



Intel[®] Omni-Path Host Fabric Interface

Installation Guide

Rev. 8.0

April 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 product plans and roadmaps are subject to change without notice.

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 may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Copyright © 2015–2020, Intel Corporation. All rights reserved.



Revision History

Date	Revision	Description
April 2020	8.0	Updated Hardware Installation with multi-socket server recommendation.
January 2020	7.0	Updated links to Intel® Omni-Path documentation set in Intel® Omni-Path Documentation Library .
October 2019	6.0	Refreshed to include updated Preface information.
April 2017	5.0	<ul style="list-style-type: none"> Added a note regarding the use of the NVIDIA* GPUDirect* RDMA Linux*kernel to Hardware Installation for HFI Without a PCI Express Riser Added a Note with a pointer to upgrading HFI firmware to Hardware Installation Checklist
December 2016	4.0	<ul style="list-style-type: none"> Added the section, Cluster Configurator for Intel® Omni-Path Fabric Added Figure 3, "Intel® Omni-Path Split PCI-E Dual Port Host Fabric Interface"
May 2016	3.0	Document updated for Revision 3.0
February 2016	2.0	Document updated for Revision 2.0
November 2015	1.0	Document updated for Revision 1.0




Contents

Revision History	3
Preface	7
Intended Audience	7
Intel® Omni-Path Documentation Library	7
How to Search the Intel® Omni-Path Documentation Set	9
Cluster Configurator for Intel® Omni-Path Fabric	10
Documentation Conventions	10
Best Practices	11
Laser Safety Information	11
Electrostatic Discharge Sensitivity (ESDS) Precautions	12
License Agreements	12
Technical Support	12
1.0 Introduction	13
1.1 Document Audience	13
1.2 How this Guide is Organized	13
1.3 Overview	13
2.0 Hardware Installation Checklist	15
3.0 Hardware Installation	16
3.1 Hardware Installation Requirements	16
3.1.1 Hardware	16
3.1.2 Cabling and Switches	16
3.2 Safety with Electricity	17
3.3 Unpacking Information	17
3.3.1 Verify the Package Contents	17
3.3.2 List of the Package Contents	17
3.3.3 Unpacking the HFI	18
3.4 Hardware Installation	19
3.4.1 Hardware Installation for HFI with PCI Express Riser	19
3.4.2 Hardware Installation for HFI Without a PCI Express Riser	21
3.5 Cabling the HFI to the Switch	22
3.6 Completing the Installation	22
3.7 LED Link and Data Indicators	23



Figures

1	Intel® Omni-Path Host Fabric Interface x8.....	18
2	Intel® Omni-Path Host Fabric Interface x16.....	18
3	PCIe Slot in a Typical Motherboard.....	20
4	Assembled PCIe Host Fabric Interface with Riser.....	20



Tables

1	HFI Model and Related Platforms.....	16
---	--------------------------------------	----



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.
<i>continued...</i>		



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.
continued...		



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.

- 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.
- In order to troubleshoot potential issues with externally-managed edge switches, Intel recommends that all fabrics contain at least one managed switch.
- 10.8.2.0 firmware introduced a feature to improve the detection of failing cables while also providing avoidance of possible downstream issues. This avoidance is accomplished by disabling cable lanes that are failing (i.e., putting the cable in a degraded mode). Equivalent protection can be achieved by enabling the 48-bit CRC feature which limits peak bandwidth of the cable by up to 3% while providing the strongest available protection against a failing cable. Intel recommends the use of 48-bit CRC when the peak bandwidth is not critical in the fabric.

Due to the improved detection of failing cables provided by 10.8.2.0 firmware, it is possible that failing cables that were not previously detected by earlier firmware will now be detected by the 10.8.2.0 firmware. It is expected that the 48-bit CRC and the 10.8.2.0 firmware detection methods will have similar detection robustness. When a failing cable is identified, it is important to replace the cable at the earliest possible convenience.

Laser Safety Information

This product may use Class 1 laser optical transceivers to communicate over the fiber optic conductors. The U.S. Department of Health and Human Services (DHHS) does not consider Class 1 lasers to be hazardous. The International Electrotechnical Commission (IEC) 825 Laser Safety Standard requires labeling in English, German, Finnish, and French stating that the product uses Class 1 lasers. Because it is impractical to label the transceivers, the following label is provided in this manual.



CLASS 1 LASER PRODUCT
LASER KLASSE 1
LUOKAN 1 LASERLAITE
APPAREIL A LASER DE CLASSE 1
TO IEC 825 (1984) + CENELEC HD 482 S1



Electrostatic Discharge Sensitivity (ESDS) Precautions

The assemblies used in the switch chassis are ESD sensitive. Observe ESD handling procedures when handling any assembly used in the switch chassis.

