Intel® Omni-Path Fabric Software

Release Notes for 10.4.2

June 2017

Contents

Tables

1.0 Overview of the Release

1.1 Introduction

This document provides a brief overview of the changes introduced into the Intel® Omni-Path Software by this release. References to more detailed information are provided where necessary. The information contained in this document is intended as supplemental information only; it should be used in conjunction with the documentation provided for each component.

These Release Notes list the features supported in this software release, open issues, and issues that were resolved during release development.

1.2 Audience

The information provided in this document is intended for installers, software support engineers, service personnel, and system administrators.

1.3 Software License Agreement

This software is provided under license agreements and may contain third-party software under separate third-party licensing. Please refer to the license files provided with the software for specific details.

1.4 If You Need Help

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.

1.5 Enhancements and Features in this Release

The following enhancements and features are new in the 10.4.2 release:

- Additional operating systems supported. See Table 1-2 for details.
- Additional hardware. See Table 1-5 for details.
- Changes to the installation path for all Intel® Omni-Path Software files. See Section 1.14.3 for details and recommended user action.
- Product Constraint described in Section 1.15.
- Support for GPUDirect* RDMA, which 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.
- Support for OpenFabrics Interfaces (OFI), a framework that includes libraries (including libfabric) and applications used to export fabric communication services to applications.

- Updated Accelerated RDMA support to include RHEL* 6.7. See Section 1.18 for details.

- Signed Kernel Modules, as required to support UEFI Secure Boot

- Technical preview of topology-aware job scheduling. See *Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide*, `opa2rm` tool for details.

- Support for NVMe over Fabric Protocol

- Support for IBM* Platform MPI and IBM* Spectrum MPI. See Table 1-3 for details.

- Virtual Fabric creation has been enhanced to better support advanced topologies, including the ability to place multicast traffic on a separate SL from unicast traffic. For details, see the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*, section 2.

1.6 Supported Features

- The list of supported operating systems is in Table 1-2.

- The list of supported hardware is in Table 1-5.

- Coexistence with Intel® True Scale Architecture. This release supports True Scale hardware serving as an InfiniBand* storage network with the Intel® Omni-Path hardware used for computing. Note that connecting a True Scale adapter card to an Omni-Path switch, or vice-versa, is not supported. For more details on this feature, refer to the *Intel® Omni-Path Fabric Host Software User Guide*.

- Supports multi-rail and multi-plane configurations. For more details, refer to the *Intel® Omni-Path Fabric Host Software User Guide*.

- Limited validation testing performed on network storage file systems:

  — NFS over TCP/IP

- Active Optical Cables (AOC). For details, see the Cable Matrix at:
  http://www.intel.com/content/www/us/en/high-performance-computing-fabrics/omni-path-cables.html

- MPI applications are provided in a stand-alone package.

- Intel® Xeon® v4 processor (codename Broadwell) support

- Intel® Xeon Phi™ support

- Monitored Intel® Omni-Path Host Fabric Interface

- DHCP and LDAP supported on Intel® Omni-Path Edge Switch 100 Series and Intel® Omni-Path Director Class Switch 100 Series hardware.

- Added support for Intel® Enterprise Edition for Lustre* software version 3.1.

- Support for the Enhanced Hypercube Routing Engine is outside the scope of Intel® OPA support. However, Intel partners may offer such support as part of their solutions. In addition there is an open source community who may be able to answer specific questions and provide guidance with respect to the Enhanced Hypercube Routing Engine.

- Support for Accelerated RDMA, also called Token ID (TID) RDMA, which is a Verbs protocol extension. See Section 1.18 for details.

- Support for SKX and SKX-F hardware.

- Supports RHEL* 6.7 and CentOS* 6.7.

- Support for active optical cables (AOC) on server platforms using integrated HFI for OPA (commonly known as "-F").
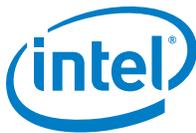
1.7 Release Packages

There are two Intel® Omni-Path Fabric Software packages:

- Basic for compute nodes
- IFS for the management node

The Basic package includes:

- Software that installs the following packages to the distribution OpenFabrics Alliance* (OFA):
    — libibumad is based on the RHEL* or SLES* release package. It contains Intel patches that support Intel® Omni-Path Architecture (Intel® OPA) technology.
    — ibacm is the latest upstream code applied with RHEL* patches.
    — hfi1-firmware, hfi1-psm, hfi1-diagtools-sw, libhfi1verbs
    — Open MPI built for verbs and PSM2 using gcc, and Intel compilers.
    — MVAPICH2 built for verbs and PSM2 using gcc, and Intel compilers.
    — mpitests
    — mpi-selector
    — GASnet
    — openSHMEM
    — srptools (includes the latest upstream code)
    — Firmware files listed in Table 1-1.
- compat-rdma which delivers kernel changes based on the OFA version. The components installed are the hfi1 driver and Intel-enhanced versions of other kernel packages. See the *Building Lustre* Servers with Intel® Omni-Path Architecture Application Note* for details.

