intel® Omni-Path Fabric Software Installation Guide*, Pre-Installation Requirements section.

The RPMs required for the operating system you are using are defined in the *Intel® Omni-Path Fabric Software Release Notes*, OS RPMs Installation Prerequisites section.

Complete the following steps before starting software installation:

1. Install HFIs in servers, install externally-managed and managed switches in racks, and connect cabling.

Details: *Intel® Omni-Path Fabric Switches Hardware Installation Guide* and *Intel® Omni-Path Host Fabric Interface Installation Guide*.

- All standard EMI protection should be adhered to whenever working with computer equipment. The same is true for switch and HFI installation.
- HFIs should be installed in PCIe\* x16 slots of the server. The same PCIe\* slot should be used on all similar servers if possible to ensure homogeneity in the cluster.
- Switches should be installed based on switch guidelines and pre-designed rack placement.
- Care must be taken when connecting the cables, please adhere to bend radius requirements and use cable trays for management. Never use zip-ties to cinch the cables together, use soft fabric Velcro\* ties instead.

2. Power up all fabric hardware and verify LEDs operate as expected.

Details: *Intel® Omni-Path Fabric Switches Hardware Installation Guide* and *Intel® Omni-Path Host Fabric Interface Installation Guide*.

### 2.1 Configuring BIOS Settings

Intel recommends that you use UEFI BIOS. For optimal performance, refer to a recommended BIOS configuration in the *Intel® Omni-Path Fabric Performance Tuning User Guide*, BIOS Settings sections for the following processors:

- Intel® Xeon® Processor E5 v3 Family and Intel® Xeon® Processor E5 v4 Family
- Intel® Xeon® Scalable Processors
- Intel® Xeon Phi™ x200 Product Family

### 2.2 Configuring OS Settings


Before you install Intel® Omni-Path software, perform the following tasks:

- Confirm Operating System (OS) versions match the versions listed in the *Intel® Omni-Path Fabric Software Release Notes*.
- Install OS RPM prerequisites listed in the *Intel® Omni-Path Fabric Software Release Notes*, OS RPMs Installation Prerequisites section.



- Configure OS settings for optimal performance as described in the *Intel® Omni-Path Fabric Performance Tuning User Guide*, Linux\* Settings section.





## 3.0 Resolve TCP/IP Host Names

For details on resolving TCP/IP Host Names, see the *Intel® Omni-Path Fabric Software Installation Guide*, Fabric Setup Prerequisites section.

Create a `/etc/hosts` file before starting Intel® Omni-Path software installation to simplify the process. In a typical installation, the server and switch names follow a local convention to indicate physical location or purpose of the node.

- If using `/etc/hosts`, update the `/etc/hosts` file on the management node (the node with IFS installed) and copy to all hosts.
- If using DNS, all Management Network and IPoIB hostnames must be added to DNS `/etc/resolv.conf` and configured on the management node.
- The `/etc/hosts` file should contain:
  - Local host, for subsequent single host verification using FastFabric TUI
  - Ethernet and IPoIB addresses and names for all hosts
  - Ethernet addresses and names of switches
  - Ethernet addresses of IPMI or remote management modules
  - Ethernet addresses of power domain

An example of these recommendations follows:

```
# /etc/hosts example
# localhost (required)
127.0.0.1 localhost localhost.localdomain localhost4 localhost4.localdomain
# Ethernet Addresses of hosts
10.128.196.14    node1
10.128.196.15    node2
10.128.196.16    node3
# IPoIB Address of hosts should be outside Ethernet network
10.128.200.14    node1-opa
10.128.200.15    node2-opa
10.128.200.16    node3-opa
# RMM IP Addresses
10.127.240.121   node1-rmm
10.127.240.122   node2-rmm
# Chassis IP Address
10.128.198.250    opaedge1
10.128.198.249    opaedge2
# OPA director switch IP Address
10.128.198.251    opadirector1
10.128.198.252    opadirector2
```

Other files that may need adjustment according to specific site requirements include:

- `/etc/hostname`
- `/etc/resolv.conf`
- `/etc/network`
- `/etc/network-scripts/ifcfg-*` (for example, `ifcfg-enp5s0f0`)



## 4.0 Install Intel® Omni-Path Software

You should configure at least one node to run the Intel® Omni-Path Management Software including Fabric Manager (FM). This node is used to configure and validate all of the other hosts, switches, and chassis fabric devices. You must install the Intel® Omni-Path Fabric Suite software on this node.

The following document and sections describe the installation procedures:

- *Intel® Omni-Path Fabric Software Installation Guide*, Download the Intel® Omni-Path Software and Unpack the Tar File sections
- *Intel® Omni-Path Fabric Software Installation Guide*, Install the Intel® Omni-Path Software section

### 4.1 Installing Drivers

Driver installation is a two-step process. First, install the IFS package on your management servers, then use FastFabric to install the drivers on the rest of the systems in the cluster in parallel.

#### NOTE

If you use an existing cluster manager, installing the drivers to the image being installed on the compute nodes instead of using FastFabric is another option for quick driver installation.

#### Procedure

The following steps provide a summary for installation:

Step	Task/Prompt	Action
1.	Install IntelOPA-IFS on management node(s) usually designated to run Subnet Manager (SM) and FastFabric Tools (including MPI applications).	<p>Install using one of the following methods:</p> <ul style="list-style-type: none"><li>• Change the directory to <code>/IntelOPA-IFS.DISTRO.VERSION</code> and use the <code>./INSTALL</code> command.<ul style="list-style-type: none"><li>— Intel recommends that you use the Intel® Omni-Path FastFabric (FF) Text-based User Interface (TUI) as the tool suite for installation, configuration, and validation of the fabric. See <i>Install Using the TUI Menus</i> in the <i>Intel® Omni-Path Fabric Software Installation Guide</i>.</li><li>— Alternatively, CLI commands can be used. See <i>Install Using CLI Commands</i> in the <i>Intel® Omni-Path Fabric Software Installation Guide</i>.</li></ul></li></ul> <p>Examples include: <code>./INSTALL -n</code> or <code>./INSTALL -a -G</code></p> <ul style="list-style-type: none"><li>• Install using the IFS-included Linux* OS repository.</li></ul>

*continued...*



Step	Task/Prompt	Action
		See Install Using Linux* Distribution Software Packages Provided by Intel in the <i>Intel® Omni-Path Fabric Software Installation Guide</i>
2.	(optional) Upgrade the HFI UEFI firmware and install the Firmware Tools.	Refer to the <i>Intel® Omni-Path Fabric Software Installation Guide</i> , Install and Upgrade Standalone Firmware section.
3.	Intel recommends that you enable servers with IPMI interfaces to support ACPI or equivalent remote power management and reset control via an Ethernet network.	
4.	Apply Technical Advisories as needed.	
5.	(optional) Install NVIDIA* software.	<p>GPUDirect* RDMA is an NVIDIA* technology that allows third party network adapters to directly read and write to CUDA* host and device memory to enhance performance for latency and bandwidth. Intel provides CUDA* and GPUDirect* RDMA enhanced versions in both the Omni-Path HFI driver and PSM2 protocol libraries. Intel also provides a custom CUDA* enabled version of OpenMPI for PSM2.</p> <p>For installation details: see the <i>Intel® Omni-Path Fabric Software Installation Guide</i>.</p> <p>For usage details: see the <i>Intel® Omni-Path Fabric Host Software User Guide</i> and the <i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>.</p>
	<b>End Task</b>	

## 4.2 Verifying HFI Speed and Bus Width Using lspci

After the IFS installation, verify the Intel® OP HFI card is configured and visible to the host OS as Gen3 x16 slot speed (values are in **bold** text):

```
lspci -d 8086:24f0 -vv |grep Width
LnkCap: Port #0, Speed 8GT/s, Width x16, ASPM L1, Exit Latency L0s
<4us, L1 <64us
LnkSta: Speed 8GT/s, Width x16, TrErr- Train- SlotClk+ DLActive- BWMgmt- ABWMgmt-
```


## 4.3 Performing Initial Fabric Verification

### Procedure

Perform the following steps:

Step	Task/Prompt	Action
1.	Verify the port state of host is Active.	Run <code>opainfo</code> . If the command fails or returns other port state than Active, verify that the SM is running using <code>systemctl status opafm</code> .
2.	Run <code>opainfo</code> on all nodes in the fabric	Use <code>opacmdall</code> or <code>pdsh</code> .
3.	Verify the OPA software version is the same on all nodes.	Run <code>opaconfig -V</code> .
4.	Verify all nodes, switches, SM, and ISLs are up.	Run <code>opafabricinfo</code> as shown in the following example.  <pre>opafabricinfo Fabric 0:0 Information: SM: node hfi1_0 Guid: 0x001175010165b116 State: Master Number of HFIs: 126</pre>

*continued...*



Step	Task/Prompt	Action
		Number of Switches: 9 Number of Links: 252 Number of HFI Links: 126 (Internal: 0 External: 126) Number of ISLs: 126 (Internal: 0 External: 126) Number of Degraded Links: 0 (HFI Links: 0 ISLs: 0) Number of Omitted Links: 0 (HFI Links: 0 ISLs: 0)
5.	Review the number of HFIs, number of switches, and external ISLs and confirm that they match the fabric design.	The number of HFIs and external ISLs provide a fabric-blocking factor. If there are any degraded links, further troubleshooting is required.
	<b>End Task</b>	

## 4.4 Editing Hosts and Allhosts Files

Edit the following files, which are used by the `opafastfabric.conf` file.