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 chapter describes the contents, intended audience, and organization of the *Intel® Omni-Path Host Fabric Interface Installation Guide*.

The *Intel® Omni-Path Host Fabric Interface Installation Guide* contains instructions for installing the HFI.

NOTE

This device has not been authorized as required by the rules of the Federal Communications Commission. This device is not, and may not be, offered for sale or lease, or sold or leased, until authorization is obtained.

1.1 Document Audience

This installation guide is intended for administrators responsible for installing the Intel® Omni-Path Host Fabric Interface in their Linux* cluster. Additional detailed information and instructions for administering an Intel® cluster can be found in the *Intel® Omni-Path Fabric Host Software User Guide*.

The *Intel® Omni-Path Host Fabric Interface Installation Guide* assumes that you are familiar with the specific hardware that you plan to use. Before installing the HFI, you should have basic knowledge of your host and target operating systems.

This document does not contain all the information you need to use basic Linux* commands or to perform all system administration tasks. For this information, see the software documentation you received with your system.

1.2 How this Guide is Organized

The *Intel® Omni-Path Host Fabric Interface Installation Guide* is organized into these sections:

- [Introduction](#) on page 13 contains an overview of the HFI and software, describes interoperability with other products, lists all related documentation, and provides Intel contact information.
- [Hardware Installation Checklist](#) on page 15 provides a high-level overview of the hardware installation procedures.
- [Hardware Installation](#) on page 16 includes instructions for installing the Intel® OP HFI.


1.3 Overview

The material in this documentation pertains to an OFED *cluster*. A cluster is defined as a collection of nodes, each attached to a fabric through the Intel® interconnect. The nodes are generally Linux*-based computers having multiple processors/cores. The Intel® OP HFI utilizes Intel® Omni-Path switches and cabling.



A cluster requires fabric management. Here are several options:

- Use the Intel® Omni-Path Embedded Fabric Manager in one or more managed Intel® Omni-Path switches.
- Use a host-based subnet manager such as the Intel® Omni-Path Fabric Suite Fabric Manager, as a part of the Intel® Omni-Path Fabric Suite download.



2.0 Hardware Installation Checklist

This section provides an overview of the hardware installation procedures. Detailed steps are found in [Hardware Installation](#) on page 16.

The following steps summarize the basic hardware installation procedure:

1. Check that the HFI hardware is appropriate for your platform. See [Table 1](#) on page 16.
2. Check to see that you have the appropriate cables and switches, as described in [Cabling and Switches](#) on page 16.
3. Following the safety instructions in [Safety with Electricity](#) on page 17. Unpack the HFI ([Unpacking Information](#) on page 17) and verify the package contents.
4. Install the HFI by following the instructions in [Hardware Installation](#) on page 19.
5. Cable the HFI to the switch, as described in [Cabling the HFI to the Switch](#) on page 22. Check that all switches are configured.
6. Follow the steps in [Completing the Installation](#) on page 22 to finish the installation.

NOTE

For information on upgrading HFI firmware, please refer to the *Boot Server Upgrade* section of the *Intel® Omni-Path Fabric Software Installation Guide*.



3.0 Hardware Installation

This section lists the requirements and provides instructions for installing the Intel® OP HFI. The components are collectively referred to as the *HFI* and the *riser card* in the remainder of this document.

3.1 Hardware Installation Requirements

This section lists hardware and software environment requirements for installing the HFI.

3.1.1 Hardware

The Intel® Omni-Path Host Fabric Interface HFI is for use with UL listed computers. The following statement is true for the HFI:

"This device complies with part 15 of the FCC Rules. Operation is subject to the following two conditions: (1) This device may not cause harmful interference, and (2) this device must accept any interference received, including interference that may cause undesired operations."

The table below displays the HFI and supported motherboards.

Table 1. HFI Model and Related Platforms

Intel Model Number	Platform	Plugs Into
Intel® Omni-Path Host Fabric Interface	PCI Express systems	x8 PCI Express slot x16 PCI Express slot

If the installation of the HFI requires the use of a riser card. See [Figure 3](#) on page 20 for an illustration of a PCI Express (PCIe) slot in a typical motherboard.

The motherboard vendor is the optimal source for information about the layout and use of PCI Express-enabled expansion slots on supported motherboards.

3.1.1.1 Form Factors

The HFI is a standard PCI Express half-height, short-form factor, and can be used with either full-height or low-profile face plates.

3.1.2 Cabling and Switches

The HFI uses standard Quad Small Form-factor Pluggable (QSFP) copper cables. These cables can be passive copper or active copper. Contact Intel Sales or your authorized Reseller for a list of qualified cables.

For cabling instructions, see [Cabling the HFI to the Switch](#) on page 22.