- Edit `/etc/opa/hosts`

This file contains all hosts except the management node running IFS.

- Edit `/etc/opa/allhosts`

This file contains the statement `include /etc/opa/hosts`. Edit the file to add the node(s) running IFS.



## 5.0 Generate Cable Map Topology Files

Two sample topology files, `detailed_topology.xlsx` and `minimal_topology.xlsx`, are provided in the release.

For complete details on the sample files, see the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Sample Topology Spreadsheet Overview section.

- Both sample topology files, `detailed_topology.xlsx` and `minimal_topology.xlsx`, are spreadsheets with 3 tabs:
  - Tab 1 Fabric is for you to customize with EXTERNAL links.
  - Tab 2 swd06 contains the internal links for an Intel® OP Edge Switch 100 Series.
  - Tab 3 swd24 contains the internal links for an Intel® OP Director Class Switch 100 Series.

### NOTE

You should **not** modify tab 2 and 3.

- `README.topology` and `README.xlat_topology` describe best practices for editing the sample topology files.

For descriptions of other sample files provided in the package, see the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, List of Files section.

## 5.1 Defining Type in the Topology Spreadsheet

All host nodes should be defined Type = FI in column F of the spreadsheet. All Edge switches should be defined as Type = SW in column L (destination from host to Edge) and column F (source for Edge to core that is also Edge switch). The following example shows links between host and Edge switch.

```
R19 opahost1 1 FI R19 opaedge1 13 SW opahost1_opaelp13 1m Cable CU
```

All links between Edge switch to core that is also an Edge switch should be defined Type = SW, as shown in the following example:

```
row1 rack01 opaedge1 1 SW row1 rack04 opaedgecore1 2 SW opaelp1_opac1p2 5M Cable Fiber
```



All Director switches should be defined as Type = CL in column L (destination from Edge switch to Director switch). Column J (Name-2) should have the destination leaf and column K should have the port number on that leaf. The following example shows a link between an Edge switch to core that is a Director switch.

```
R19 opaedge1 5 SW R72 opadirector1 01 L105B 11 CL opaelp5opad1L105Bp11 30m Fiber
```

All 24-leaf chassis Director switches should be defined as shown in the following example:

```
Core Name:opadirector1 Core Group:row1 Core Rack:rack72 Core Size:1152 Core Full:0
```

Set Core Full to 0 if the Director switch is not fully populated with all the leafs and spines. If it is fully populated, set Core Full to 1.

For complete details, see the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Sample Topology Spreadsheet Overview section.

## 5.2 Creating the Topology File

To create a topology file, perform the following steps:

1. Copy and save either `detailed_topology.xlsx` or `minimal_topology.xlsx` located in `/usr/share/opa/samples` from the Fabric Manager node to your local PC for editing.
2. Edit tab 1 in the spreadsheet to reflect your specific installation details. Save tab 1 as `<topologyfile>.csv` and copy this CSV file back to the Fabric Manager node.

---

### NOTE

The cable label field can be up to 57 characters.

---

3. Generate the topology file in XML format using the following command and your customized CSV file as the source:

```
# opaxlattopology <topologyfile>.csv <topologyfile>.xml
```

If there are Director switches defined in the CSV file, then `opaxlattopology` includes all the ISL (internal chassis links between leafs and spines) in the XML file.

For more details, see the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, `opaxlattopology` section.



## 6.0 Configure FastFabric

The list of configuration files that are used by FastFabric are contained in the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Configuration Files for FastFabric section.

The `opafastfabric.conf` file provides default settings for most of the FastFabric command line options.

### 6.1 Formatting for IPoIB Host Names

By default, FastFabric uses the suffix OPA for the IPoIB host name. You can change this to a prefix and you can also change from opa to another convention such as ib, as the customer requires in `/etc/opa/opafastfabric.conf`.

The following examples show how to change opa to ib as a prefix or suffix.

For suffix:

```
export FF_IPoIB_SUFFIX=${FF_IPoIB_SUFFIX:--opa} to export FF_IPoIB_SUFFIX=${FF_IPoIB_SUFFIX:--ib}
```

For prefix:

```
export FF_IPoIB_PREFIX=${FF_IPoIB_PREFIX:-opa-} to export FF_IPoIB_PREFIX=${FF_IPoIB_PREFIX:-ib-}
```

### 6.2 Specifying Test Areas for opaallanalysis

By default, `opaallanalysis` includes the fabric and chassis. These can be modified to include host SM, embedded SM, and externally-managed switches in `/etc/opa/opafastfabric.conf` as follows:

```
# pick appropriate type of SM to analyze
#export FF_ALL_ANALYSIS=${FF_ALL_ANALYSIS:-fabric chassis hostsm esm}
export FF_ALL_ANALYSIS=${FF_ALL_ANALYSIS:-fabric chassis hostsm}
```

### 6.3 Modifying the Location of mpi\_apps Directory

By default, `opafastfabric` uses `mpi_apps` located in `/usr/src/opa/mpi_apps`. If a different path is set up for `mpi_apps`, then modify the following in `/etc/opa/opafastfabric.conf`:

```
export FF_MPI_APPS_DIR=${FF_MPI_APPS_DIR:-/usr/src/opa/mpi_apps}
```



---

**NOTE**

The default source location for `mpi_apps` is `/usr/src/opa/mpi_apps`, but the default compilation directory for `mpi_apps` is `$HOME/mpi_apps`.

---





## 7.0 Configure Managed Intel® Omni-Path Edge Switches

For a complete description of the configuration process, refer to the *Intel® Omni-Path Fabric Software Installation Guide*, Configure the Chassis section.

For more information about the managed switch, refer to the *Intel® Omni-Path Fabric Switches Hardware Installation Guide* and *Intel® Omni-Path Fabric Switches Release Notes*.

### Procedure

The following steps provide a summary for configuring the Edge Switch:

Step	Task/Prompt	Action
1.	Install the driver file CDM v2.12.00 WHQL Certified.exe	Download the file from: <a href="http://www.ftdichip.com/Drivers/VCP.htm">http://www.ftdichip.com/Drivers/VCP.htm</a>
2.	Set up USB serial port terminal emulator .	Use the following serial options: <ul style="list-style-type: none"> <li>• Speed: <b>115200</b></li> <li>• Data Bits: <b>8</b></li> <li>• Stop Bits: <b>1</b></li> <li>• Parity: <b>None</b></li> <li>• Flow Control: <b>None</b></li> </ul>
3.	Set up the switch TCP/IP address, gateway, netmask, and other options using a terminal emulator.	Notes: <ul style="list-style-type: none"> <li>• The changes are effective immediately.</li> <li>• For details, refer to the <i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>.</li> </ul>
	a) Set the chassis IP address.	Use the command: <code>setChassisIpAddr -h ipaddress -m netMask</code> where <i>ipaddress</i> is the new IP address in dotted decimal format (xxx.xxx.xxx.xxx), and <i>netMask</i> is the new subnet mask in dotted decimal format.
	b) Change the chassis default gateway IP address.	Use the command: <code>setDefaultRoute -h ipaddress</code> where <i>ipaddress</i> is the new default gateway IP address in dotted decimal format.
4.	Edit the chassis file.	Use the command: <code>opagenchassis &gt;&gt; /etc/opa/chassis</code> The chassis file contains the node name of managed switches corresponding to TCP/IP addresses as defined in the <code>/etc/hosts</code> file.
5.	Run the <code>opafastfabric</code> TUI.	Type <code>opafastfabric</code> and press <b>Enter</b> .
6.	Access the <b>FastFabric OPA Chassis Setup/Admin menu</b> .	Press <b>1</b> .
7.	Select menu items.	Select menu items 0–6 and press <b>P</b> to perform the operations.
8.	<b>Edit the Configuration and Select/Edit Chassis File</b> (menu item 0)	
continued...		



Step	Task/Prompt	Action
	a) Skip opafastfabric.conf.	No changes needed.
	b) Skip ports.	No changes needed.
	c) For chassis file, open the file in the editor.	Review the list of chassis selected. The setup of this file should have occurred above when setting up the Management Node by editing /etc/opa/chassis with the name corresponding to the Ethernet IP address of the chassis.
9.	<b>Verify Chassis via Ethernet Ping</b> (menu item 1)	Should pass without error.
10.	<b>Update Chassis Firmware</b> (menu item 2)	Specify the location for the firmware file to use.
11.	<b>Set Up Chassis Basic Configuration</b> (menu item 3)	Provide answers as follows: 1. Password: Press <b>Enter</b> (no password). 2. Syslog (y) a. Syslog server (n) b. TCP/UDP port number (n) - Use default. c. Syslog facility (n) - Use default. 3. NTP (n) - Customer to assign. 4. Timezone and DST (y) Use local timezone of server (y). 5. Do you wish to configure OPA Node Desc to match Ethernet chassis name? (y) - Enter y. 6. Do you wish to configure the Link CRC Mode? (n)
12.	<b>Set Up Password-less ssh/scp</b> (menu item 4)	Set up as required or press <b>Enter</b> to accept the default Chassis password.
13.	<b>Reboot the Chassis</b> (menu item 5)	Should pass without error.
14.	<b>Get Basic Chassis Configuration</b> (menu item 6)	Expected Summary output at end is shown below. Note that count should match the number of Edge switches.  <pre> Edgeswitch1: Firmware Active       : 10.x.x.x.x Firmware Primary      : 10.x.x.x.x Syslog Configuration  : Syslog host set to: 0.0.0.0 port 514 facility 22 NTP clock Time Zone             : Current time zone offset is: -5 LinkWidth Support     : 4X Node Description      : switch1 Link CRC Mode         : 48b_or_14b_or_16b </pre>
15.	Review the results.	Use an editor to view the following files: <ul style="list-style-type: none"> <li>• /root/test.res</li> <li>• /root/test.log</li> </ul>
	<b>End Task</b>	



## 8.0 Configure Intel® Omni-Path Director Class Switch 100 Series

---

Most Intel® OP Director switches are supplied with two Management Modules (MMs) for redundancy. In addition, Intel® OP Director switches have the following additional features:

- The switch has two Ethernet ports (one for each MM) and requires two Ethernet cables.
- The switch requires three IP addresses: one for each MM and one for the chassis, which is bound to the MM that is currently Master.
- It is useful to understand all reboot modes: `reboot all|-s|-m [slot #]` and how that causes failover.
- Default IP addresses of the Management Modules are:  
Chassis IP address: 192.168.100.9  
Management Module M201: 192.168.100.10  
Management Module M202: 192.168.100.11


The chassis file, located in `/etc/opa/chassis`, contains the node name of Intel® OP Director switches corresponding to TCP/IP addresses as defined in the `/etc/hosts` file. The chassis IP address is configured using the procedure for configuring managed switches, as described in [Configure Managed Intel® Omni-Path Edge Switches](#).