3.2 Safety with Electricity

Observe these guidelines and safety precautions when working around computer hardware and electrical equipment:

- Locate the power source shutoff for the computer room or lab where you are working. This is where you will turn OFF the power in the event of an emergency or accident. Never assume that power has been disconnected for a circuit; always check first.
- Do not wear loose clothing. Fasten your tie or scarf, remove jewelry, and roll up your sleeves. Wear safety glasses when working under any conditions that might be hazardous to your eyes.
- Shut down and disconnect the system's power supply from AC service before you begin work, to insure that standby power is not active. Power off all attached devices such as monitors, printers, and external components. Note that many motherboards and power supplies maintain standby power at all times. Inserting or removing components while standby is active can damage them.
- Use normal precautions to prevent electrostatic discharge, which can damage integrated circuits.

3.3 Unpacking Information

This section provides instructions for safely unpacking and handling the Intel® Omni-Path Host Fabric Interface. To avoid damaging the HFI, always take normal precautions to avoid electrostatic discharge.

3.3.1 Verify the Package Contents

The HFI should arrive in good condition. Before unpacking, check for any obvious damage to the packaging. If you find any damage to the packaging or to the contents, please notify your reseller immediately.

3.3.2 List of the Package Contents

The package contents for the HFI are:

- Intel® Omni-Path Host Fabric Interface
- Additional short bracket
- Additional standard bracket
- Readme First

Standard PCIe risers can be used. Risers are typically supplied by your system or motherboard vendor.

Figure 1. Intel® Omni-Path Host Fabric Interface x8





Figure 2. Intel® Omni-Path Host Fabric Interface x16



3.3.3 Unpacking the HFI

Follow these steps when unpacking the HFI:

1. When unpacking, ground yourself before removing the Intel® OP HFI from the anti- static bag.
2. Grasping the Intel® OP HFI by its face plate, pull the HFI out of the anti-static bag.



Handle the HFI only by its edges or the face plate. Do not allow the HFI or any of its components to touch any metal parts.

3. After checking for visual damage, store the HFI in its anti-static bag until you are ready to install.

3.4 Hardware Installation

This section contains hardware installation instructions for the HFI. When possible on a multi-socket server, Intel recommends that you install the HFI on the first socket. This eliminates unnecessary inter-socket communication for many applications and workloads that schedule tasks using cores on the first CPU socket by default.

3.4.1 Hardware Installation for HFI with PCI Express Riser

Most installations will be in 1U and 2U chassis, which requires using a PCIe right angle riser card. This results in an installation of the HFI that is parallel to the motherboard.

A taller riser card can be used if necessary. The HFI can connect to any of the standard compatible PCI Express riser cards.

3.4.1.1 Installation Steps


NOTE

Change the server BIOS PCIe setting to **AUTO**. The Intel® OP HFI card will train itself to PCIe **GEN 3**. For details on changing server BIOS, refer your server vendor documentation.

To install the Intel® OP HFI with a PCIe riser card:


1. Shut down the power supply to the system into which you will install the Intel® OP HFI.
2. Take precautions to avoid electrostatic damage (ESD) to the cards by properly grounding yourself or touching the metal chassis to discharge static electricity before handling the cards.
3. Remove the cover screws and cover plate to expose the system's motherboard. For specific instructions on how to do this, follow the hardware documentation that came with your system.
4. Locate the PCIe slot on your motherboard. Note that the PCIe slot has two separate sections, with the smaller slot opening located towards the front (see the figure below). These two sections correspond to the shorter and longer connector edges of the HFI and riser.

Figure 3. PCIe Slot in a Typical Motherboard



5. Determine if a blanking panel is installed in your chassis. If it is, remove it so that the InfiniBand* connector will be accessible. Refer to your system vendor instructions for how to remove the blanking panel.
6. Remove the Intel® OP HFI from the anti-static bag.
7. Locate the face plate on the connector edge of the card.
8. Connect the Intel® OP HFI and PCIe riser card together, forming the assembly that you will insert into your motherboard. First, visually line up the HFI slot connector edge with the edge connector of the PCIe riser card (see the figure below).

Figure 4. Assembled PCIe Host Fabric Interface with Riser



9. Holding the Intel® OP HFI by its edges, carefully insert the card slot connector into the PCIe riser card edge connector, as show in the figure above. The result is a combined L- shaped assembly of the PCIe riser card and Intel® OP HFI. This assembly is what you will insert into the PCIe slot on the motherboard in the next step.



10. Turn the assembly so that the riser card connector edge is facing the PCIe slot on the motherboard, and the face plate is toward the front of the chassis.
11. Holding this assembly above the motherboard at about a 45 degree angle, slowly lower it so that the connector on the face plate clears the blanking panel opening of the chassis from the inside. Slowly align the connector edge of the riser card with the PCIe slot of the motherboard. The short section of the connector must align with the short section of the slot.
12. Insert the riser assembly into the PCIe slot, ensuring good contact. The Intel® OP HFI should now be parallel to the motherboard and about one inch above it (see [Figure 4](#)).
13. Secure the face plate to the chassis. The Intel® OP HFI has a screw hole on the side of the face plate that can be attached to the chassis with a retention screw. The securing method may vary depending on the chassis manufacturer. Refer to the system documentation for information about mounting details such as mounting holes, screws to secure the card, or other brackets.

The Intel® OP HFI with PCIe riser card is now installed. Next, install the cables as described in [Cabling the HFI to the Switch](#) on page 22. Then test your installation by powering up and verifying link status (see [Completing the Installation](#) on page 22).

3.4.2 Hardware Installation for HFI Without a PCI Express Riser

NOTE

If you intend to use the NVIDIA* GPUDirect* RDMA Linux* kernel feature provided within the Intel® Omni-Path Fabric Software stack, you need to choose a PCIe card on *the same root complex* that you intend to install (or have installed) an NVIDIA* TESLA* card.

GPUDirect* RDMA is a technology that enables a direct path for data exchange between a graphics processing unit (GPU) and a third-party peer device using standard features of PCI Express. For more information, see the NVIDIA* CUDA* toolkit documentation:

<http://docs.nvidia.com/cuda/gpudirect-rdma/index.html>

Installing the Intel® OP HFI without a PCI Express riser card requires a 3U or larger chassis.

To install the Intel® OP HFI without a riser card:

1. Shut down the power supply to the system where you will install the Intel® OP HFI.
2. Take precautions to avoid electrostatic discharge (ESD) damage to the cards by properly grounding yourself or touching the metal chassis to discharge static electricity before handling the cards.
3. If you are installing the Intel® OP HFI into a covered system, remove the cover screws and cover plate to expose the system's motherboard. For specific instructions on how to do this, follow the hardware documentation that came with your system.
4. Locate the PCIe slot on your motherboard.



5. Remove the Intel® OP HFI from the anti-static bag. Hold the card by the top horizontal section of the bracket, and the top rear corner of the card. Be careful not to touch any of the components on the printed circuit card.
6. Without fully inserting, gently align and rest the gold fingers of the HFI on top of the motherboard's PCIe slot.
7. Insert the card by pressing firmly and evenly on the top of the horizontal bracket and the top rear corner of the card simultaneously. The card should insert evenly into the slot. Be careful not to push, grab, or put pressure on any other part of the card, and avoid touching any of the components.
8. Secure the face plate to the chassis. The Intel® OP HFI has a screw hole on the side of the face plate that can be attached to the chassis with a retention screw. The securing method may vary depending on the chassis manufacturer. Refer to the system documentation for information about mounting details such as mounting holes, and screws to secure the card or other brackets.

Next, install the cables, as described [Cabling the HFI to the Switch](#) on page 22. Then test your installation by powering up the system (see [Completing the Installation](#) on page 22).

3.5 Cabling the HFI to the Switch

Follow the recommendations of your cable vendor for cable management and proper bend radius.

To install the cable(s):

1. Check that you have removed the protector plugs(if any) from the cable connector ends.
2. Different vendor cables might have different latch mechanisms. Determine if your cable has a spring-loaded latch mechanism.
 - If your cable is spring-loaded, grasp the metal shell and pull on the plastic latch to release the cable. To insert, push and the cable snaps into place. You will hear a short "click" sound from the cable connector when it snaps in.
 - If your cable latch mechanism is not spring-loaded, push on the metal case, then push the plastic latch to lock the cable in place.
3. Connect the other end of the cable to a switch.

3.6 Completing the Installation

To complete the hardware installation:

1. Complete any other installation steps for other components.
2. Replace the cover plate and back panel.
3. Verify that the power cable is properly connected.
4. Turn on the power supply and boot the system normally.
5. Watch the LED indicator. The LEDs will flash only once, briefly, at power-up. The LED is functional only after the Intel® Omni-Path Fabric software has been installed, the driver has been loaded, and the system is connected to a switch.



NOTE

For information on upgrading HFI firmware, please refer to the *Boot Server Upgrade* section of the *Intel® Omni-Path Fabric Software Installation Guide*.

3.7 LED Link and Data Indicators

The LED functions are as follows:

- Link Down = Off
- Link Up Initialized = On solid green
- Link Up Active (No Traffic) = On solid green
- Link Up: Active (Slow Packet rate <10K/S) = BLINK: 384ms On, 384ms Off
- Link Up: Active (Fast Packet rate >10K/S) = BLINK: 128ms On, 128ms Off