Configure the MM IP addresses using a serial connection as described in the following procedure:

1. Ensure that the module is connected to a COM port on a serial terminal device through the USB port.
2. Get to a `[boot]` : prompt by following either step a or b:
  - a. If the management module is running and displays `->` prompt, type the following command at the console: `reboot now` and press **ENTER**.
  - b. If the management module is not running, power on the switch.
3. When the system displays `image1`, press the spacebar to interrupt the autoloading sequence before the counter expires (within 5 seconds).
4. At the prompt, enter the command: `moduleip <ip_address>`  
The module reboots itself within 5 seconds and comes back with the new IP assigned to it. This module becomes the slave and the other MM becomes the master.

Repeat these steps for the second management module.

For more information, refer to the *Intel® Omni-Path Fabric Switches Hardware Installation Guide* and *Intel® Omni-Path Fabric Switches Release Notes*.



## 9.0 Configure Externally-Managed Intel® Omni-Path Edge Switches

For a complete description of the install process, refer to the *Intel® Omni-Path Fabric Software Installation Guide*, Configure the Externally-Managed Switches section.

For more information about the managed switch, refer to the *Intel® Omni-Path Fabric Switches Hardware Installation Guide* and *Intel® Omni-Path Fabric Switches Release Notes*.

The 100SWE48QF Edge switches do not have an Ethernet\* interface. Setup of these switches is performed using FastFabric via in-band commands.

### Procedure

The following steps describe the preferred approach for configuring the externally-managed Edge Switch:


Step	Task/Prompt	Action
1.	Edit the switches file for externally-managed switches.	Use the command: <code>opagenswitches &gt;&gt; /etc/opa/switches</code> The switches file contains a list of all the externally-managed switches in the fabric.
	a) Edit the switches file to replace the default switch name with the actual name that corresponds to the GUID for each switch.	For example: Default: <code>0x00117501026a5683:0:0,OmniPth00117501ff6a5602,2</code> Edited: <code>0x00117501026a5683:0:0,opaextmanagededge1,2</code>
2.	Run the <code>opafastfabric</code> TUI.	Type <code>opafastfabric</code> and press <b>Enter</b> .
3.	Access the <b>Externally Managed Switch Setup/Admin menu</b> .	Press <b>2</b> .
4.	Select menu items.	Select menu items 0–9 and press <b>P</b> to perform the operations.
5.	<b>Edit the Configuration and Select/Edit Switch File</b> (menu item 0)	
	a) Skip <code>opafastfabric.conf</code> .	No changes needed.
	b) Skip <code>ports</code> .	No changes needed.
	c) For <code>switches</code> file, open the file in the editor.	Review the list of switches selected. The <code>switches</code> file specifies: <ul style="list-style-type: none"><li>switches by node GUID</li><li>(optional) <code>hfi:port</code></li><li>(optional) Node Description (nodename) to be assigned to the switch</li><li>(optional) distance value indicating the relative distance from the FastFabric node for each switch</li></ul>

*continued...*



Step	Task/Prompt	Action
		<p>The following snippet shows the switches file format and an example:</p> <pre>nodeguid:hfi:port,nodename,distance 0x00117501026a5683:0:0,opaextmanagededge1,2</pre>
6.	<b>Generate or Update Switch File</b> (menu item 1)	
	a) Regenerate the switch file.	Type <b>n</b> if the switch file was generated in step 1. Type <b>y</b> if this is the first time or additional externally-managed switches have been added or replaced.
	b) Update switch names.	Type <b>y</b> . Note that this step may take a few minutes.
7.	<b>Test for Switch Presence</b> (menu item 2)	Should pass without error.
8.	<b>Verify Switch Firmware</b> (menu item 3)	Should pass without error.
9.	<b>Update Switch Firmware</b> (menu item 4)	Specify the location for the FW file (.emfw) to use.
10.	<b>Set Up Switch Basic Configuration</b> (menu item 5)	<p>Set up switch basic configuration and set the node description as shown below:</p> <pre>Performing Switch Admin: Setup Switch basic configuration Executing: /usr/sbin/opa-switch-admin -L /etc/opa/switches configure Do you wish to configure the switch Link Width Options? [n]: Do you wish to configure the switch Node Description as it is set in the switches file? [n]: y Do you wish to configure the switch FM Enabled option? [n]: Do you wish to configure the switch Link CRC Mode? [n]: Executing configure Test Suite (configure) Fri Jan 15 11:11:12 EST 2016 ... Executing TEST SUITE configure CASE (configure). 0x00117501026a5683:0:0,OmniPth00117501ff6a5602.i2c .extmgd.switchconfigure) configure switch 0x00117501026a5683:0:0,OmniPth00117501ff6a5602 ... TEST SUITE configure CASE (configure). 0x00117501026a5683:0:0,OmniPth00117501ff6a5602.i2c .extmgd.switchconfigure) configure switch 0x00117501026a5683:0:0,OmniPth00117501ff6a5602 PASSED TEST SUITE configure: 1 Cases; 1 PASSED</pre>
11.	<b>Reboot Switch</b> (menu item 6)	
12.	<b>Report Switch Firmware &amp; Hardware Info</b> (menu item 7)	<p>Review results for redundant power and FAN status. Expected summary output at end should be similar to the following (count should match number of externally-managed Edge switches):</p> <pre>0x00117501026a5683:0:0,opaextmanagededge1: F/W ver:10.x.x.x.x H/W ver:003-01 H/W pt num:H89344-003- 01 Fan status:Normal/Normal/Normal/Normal/Normal/Normal PS1 Status:ONLINE PS2 Status: ONLINE Temperature status:LTC2974:33C/MAX_QSFP:40C/PRR_ASIC:40C</pre> <p>Any non-redundant or failed fans or power supplies found during this step are also reported in /root/punchlist.csv.</p>
13.	<b>Get Basic Switch configuration</b> (menu item 8)	<p>Expected summary output at end should be similar to the following (count should match number of externally-managed Edge switches):</p> <pre>Link Width           : 1,2,3,4 Link Speed           : 25Gb FM Enabled           : No Link CRC Mode        : None vCU                  : 0 External Loopback Allowed : Yes Node Description      : Edgeswitch1</pre>

continued...



Step	Task/Prompt	Action
14.	<b>Report Switch VPD Information</b> (menu item 9)	Save the <code>test.res</code> output for future reference.
15.	Review the results.	Use an editor to view the following files: <ul style="list-style-type: none"><li>• <code>/root/test.res</code></li><li>• <code>/root/test.log</code></li></ul>
	<b>End Task</b>	



## 10.0 Configure Host Setup

For more information about the installation process, refer to the *Intel® Omni-Path Fabric Software Installation Guide* and *Intel® Omni-Path Fabric Software Release Notes*.

### Procedure

The following steps provide a summary for configuring the Hosts:

Step	Task/Prompt	Action
1.	Identify switch names.	Make sure all hosts are booted. This is required to identify switch names. If hosts are not available, you can perform all configuration steps except setting the switch names.
2.	Run the <code>opafastfabric</code> TUI.	Type <code>opafastfabric</code> and press <b>Enter</b> .
3.	Access the <b>FastFabric OPA Host Setup</b> menu.	Press <b>3</b> .
4.	Select menu items.	Select items 0–8 and press <b>P</b> to perform the operations. Perform the following operations as needed: <ul style="list-style-type: none"> <li>• <b>Edit Configuration and Select/Edit Host File</b> (menu item 0)</li> <li>• <b>Verify Hosts Pingable</b> (menu item 1)</li> <li>• <b>Set Up Password-Less SSH/SCP</b> (menu item 2)</li> <li>• <b>Copy /etc/hosts to All Hosts</b> (menu item 3)</li> <li>• <b>Show uname -a for All Hosts</b> (menu item 4)</li> </ul>
5.	<b>Install/Upgrade OPA Software</b> (menu item 5)	This installs IntelOPA-Basic on all compute nodes defined in <code>/etc/opa/hosts</code> . Be sure to exclude management node(s) with IFS installed and the node where you are running <code>opafastfabric</code> . <ol style="list-style-type: none"> <li>1. Provide the path to <code>IntelOPA-Basic.DISTRO.VERSION.tgz</code> when prompted.</li> <li>2. Enter directory to get <code>IntelOPA-Basic.DISTRO.VERSION.tgz</code> from (or none) : <code>/root</code>.</li> </ol>
6.	<b>Configure IPoIB IP Address</b> (menu item 6)	This performs the IPoIB ping test.
7.	<b>Build Test Apps and Copy to Hosts</b> (menu item 7)	This Builds Test Apps and Copy to Hosts. <ol style="list-style-type: none"> <li>1. Choose an MPI when prompted: Please Select MPI Directory</li> <li>2. Select an MPI with <code>-hfi</code> extension, so it will build with PSM2. For example: <code>/usr/mpi/gcc/openmpi-x.x.x-hfi</code>.</li> <li>3. When prompted to build base sample applications, select <code>yes</code>.</li> </ol>
8.	<b>Reboot Hosts</b> (menu item 8)	
	<b>End Task</b>	

## 11.0 Verify Cable Map Topology

---

This section describes how to use the fabric `topology.xml` file created in [Defining Type in the Topology Spreadsheet](#) to verify that fabric topology (cabling) is consistent with the cable map.

The command `opareport -o verify* -T <topologyfilename>.xml` compares the live fabric interconnect against the topology file created based on the cable map. These commands test links, switches, and SM topology. If successful, the output reports a total of 0 Incorrect Links found, 0 Missing, 0 Unexpected, 0 Misconnected, 0 Duplicate, and 0 Different.

```
# opareport -o verifyfis -T <topologyfilename>.xml
# opareport -o verifyextlinks -T <topologyfilename>.xml
# opareport -o verifyall -T <topologyfilename>.xml
```

In most cases, links reported with errors are either due to incorrect cabling to the wrong port or the `topology.csv` file has incorrect source and port destinations.

Verify the physical interconnect against the cable map using `opaextractsellinks` as shown in the following examples:

- List all the links in the fabric: `opaextractsellinks`
- List all the links to a switch named `OmniPth00117501ffffffff`:  
`opaextractsellinks -F "node:OmniPth00117501ffffffff"`
- List all the connections to end-nodes: `opaextractsellinks -F "nodetype:FI"`
- List all the links on the second HFI's fabric of a multi-plane fabric:  
`opaextractsellinks -h 2`

After all topology issues have been resolved, copy the `topologyfile.xml` from the local working directory to `cat /etc/opa/topology.0\0.xml`.

Refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide* for more information about using `opareport` in general, and using `opareport` for Advanced Topology Verification.





## 12.0 Verify Server and Fabric

Validation of servers and the fabric is initiated from the Management Node using the FastFabric TUI using `opafastfabric` on hosts defined in `/etc/opa/allhosts`.

### Procedure

Perform the following steps to verify servers and fabric:


Step	Task/Prompt	Action
1.	Run the <code>opafastfabric</code> TUI.	Type <code>opafastfabric</code> and press <b>Enter</b> .
2.	Access the <b>Host Verification/Admin</b> menu.	Press <b>4</b> .
3.	Select menu items.	Select items 3, 4, 6, and 8 and press <b>P</b> to perform the operations.
4.	<b>Perform Single Host Verification</b> (menu item 3)	
	a) Would you like to specify tests to run? [n]:	Enter <b>y</b> .
	b) View Load on hosts prior to verification? [y]:	Enter <b>y</b> . This option checks CPU load by running <code>/usr/sbin/opacheckload -f /etc/opa/allhosts</code> .
5.	<b>Verify OPA Fabric Status and Topology</b> (menu item 4)	Accept the default for all prompts. <i>Note:</i> This option goes through a fabric error and topology verification. Edit <code>/root/linkanalysis.res</code> to view results.
6.	<b>Verify Hosts Ping via IPoIB</b> (menu item 6)	This option pings all IPoIB interfaces.
7.	<b>Check MPI Performance</b> (menu item 8)	Accept the default for all prompts. <i>Note:</i> This option tests Latency and Bandwidth deviation between all hosts. Edit <code>/root/test.log</code> to view results.
	<b>End Task</b>	

For more information, refer to the *Intel® Omni-Path Fabric Software Installation Guide*.

### NOTE

A punchlist file is generated during execution of the FastFabric TUI and CLI commands, which can be used to track issues identified by the Intel® OPA tools. The punchlist file is located in `$FF_RESULT_DIR /punchlist.csv`, typically `/root/punchlist.csv`.

Two additional files, `/root/test.res` and `/root/test.log`, are created during OPA test commands and are useful for tracking test failures and issues.



## 13.0 Apply Best Known Methods for Site Installation

This section contains commands useful for configuring and debugging issues during fabric installation.

### 13.1 Enabling Intel® Omni-Path Fabric Manager GUI

By default, the Intel® Omni-Path Fabric Suite Fabric Manager GUI is disabled after installation of the IFS software. To quickly enable for early debug, use the following steps. For complete details, refer to the *Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide*.

#### NOTE

This method bypasses the SSH key authorization and is not intended for end customer installs.

#### Procedure

Perform the following steps to enable the Fabric Manager GUI:

Step	Task/Prompt	Action
1.	Edit /etc/opa-fm/opafm.xml file on the Management Node.	Make the two changes shown in <b>bold</b> for SslSecurityEnabled and default FE startup:  <pre>&lt;SslSecurityEnabled&gt;<b>0</b>&lt;/SslSecurityEnabled&gt;  &lt;!-- Common FE (Fabric Executive) attributes --&gt; &lt;Fe&gt; &lt;!-- The FE is required by the Intel Omni-Path FM GUI. --&gt; &lt;!-- To enable the FE, configure the SslSecurity parameters in this file --&gt; &lt;!-- as desired. --&gt; &lt;!-- For Host FM then set Start to 1. --&gt; &lt;!-- For Embedded FM the Start parameter in this file is not used; --&gt; &lt;!-- enable the FE via the smConfig and smFmStart chassis CLI commands. --&gt; &lt;Start&gt;<b>1</b>&lt;/Start&gt; &lt;!-- default FE startup for all instances --&gt; &lt;!-- Overrides of the Common.Shared parameters if desired --&gt; &lt;!-- &lt;SyslogFacility&gt;Local6&lt;/SyslogFacility&gt; --&gt;</pre>
2.	Restart the Fabric Manager to enable the changes and start the FE process required by the Fabric Manager GUI.	Type: <code>systemctl restart opafm</code>
3.	Download and install the Fabric Manager GUI application to a Windows* PC or Linux* system.	Refer to the <i>Intel® Omni-Path Fabric Software Installation Guide</i> , Install Intel® Omni-Path Fabric Suite Fabric Manager GUI section.
4.	Start the Fabric Manager GUI application.	
5.	Open the <b>Configuration</b> tab.	Enter the hostname or IP address of the Management Node running the Fabric Manager in your system into the FE Connection.
6.	Uncheck the Secure tab.	
7.	Click <b>Apply</b> to run the connection test.	
8.	Click <b>Run</b> to start the Fabric Manager GUI application.	
	<b>End Task</b>	



## NOTE

The Fabric Manager GUI does not operate through network proxies. Network firewall access may also need to be disabled. For a quick go/no-go verification, complete the connection test in the configuration tab as previously described.

## 13.2 Review Server and Fabric Verification Test Results

During fabric validation, unexpected loads on Host CPUs may result in inconsistent performance results. As a debug step, isolate the issue using the following:

Use the OPA tool to verify CPU host load. By default, it captures the top ten most heavily loaded hosts.

```
# /usr/sbin/opacheckload -f /etc/opa/allhosts
```


After the high load hosts have been identified, the next step is to root cause the issues.

### Procedure

Perform the following steps to verify servers and fabric:

Step	Task/Prompt	Action
1.	Check for HFI PCIe width or speed issues.	<p>Are HFI cards operating in a degraded mode, narrow width, or less than PCIe Gen3? Use <code>lspci</code> or <code>opahfirev</code> to verify the PCIe operating speed and bus width:</p> <ul style="list-style-type: none"> <li><code>lspci</code>: <a href="#">Verifying HFI Speed and Bus Width Using lspci</a></li> <li><code>opahfirev</code>: <a href="#">Decoding the Physical Configuration of an HFI</a></li> </ul> <p>Possible sources for narrow PCIe width:</p> <ol style="list-style-type: none"> <li>Be aware that OPA does support different width PCIe cards, including dual HFI cards using two x8 slices of a x16 physical connector. <code>opahfirev</code> is very useful for detecting this configuration.</li> <li>HFI Card partial insertion into x16 slots. Initially this appears to be a narrow width issue but re-inserting the card often resolves the issue. This may occur after a server is shipped. This step has resolved most width issues.</li> <li>Server physical configuration: Many servers support different PCIe logical widths based on riser card configuration. The slot may be physically x16 but internally limited to x8. Check other servers of the same configuration in the fabric. Check the server configuration. <i>This is also a common issue.</i></li> <li>Swap the HFI to another server to determine if the problem follows the card or the server.</li> </ol>
2.	Identify the key CPU load processes.	Use the Linux* <code>top</code> command.

*continued...*



Step	Task/Prompt	Action
		<i>Note:</i> <code>opatop</code> may be useful for checking for loads that vary over time. Use the <code>r</code> (rev), <code>f</code> (forward), and <code>L</code> (live) options to look through PM snapshots of system activity. This is also helpful for monitoring application startup versus run time loads. The PM captures high resolution statistics, with very low system overhead, over periods up to two days. The tools that harvest the PM stats are <code>opatop</code> and the FM GUI.
3.	Check for high CPU percent processes.	Examples of some issues: <b>Screen savers</b> - when a Linux* GUI is enabled on hosts, the screen that runs when the user interface is idle may have a high CPU load. <b>Test applications</b> - look for MPI jobs or similar applications running in the background. This is a common issue particularly in a shared fabric bring-up environment. Use <code>kill -p process</code> to stop orphan applications or reboot the server to debug the issue.
4.	Review BIOS settings to isolate nodes with different or incorrect settings.	Refer to <a href="#">Configuring BIOS Settings</a> on page 15.
	<b>End Task</b>	

## 13.3 Debugging Intel® Omni-Path Physical Link Issues

After you have run the FastFabric tool suite and identified issues with links, it is useful to start root-causing the issues. This section focuses on Intel® Omni-Path Fabric physical links and not PCIe bus link issues.

OPA reporting tools are robust, but it can be confusing for new users to understand the difference between error counters and actual failures.

From an installation perspective, it is important to watch for physical issues with cabling, both copper and optical. In general, bend radius, cable insertion issues, and physical compression or damage to cables can result in transmission issues. OPA recovers from many issues transparently.

The following information can help you root-cause solid failures as well as marginal links. Most often the issue is resolved simply by re-installing a cable and verifying that it clicks into the connector socket on the HFI or switch.

- View the QSFP/cable details of a specific switch port using the command:

```
opasmaquery -o cableinfo -d 10 -l <lid> -m <switch portnumber>
```

- To debug a particular switch, a useful technique is to get a snapshot of it using the command:

```
opareport -o snapshot -F portguid:value
```

### 13.3.1 OPA Link Transition Flow

To debug link issues, it is helpful to understand the four key link states, starting from Offline and running properly in the final Active state.



---

**NOTE**

The Fabric Manager, `opafm`, must be running to transition physical links from the Init state to the Active state. If you subsequently stop the Fabric Manager when a link is in the Active state, the link remains active. You can safely make changes to the `opafm.xml` file for the Fabric Manager and restart the service without dropping active links. As of the 10.0.0.696 software release, by default, the `opafm` service is not configured for autostart after IFS FULL installation.

---

PortState:

- Offline: link down. QSFP not present or not visible to the HFI driver.
- Polling: physical link training in progress. At this point you do not know whether the other end of the QSFP is connected to a working OPA device.
- Init: Link training has completed, both sides are present. Typically waiting for the Fabric Manager to enable the link.
- Active: Normal operating state of a fully functional link.

### 13.3.2 Verify the Fabric Manager is Running

From the Management Node, run the following command to report all HFIs and Switches.

```
# opafabricinfo
```

If it fails, try the following steps:

- Check status of the Fabric Manager process using the command:

```
# systemctl status opafm
```

- Restart the Fabric Manager using the command:

```
# systemctl start opafm
```

### 13.3.3 Check the State of All Links in the System

The `opaextractsellinks` command generates a CSV output representing the entire link state of the fabric.

- To generate the CSV output, run the following:

```
# opaextractsellinks > link_status.csv
```

- For links with errors, run the `opaextracterror` command.

```
# opaextracterror > link_status.csv
```



### 13.3.4 Check the State of HFI Links from a Server

If you are debugging server link issues, the `opainfo` command may be useful for a single server view.

`opainfo` captures a variety of data useful for debugging server related link issues. Multiple OPA commands can be used to extract individual data elements, however, this command is unique in the combination of data it provides.

- PortState: see [OPA Link Transition Flow](#) on page 36.
- LinkWidth: a fully functional link should indicate Act:4 and En:4.
- QSFP: Physical cable information for the QSFP, in this case a 5M Optical (AOC) Finisar cable.
- Link Quality: Range = 0 - 5 where 5 is Excellent.

```
# opainfo
hfi1_0:1 PortGID:0xfe80000000000000:001175010165b19c
  PortState: Active
  LinkSpeed Act: 25Gb En: 25Gb
  LinkWidth Act: 4 En: 4
  LinkWidthDnGrd ActTx: 4 Rx: 4 En: 3,4
  LCRC Act: 14-bit En: 14-bit,16-bit,48-bit Mgmt: True
  LID: 0x00000001-0x00000001 SM LID: 0x00000002 SL: 0
  QSFP: PassiveCu, 1m FCI Electronics P/N 10131941-2010LF Rev 5
  Xmit Data: 22581581 MB Pkts: 5100825193
  Recv Data: 18725619 MB Pkts: 4024569756
  Link Quality: 5 (Excellent)
```

### 13.3.5 Link Width, Downgrades, and `opafm.xml`

By default, OPA links run in x4 link width mode. OPA has a highly robust link mechanism, as compared to InfiniBand\*, and it allows links to run in reduced widths with no data loss.

Three things to know:

1. By default, the `opafm.xml` configuration file requires links to start up in x4 link width mode. This is configurable separately for HFI and ISL links using the **WidthPolicy** parameter.
2. Link downgrade ranges are also configurable in the `opafm.xml` file, using the **MaxDroppedLanes** parameter.
3. Default configuration example - A link that successfully starts up in x4 width and subsequently downgrades to x3 width continues to operate. If the link is restarted, by a server reboot, for example, and attempts to run by less than x4 width, then the link is disabled by the Fabric Manager and does not enter the Active state.

The `opainfo` command for HFIs is useful for checking the link width and link downgrade configuration on servers.

For a system view of all links that are running in less than x4 width mode, use the command:

```
# opareport -o errors -o slowlinks
```



### 13.3.6 How to Check Fabric Connectivity

For large fabrics, follow the flow described in [Defining Type in the Topology Spreadsheet](#) on page 21.

### 13.3.7 Physical Links Stability Test Using opacabletest

Intel® Omni-Path Architecture uses a quality metric for reporting status (opainfo). The quality metric ranges from 5 (excellent) to 1 (poor). For a more quantitative metric, use `cabletest` to generate traffic from on the HFI and ISL links, and `opaextractperf` and `opaextracterrors` to harvest the data.

#### Procedure

Perform the following steps:

Step	Task/Prompt	Action
1.	Prior to running the cable tests, perform the following:	<i>Note:</i> Use <code>opatop</code> to monitor fabric and link utilization.
	a) Clear error counters using <code>opareport -o none -clearall</code> and check the error counters after the test.	Use: <code>/usr/sbin/opacabletest -A -n 3 -f '/etc/opa/allhosts' stop_fi stop_isl # opareport -o none -clearall</code>
	b) Check to make sure there are no errors in fabric.	Use: <code>opareport -o errors</code>
2.	Start and stop cable test on the management node either from the <code>opafastfabric</code> TUI or using CLI commands	
	a) Using <code>opafastfabric</code> TUI	<ol style="list-style-type: none"> <li>1. Type <code># opafastfabric</code> at the prompt to start the FastFabric TUI.</li> <li>2. Select menu item <b>4</b> for Host Verification/Admin</li> <li>3. Select menu item <b>a</b> for Start or Stop Bit Error Rate Cable Test</li> <li>4. Press <b>P</b> to perform the operations.</li> </ol>
	b) Using CLI commands	Manually, run each test for a reasonable time, typically 5 - 15 minutes.
	• Run the cable test for hosts.	<code>/usr/sbin/opacabletest -A -n 3 -f '/etc/opa/allhosts' start_fi</code>
	• Run the cable test for ISLs.	<code>/usr/sbin/opacabletest -A -n 3 -f '/etc/opa/allhosts' stop_fi start_isl</code>
	• Stop the cable test and collect performance counters and statistics in csv files for viewing.	Use the following commands: <pre># /usr/sbin/opacabletest -A -n 3 -f '/etc/opa/allhosts' stop_isl stop_fi # opaextractperf &gt; link_stability_perf.csv # opaextracterrors &gt; link_stability_counters.csv</pre>
3.	For large fabrics, check stability using a long run of <code>opacabletest</code> .	Typically, you should run this for 4-8 hours. <i>Note:</i> Short runs of 10-15 minutes are fine for initial validation.
	<b>End Task</b>	



### How to Interpret the Results

The `opaextracterrors` command is a misnomer, it captures interesting statistics for evaluating links, but most of the content is not indicative of failures. The OPA fabric has robust end-to-end recovery mechanisms that handle issues.

Suggest looking specifically at the following columns:

- `LinkWidthDnGradeTxActive` - expect to see x4 Width
- `LinkWidthDnGradeRxActive` - expect to see x4 Width
- `LinkQualityIndicator` - 5 is excellent, 4 is acceptable, 3 is marginal and clearly an issue.
- `LinkDowned` - when an HFI is reset, the link down count increases, so rebooting a server results in small increments. If you see a link with significantly higher counts than its reboot expectations, then take a look at the server `/var/log/messages` file to determine whether the server is rebooting or the link is re-initializing.

For the other error counters, run a column sort and look for high error counts (greater than 100x) versus other links and take a look at the link types. Optical links have higher retry rates. This is not typically an issue unless they far exceed their peers.

The output is useful for verifying that every link is being tested. Unusual fabric `opaextractperf` topologies may result in non-optimum cabletest results. One workaround is to separately run `isl` and `fi` (HFI) link tests, then look at the total error results.

### 13.3.8 How to Debug and Fix Physical Link Issues

Check the topology before and after each of the debug steps using:

```
# opareport -o verifyall -T test_topology.xml
```

If the original issue was marginal operation rather than a hard failure, then re-run `cabletest` and analyze the `opaextracterrors` results to verify whether the issues were resolved.

At this point, you have a list of links with issues. Intel recommends the following approach for physical link resolution:

1. Unplug and re-insert each end of a physical cable. Check that the cable actually clicks into place. It may be useful to do this step separately for each end of the cable. Re-run `opacabletest` and verify whether the issue has been resolved or not.

---

#### NOTE

This step has resolved more link issues in fabric installs than all others.

---

2. Swap the questionable cable with a known good cable to isolate whether it is an HFI/Switch issue or cable issue.
3. If step 2 worked, then install the questionable cable into another location and verify whether it works.
4. If the issue is corrected, then the issue may be a mechanical latching issue on the HFI/Switch connector.





5. If the original issue was marginal operation rather than a hard failure, then re-run `opacabletest` and analyze the `opaextracterrors` results to verify whether the issues were resolved.
6. Re-run the physical links stability test using `opacabletest`.

### 13.3.9 Link Debug CLI Commands

Task	CLI Command
Identify fabric errors.	<code>opareport -o errors</code>
Identify slow links (< x4 width).	<code>opareport -o slowlinks</code>
Obtain lid of the switch.	Use <code>opaextractlids</code> for the links.
Obtain the port and links.	Use <code>opaextractsellinks</code> for the links.
If a link is not coming up as Active, first bounce the link, then check the link state.	<code>opaportconfig -l &lt;lid&gt; -m &lt;port&gt; bounce</code>
Get detailed link info for all nodes connected to an Edge switch or leaf and their neighbor.	<code>opareport -M -m -A -s -o comps -d 20 -F lid:&lt;lid of edge switch&gt;:node</code>
Find links that are not plugged in or not seen by the interface. Find all links stuck in the Offline state.	<code>opareport -A -m -F portphysstate:offline -o comps -d 5</code>
Find all links stuck in the Polling state. <i>Note:</i> A link stuck in Polling may indicate that the other end of the cable is not inserted correctly. In this case, typically, one end is Polling and the other end is Offline.	<code>opareport -A -m -F portphysstate:polling -o comps -d 5</code>
Identify bad links.	<code>opaextractbadlinks</code>
Disable all bad links and append <code>/etc/opa/disabled.0:0.csv</code> with a list of all bad links disabled. <i>Note:</i> As a debug step, disabling bad links should be temporary.	<code>opaextractbadlinks   opadisableports</code>
Enable links previously disabled.	<code>cat /etc/opa/disabled.0:0.csv   opaenableports</code>
To bounce a link, simulate a cable pull and re-insert on a server. <i>Note:</i> It may take up to 60 seconds for the port to re-enter the active state.	<code>opaportconfig bounce</code>
Check status of local hfi ports.	<code>opainfo</code>
<code>opaportconfig</code> and <code>opaportinfo</code> are key commands for port debugging.	Run the commands with the <code>-help</code> option to see available parameters.
Disable a set of links by extracting them to a csv file using <code>opaextractsellinks</code> .	<p>In the following example, links are extracted to <code>linkstodisable.csv</code>.</p> <ol style="list-style-type: none"> <li>To disable a set of links, run: <code>opadisableports &lt; linkstodisable.csv</code> By default, all disabled links are appended to the file <code>/etc/opa/disabled\1\1.csv</code>.</li> <li>To enable the disabled ports, run: <code>opaenableports &lt; /etc/opa/disabled\1\1.csv</code></li> </ol>

continued...



Task	CLI Command
	<p>After enabling the ports, the file <code>/etc/opa/disabled\1\1.csv</code> purges the links that are enabled.</p> <p><i>Note:</i> For each listed link, the switch port closer to this node is disabled.</p> <p>3. Run <code>opaportinfo -l &lt;lid of switch&gt; -m &lt;port number&gt;</code>.</p> <p>4. Check the port state by running: <code>opaportinfo -l 3 -m 0x10</code></p>

#### NOTE

Be sure to exclude the SM node on the Edge switch you are on and run `disableports` from the `linkstodisable` file to prevent cutting off this node from the fabric.

## 13.4 Using opatop to View Bandwidth and Error Summary

Use the the Fabric Performance Monitor TUI, `opatop`, to look at the bandwidth and error summary of HFIs and switches. This section provides a high-level overview of `opatop`.

1. To start up the Fabric Performance Monitor TUI, type `opatop` at the prompt.
2. Select 1 for HFIs or 2 for SW.

Intel recommends that you select 2 to view SWs bandwidth and error summary.

#### NOTE

In this display, HFIs show as `Send/Rcv` and ISLs show as `Int`.

3. On the Group Information screen:
  - Select `P` for Group Performance/Bandwidth Utilization.
  - Select `S` for Group Statistics/Error Category Stats.
4. Use `u` to move to an upper level.

For additional details, see the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Monitoring Fabric Performance section.

## 13.5 Using the Beacon LED to Identify HFI and Switch Ports

The LED beaconing flash pattern can be turned ON/OFF with the `opaportconfig` command. This can be used to identify the HFI and switches/ports installed in racks that need attention.



- For HFI:

```
opaportconfig -l 0x001 ledoff
Disabling Led at LID 0x00000001 Port 0 via local port 1 (0x0011750101671ed9)

opaportconfig -l 0x001 ledon
Enabling LED at LID 0x00000001 Port 0 via local port 1 (0x0011750101671ed9)
```

- For Switch port:

```
opaportconfig -l 0x002 -m 40 ledon (where -m 40 is port number)
Enabling LED at LID 0x00000002 Port 40 via local port 1 (0x0011750101671ed9)

opaportconfig -l 0x002 -m 40 ledoff
Disabling Led at LID 0x00000002 Port 40 via local port 1 (0x0011750101671ed9)
```

## 13.6 Decoding the Physical Configuration of an HFI

The `opahfirev` command provides a quick snapshot of an Intel® Omni-Path Host Fabric Interface (HFI), providing both PCIe status and physical configuration state, complementary to the `opainfo` command.

```
# opahfirev
#####
phsmpriv07.ph.intel.com - HFI 0000:81:00.0
HFI: hfil_0
Board: ChipABI 3.0, ChipRev 7.17, SW Compat 3
SN: 0x0063be82
Location:Discrete Socket:1 PCISlot:00 NUMANode:1 HFI0
Bus: Speed 8GT/s, Width x16
GUID: 0011:7501:0163:be82
SiRev: B1 (11)
TMM: 10.4.0.0.146
#####
```

Note the field for Thermal Monitoring Module (TMM) firmware version is an optional micro-controller for thermal monitoring on vendor-specific HFI adapters using the SMBus. For more information on the `opatmmtool`, see the *Intel® Omni-Path Fabric Host Software User Guide*.

- Check the current TMM firmware version using: `opatmmtool -fwversion`.
- Check the TMM firmware version in the `hfil_smbus.fw` file using:

```
opatmmtool -f /lib/firmware/updates/hfil_smbus.fw fileversion
```

- If the `fwversion` is less than `fileversion`, then update the TMM firmware version using:

```
opatmmtool -f /lib/firmware/updates/hfil_smbus.fw update
```

- After the TMM is updated, restart the TMM using:

```
opatmmtool reboot
```



---

**NOTE**

Data traffic is not interrupted during a reboot of the TMM.

---

## 13.7 Programming and Verifying Option ROM EEPROM Device

This section describes how to program and verify the Option ROM EEPROM device on an Intel® Omni-Path Host Fabric Interface (HFI).

### Before you Begin

- If you have an Intel® Omni-Path Host Fabric Interface, continue with this process.
- If you have an HFI from another manufacturer, contact the manufacturer's support team for the supported file versions and other specific instructions.

### Overview

The Option ROM contains three files that are pre-installed on the HFI at the time of manufacture.

- HFI1 UEFI Option ROM: `HfiPcieGen3_x.x.x.x.x.efi`
- UEFI UNDI Loader: `HfiPcieGen3Loader_x.x.x.x.x.rom`
- HFI1 platform file: `hfil_platform.dat`

---

**NOTE**

The `hfil_platform.dat` file is unique to the HFI hardware and should not need to be updated. If you do need to update this file, contact Intel Customer Support.

---

The files and the tools required to program the Option ROM EEPROM are provided at <https://www.intel.com/content/www/us/en/design/products-and-solutions/networking-and-io/fabric-products/omni-path/downloads.html>, Latest Release Library, Drivers and Software, Intel® Omni-Path UEFI Firmware.

However, the `hfil_platform.dat` file is unique to the HFI hardware and is not provided on the Intel Resource & Design Center. For Intel® HFI cards, this file may be obtained by contacting Intel Customer Support. For non-Intel HFIs, contact your hardware supplier.

To find the file locations, enter the following Linux\* commands:

```
find / -name hfil_platform.dat
find / -name <UEFI UNDI file name>.rom
find / -name <HFI1 UEFI file name>.efi
```

### Single Rail (one HFI) Example

The following example shows how to program a single HFI with three partitions.

1. Enter the following commands, replacing `x.x.x.x.x` with the versions provided in the release:



In normal mode user can only update HfiPcieGen3Loader.x.x.rom and HfiPcieGen3\_x.x.x.efi using only one command.

```
# hfil_eprom -u /usr/share/opa/bios_images/*
```

Individual files can be only updated in service mode.

```
# hfil_eprom -S -w -c /lib/firmware/updates/hfil_platform.dat
# hfil_eprom -S -w -o /usr/share/opa/bios_images/
HfiPcieGen3Loader_x.x.x.x.x.rom
# hfil_eprom -S -w -b /usr/share/opa/bios_images/HfiPcieGen3_x.x.x.x.x.efi
```

## NOTE

Ensure you select the correct partitions for your files.

2. Verify programmed versions using the following command:

```
# hfil_eprom -V
```

3. Reboot the server for the firmware updates to take effect.

## Dual Rail (two HFIs) Example

The following example shows how to program two HFIs with three partitions for each HFI.

1. Obtain the device assignment using the following Linux\* command:

```
# lspci | grep HFI
05:00.0 Fabric controller: Intel Corporation Omni-Path HFI Silicon 100 Series
[discrete] (rev 11)
81:00.0 Fabric controller: Intel Corporation Omni-Path HFI Silicon 100 Series
[discrete] (rev 11)
```

2. Program the devices using the following commands, replacing x.x.x.x.x with the versions provided in the release:

## NOTE

Ensure you use the correct device assignment.

In normal mode user can only update HfiPcieGen3Loader.x.x.rom and HfiPcieGen3\_x.x.x.efi using only one command.

```
# hfil_eprom -d /sys/bus/pci/devices/0000:05:00.0/resource0 -u /usr/share/opa/
bios_images/*

hfil_eprom -d /sys/bus/pci/devices/0000:81:00.0/resource0 -u /usr/share/opa/
bios_images/*
```

Individual files can be only updated in service mode.

```
# hfil_eprom -d /sys/bus/pci/devices/0000:05:00.0/resource0 -S -w -c /lib/
firmware/updates/hfil_platform.dat
# hfil_eprom -d /sys/bus/pci/devices/0000:05:00.0/resource0 -S -w -o /usr/
share/opa/bios_images/HfiPcieGen3Loader_x.x.x.x.x.rom
```



```
# hfi1_eprom -d /sys/bus/pci/devices/0000:05:00.0/resource0 -S -w -b /usr/
share/opa/bios_images/HfiPcieGen3_x.x.x.x.x.efi

# hfi1_eprom -d /sys/bus/pci/devices/0000:81:00.0/resource0 -S -w -c /lib/
firmware/updates/hfi1_platform.dat
# hfi1_eprom -d /sys/bus/pci/devices/0000:81:00.0/resource0 -S -w -o /usr/
share/opa/bios_images/HfiPcieGen3Loader_x.x.x.x.x.rom
# hfi1_eprom -d /sys/bus/pci/devices/0000:81:00.0/resource0 -S -w -b /usr/
share/opa/bios_images/HfiPcieGen3_x.x.x.x.x.efi
```

3. Verify programmed versions using the following commands:

```
# hfi1_eprom -d /sys/bus/pci/devices/0000:05:00.0/resource0 -V
# hfi1_eprom -d /sys/bus/pci/devices/0000:81:00.0/resource0 -V
```

4. Reboot the server for the firmware updates to take effect.

## 13.8 Verifying Fabric Manager Sweep

By default, Fabric Manager sweeps every five minutes as defined in the `/etc/opa-fm/opa-fm.xml` file. Sweeps are triggered sooner if there are fabric changes such as hosts, switches, or links going up or down. Edit `/var/log/messages` and search for `CYCLE START`. Each cycle start has a complementary cycle end. Any links with errors are noted during this sweep cycle.

An example of a clean SM sweep follows:

```
Feb 16 16:12:08 hds1fnb8261 fm0_sm[3946]: PROGR[topology]: SM: topology_main: TT:
DISCOVERY CYCLE START - REASON: Scheduled sweep interval
Feb 16 16:12:08 hds1fnb8261 fm0_sm[3946]: PROGR[topology]: SM: topology_main:
DISCOVERY CYCLE END. 9 SWs, 131 HFIs, 131 end ports, 523 total ports, 1 SM(s),
1902 packets, 0 retries, 0.350 sec sweep
```

Compare the sweep results with `opafabricinfo` and the fabric topology.

## 13.9 Verifying PM Sweep Duration

To show the PM sweep duration, perform the following steps:


1. Open `opatop`, then select `i`.

```
opatop: Img:Tue Feb 16 01:54:43 2016, Hist Now:Tue Feb 16 09:53:26 2016
Image Info:
Sweep Start: Tue Feb 16 01:54:43 2016
Sweep Duration: 0.001 Seconds

Num SW-Ports:      3  HFI-Ports:      2
Num SWs:           1  Num Links:      2  Num SMs:           2

Num Fail Nodes:    0  Ports:          0  Unexpected Clear Ports: 0
Num Skip Nodes:    0  Ports:          0
```

2. Select `r` to traverse the previous sweep duration time from history files. By default, PM sweeps every ten seconds. The latest ten image files (100 sec) are stored in RAM and up to 24 hours of history is stored in `/var/usr/lib/opa-fm`.



## 13.10 Checking Credit Loop Operation

For details on credit loops, see the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* QoS Operation section.

To verify that a fabric does not have a credit loop issue, use:

```
# opareport -o validatecreditloops
```

The output should report similar to the following where no credit loops are detected:

```
Fabric summary: 135 devices, 126 HFIs, 9 switches,
504 connections, 16880 routing decisions,
15750 analyzed routes, 0 incomplete routes
Done Building Graphical Layout of All Routes
Routes are deadlock free (No credit loops detected)
```

## 13.11 Modifying the Fabric Manager Routing Algorithm

If long Fabric Manager (FM) sweep times are observed or FM sweeps do not finish when a large number of nodes are bounced, consider changing the FM routing algorithm to `fattree` from the default `shortestpath`. You can do this by updating the `/etc/opa-fm/opafm.xml` file as shown in the following example:

```
<!-- ***** Fabric Routing ***** -->
<!-- The following Routing Algorithms are supported -->
<!-- shortestpath - pick shortest path and balance lids on ISLs -->
<!-- dgshortestpath - A variation of shortestpath that uses the -->
<!-- RoutingOrder parameter to control the order in which -->
<!-- switch egress ports are assigned to LIDs being routed -->
<!-- through the fabric. This can provide a better balance -->
<!-- of traffic through fabrics with multiple types of end -->
<!-- nodes. -->
<!-- See the <DGShortestPathTopology> section, below, for -->
<!-- more information. -->
<!-- fattree - A variation of shortestpath with better balancing -->
<!-- and improved SM performance on fat tree-like fabrics. -->
<RoutingAlgorithm>fattree</RoutingAlgorithm>
```

## 14.0 Run Benchmark and Stress Tests

---

For details on the tests provided with the Intel® Omni-Path Software, refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, MPI Sample Applications.

For optimal performance when running benchmark tests, configure systems according to the *Intel® Omni-Path Fabric Performance Tuning User Guide*.

### 14.1 Running Bandwidth Test

From \$FF\_MPI\_APPS\_DIR run:

```
# ./run_bw3
```

This test uses hosts defined in the `mpi_hosts` file, however, only the first two hosts in the file are used.

### 14.2 Running Latency Test

From \$FF\_MPI\_APPS\_DIR run:

```
# ./run_lat3
```

This test uses hosts defined in the `mpi_hosts` file, however, only the first two hosts in the file are used.

### 14.3 Running MPI Deviation Test

From \$FF\_MPI\_APPS\_DIR run:

```
# ./run_deviation 20 20 50
```

This test uses hosts defined in the `mpi_hosts` file.

### 14.4 Running run\_mpi\_stress

The default traffic pattern is "all-to-all" for this test.

Refer to *Intel® Omni-Path Fabric Suite FastFabric User Guide* for detailed information.

The test is located in \$FF\_MPI\_APPS\_DIR.



**NOTE**

These steps assume that you have defined hosts in the `/usr/src/opa/mpi_apps/mpi_hosts` file.

Step	Task/Prompt	Action
1.	Clear error counters.	# opareport -o none --clearall
2.	Confirm no errors exist.	# opareport -o error
3.	Run <code>mpi_stress</code> test using a 60 minute duration.	# ./run_mpi_stress all -t 60
4.	Run <code>opatop</code> to monitor the link utilization during the test.	
5.	Check error counts after the test.	# opareport -o errors
6.	View the log file that is available for analysis in <code>\$FF_MPI_APPS_DIR/logs</code> .	The log filename format is <code>mpi_stress.date_time</code> .
7.	Extract the log file in CSV format for errors and performance.	# opaextracterror # opaextractperf
	<b>End Task</b>	

**Example**

The following example demonstrates this test being run on four nodes.

```
# ./run_mpi_stress all -t 60
Running MPI tests with 4 processes
logfile /usr/src/opa/mpi_apps/logs/mpi_stress.12Apr17150628
OpenMPI Detected, running with mpirun.
Running Mpi Stress ... Running Mpi Stress ...
Using hosts list: /usr/src/opa/mpi_apps/mpi_hosts
Hosts in run:
node1
node2
node3
node4

+ /usr/mpi/gcc/openmpi-1.10.4-hfi/bin/mpirun -np 4 -map-by node
--allow-run-as-root -machinefile /usr/src/opa/mpi_apps/mpi_hosts
-mca plm_rsh_no_tree_spawn 1
/usr/mpi/gcc/openmpi-1.10.4-hfi/tests/intel/mpi_stress -t 60
```



## 15.0 Take State Dump of a Switch

### NOTE

Taking a state dump is a disruptive process and requires reboot of the switch after the state dump is taken. A state dump should only be taken if required to debug an issue.

You can take a state dump of any switch in the fabric, using its LID.

### Prerequisites

- Find the LID of the switch whose state you want to dump by running the `opaextractlids|grep switch name` command.
- Contact Intel Customer Support to get the correct username and password for the `supportLogin` command.

### Procedure

The following example describes how to take a state dump of a switch.

Step	Task/Prompt	Action
1.	Log in to a managed switch.	The default username and password are admin and adminpass.
2.	Run the <code>supportLogin</code> command using the support username and password obtained in <b>Prerequisites</b> .	<pre>-&gt; supportLogin username: support password:</pre>
3.	Run <code>ismTakeStateDump</code> dependent on your type of switch shown below.	
	a) For a local managed switch	Run the <code>ismTakeStateDump</code> command.
	b) For a remote managed or externally-managed switch	Run the command below, where <code>&lt;lid&gt;</code> identifies the desired switch.  <pre>-&gt; ismTakeStateDump -lid &lt;lid&gt; Dumping state of the switch at lid 4 to /firmware/prr- LID0004.gz</pre>
4.	From the Management Node, SFTP to the managed switch used for running the state dump command to retrieve the log	<pre>sftp admin@&lt;managed switch&gt; with password adminpass. admin@10.228.222.20's password: Connected to 10.228.222.20. sftp&gt; dir admin      operator      prr-LID0004.gz  prr- LID0005.gz  prr-LID0015.gz get prr-LID0004.gz</pre>
5.	Reboot the switch on which the state dump was taken to clear the state dump.	For externally-managed switches, use FastFabric to reboot the switch.
	<b>End Task</b>	



## 16.0 Apply Best Known Methods for OPA Commands

OPA commands should be issued from the Management Node where the IFS Full package was installed.

### 16.1 Retrieving the Host Fabric Interface Temperature

To retrieve the HFI temperature, use the command:

```
# cat /sys/class/infiniband/hfi1_X/tempsense
```

where *X* represents the device number.

When you send the command, the information is acquired at that specific time. Do not be concerned with the file's date/time.

An example of the output and the definition for each group of numbers follows:

```
# cat /sys/class/infiniband/hfi1_0/tempsense
68.50 0.00 105.00 105.00 0 0 0
```

**Table 1. HFI Temperature Output Definitions**

Example Value	Definition
68.50	Actual temperature Temperature steps are 0.25 °C increments.
0.00	Low limit
105.00	Upper limit
105.00	Critical limit
0	Low limit flag 1 = flag is set.
0	Upper limit flag 1 = flag is set.
0	Critical limit flag 1 = flag is set.

### 16.2 Reading Error Counters

To use the default thresholds defined in the `/etc/opa/opamon.conf` file, use the command:

```
# opareport -o errors
```



To run against a different threshold file, for example `/etc/opa/opamon.si.conf`, use the command:

```
# opareport -o errors -c /etc/opa/filename.conf
```

## 16.3 Clearing Error Counters

To clear error counters, use the command:

```
# opareport -o none --clearall
```

## 16.4 Loading and Unloading Intel® Omni-Path Host HFI Driver

If the configuration is changed, you may need to reload the HFI driver.

- To unload the HFI driver, use the command:

```
# modprobe -r hfi1
```

- To load the HFI driver, use the command:

```
# modprobe hfi1
```

---

### NOTE

The HFI driver should not be reloaded on SM nodes due to unloading other required dependencies and restarting them. On all other nodes, you may have to restart the IPoIB interface with `ifup` after the HFI driver is reloaded.

---

## 16.5 Analyzing Links

To include the link quality of the local HFI port, use the command:

```
# opainfo
```

To include links with lower quality, use the command:

```
# opareport -o errors
```

To output the ports with a link quality less than or equal to *value*, use the command:

```
# opareport -o links -F linkqualLE:value
```

To output the ports with a link quality greater than or equal to *value*, use the command:

```
# opareport -o links -F linkqualGE:value
```



To output the ports with a link quality equal to *value*, use the command:

```
# opareport -o links -F linkqual:value
```

**Table 2. Link Quality Values and Description**

Link Quality Value	Description
5	Working at or above preferred link quality, no action needed.
3	Working on low end of acceptable link quality, recommend corrective action on next maintenance window.
2	Working below acceptable link quality, recommend timely corrective action.
1	Working far below acceptable link quality, recommend immediate corrective action.
0	Link down
<i>Note:</i> Corrective action entails diagnosing the hardware (links/cables and ports/devices). For example: Are the cables bad or improperly placed? Is the HFI/switch responsive? Does rebooting the device/sever fix the issue?	

## 16.6 Tracing a Route between Two Nodes

To trace a route between two nodes, use the command:

```
# opareport -o route -S nodepat:"hds1fnb6101 hfil_0" -D nodepat:"hds1fnb6103 hfil_0"
```

To trace using LID, use the command:

```
# opareport -o route -S lid:5 -D lid:8
```

## 16.7 Analyzing All Fabric ISLs Routing Balance

To analyze the routing balance of all fabric ISLs, use the command:

```
# opareport -o treepathusage
```

## 16.8 Dumping Switch ASIC Forwarding Tables

To display all switch unicast forwarding tables DLIDs and Egress ports, use the command:

```
# opareport -o linear
```

To display multicast groups and members, use the command:

```
# opareport -o mcast
```



## 16.9 Configuring Redundant Fabric Manager Priority

This section describes several configuration methods.

### 16.9.1 FM Priority from a Local or Remote Terminal

Perform the following steps:

1. Edit the `/etc/opa-fm/opafm.xml` file.
2. Select the `<Priority>0</Priority>` value and change 0 to the number you want (0 - 15).
3. Save the file.
4. Start or restart the Fabric Manager to load the new file, using the command:

```
# systemctl restart opafm
```

---

#### NOTE

If you set a Fabric Manager to a higher priority, it becomes the master Fabric Manager automatically. The sticky finger option is disabled by default.

---

### 16.9.2 FM Elevated Priority

Perform the following steps:

1. Edit the `/etc/opa-fm/opafm.xml` file.
2. Select the `<ElevatedPriority>0</ElevatedPriority>` value and change 0 to the number you want (0 - 15).
3. Save the file.
4. Start or restart the Fabric Manager to load the new file, using the command:

```
# systemctl restart opafm
```

### 16.9.3 Configuration Consistency for Priority/Elevated Priority

Priority and Elevated Priority are not part of the `opafm.xml` configuration consistency checksum calculation. This makes standby Fabric Managers with mismatched configuration inactive because they are not valid to take over as Master in case of failover.

Having different values for Priority and Elevated Priority settings for SM instances is allowed and failover works as documented per Priority/ElevatedPriority settings. In normal failover without elevated priority, if the original Master Fabric Manager goes down, the Standby Fabric Manager becomes Master. When the original Master comes back up, it again takes over as Master.

---

#### NOTE

In sticky failover, Elevated Priority is used and with sticky failover enabled, when the original Master comes back up, it does NOT take over.

---



#### **16.9.4 FM states from the Management Node**

Run the `opafabricinfo` command to view the new active master SM.



## **17.0 Perform Final Fabric Checks**

---

After addressing all issues, perform final fabric checks as described in [Verify Server and Fabric](#) on page 33.





## 18.0 Customizing MPI

---

This section provides information for customizing and rebuilding MPI Libraries and MPI Applications, as needed.

### 18.1 Building MPI Libraries

Intel® Omni-Path Fabric Suite (IFS) comes complete with prebuilt versions of OpenMPI and MVAPICH2 that are enabled to use PSM2. The source code for the provided MPI versions are included with IFS and can be used with the FastFabric TUI to rebuild the MPI library along with the selection of a few options, such as compiler (refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Managing the Host Configuration). The src rpms for the MPI libraries are installed under `/usr/src/opa/MPI/` along with a `do_build` script that FastFabric uses to prompt the user to rebuild the MPI library.

If you want more advanced customization of the MPI options, refer to the `do_openmpi_build` and `do_mvapich2_build` scripts. Each of these scripts provide build options to the respective source rpms allowing you to rebuild the MPI library along with some standard test programs and benchmarks. Also, you can use these scripts as samples of the options needed to rebuild OpenMPI and MVAPICH2. For further details, the OpenMPI and MVAPICH2 libraries are publically documented in <https://www.open-mpi.org/> and <http://mvapich.cse.ohio-state.edu/>, respectively.

Intel® MPI is shipped separately from IFS as part of Intel® Parallel Studio. For information on Intel® MPI, refer to <https://software.intel.com/en-us/parallel-studio-xe>. For more information on using Intel® MPI with Intel® OPA, see the *Intel® Omni-Path Fabric Host Software User Guide*, Using MPI.

### 18.2 Building and Running MPI Applications

IFS comes with the source code for a number of MPI benchmarks. You can use the FastFabric TUI to rebuild the MPI benchmarks and distribute them to the compute nodes in the cluster (refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Managing the Host Configuration and MPI Sample Applications). The source code for the benchmarks are provided in `/usr/src/opa/mpi_apps` together with a `makefile` for building them as well as some sample `run_*` scripts. You may use these benchmarks and scripts to verify basic fabric and cluster performance during deployment. As such, they directly invoke the appropriate `mpirun` command for the selected MPI library using a local file listing the hostname for each rank in the job.

Typical production clusters use a job scheduler to select nodes for a given job, launch the MPI application, and monitor for its completion. Documentation on job schedulers is beyond the scope of this document. Consult with your system integrator for further details on using the job scheduler installed on the cluster.