    Note: In the Intel® Omni-Path Software package for RHEL* 7.2, the hfi1 driver and ifs-kernel-updates are supplied as a smaller package.

The IFS package includes the Basic package plus:

- Fabric Manager, which allows comprehensive control of administrative functions using a mature Subnet Manager. Fabric Manager simplifies subnet, fabric, and individual component management, easing the deployment and optimization of large fabrics.
- Fabric Suite FastFabric Toolset, which enables rapid, error-free installation and configuration of Intel® OPA host software and management software tools, as well as simplified installation, configuration, validation, and optimization of HPC fabrics. For details, refer to the Fabric Suite FastFabric documentation.

1.8 Firmware Files

This release of the Intel® Omni-Path Software contains the firmware files listed in Table 1-1. Intel provides UEFI firmware for discrete Intel® Omni-Path HFI cards and includes a platform file specific to Intel® Omni-Path HFI cards.

Table 1-1.   Firmware Files

Description	File Name	Version
HFI1 UEFI Option ROM	HfiPcieGen3_1.4.2.0.0.efi	1.4.2.0.0
UEFI UNDI Loader	HfiPcieGen3Loader_1.4.2.0.0.rom	1.4.2.0.0
HFI1 SMBus Microcontroller Firmware (Thermal Monitor)	hfi1_smbus.fw	10.4.0.0.146
Intel® Omni-Path HFI platform file ***Note:*** If you have a non-Intel HFI, contact the manufacturer's support team for details.	hfi1_platform.dat	HFI_TYPE1 v1.0.1.0

1.9 Operating Systems

This release of the Intel® Omni-Path Software supports the operating systems listed in Table 1-2.

Table 1-2.   Supported Operating Systems

Operating System	Update/ SP	Kernel Version
Red Hat* Enterprise Linux* (RHEL*) 6.7 X86_64	Update 7	2.6.32-573.el6.x86_64
CentOS* 6.7 X86_64	Update 7	2.6.32-573.el6.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.2 X86_64	Update 2	3.10.0-327.el7.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.3 X86_64	N/A	3.10.0-514.el7.x86_64
CentOS* 7.2 X86_64	N/A	3.10.0-327.el7.x86_64
CentOS* 7.3 X86_64	N/A	3.10.0-514.el7.x86_64
Scientific Linux* 7.2 X86_64	N/A	3.10.0-327.el7.x86_64
Scientific Linux* 7.3 X86_64	N/A	3.10.0-514.el7.x86_64
SUSE* Linux* Enterprise Server (SLES*) 12.1 X86_64	Service Pack 1	3.12.49-11.1-default
SUSE* Linux* Enterprise Server (SLES*) 12.2 X86_64	Service Pack 2	4.4.21-69-default

Note: The Intel® Xeon Phi™ Processor x200 product family (Knights Landing CPU-based servers) supports the following operating systems: RHEL* 7.3, CentOS* 7.3, and SLES* 12.2.

Note: PSM2 GPUDirect* RDMA is supported on RHEL* 7.2, RHEL* 7.3, and SLES* 12.2.

1.10 Parallel File Systems

The following parallel file systems have been tested with this release of the Intel® Omni-Path Software:

- Intel® Enterprise Edition for Lustre* software v3.1
    - — RHEL* versions supported by Intel® Omni-Path Software.
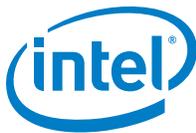
- IBM* Spectrum Scale version 4.2.1.0
  — RHEL* 7.2.

Refer to the *Intel® Omni-Path Fabric Performance Tuning User Guide* for details on optimizing parallel file system performance with Intel® Omni-Path Software.

1.11 MPI Libraries

1.11.1 Supported MPI Libraries

The table below lists the different MPI libraries supported by Intel® Omni-Path Fabric Software with the corresponding version, fabric support, and compiler used. Note that the second column indicates whether the MPI library is included in the Intel® Omni-Path Software package or not.

Table 1-3. Supported MPI Libraries

MPI Implementation	Included in IFS package?	Runs Over	Compiled With
Open MPI 1.10.4	Yes	Verbs	GCC
	Yes	PSM2	GCC, Intel
Open MPI 1.10.4-cuda	No	Verbs	N/A
	Yes	PSM2	GCC, Intel
MVAPICH2-2.2	Yes	Verbs	GCC
	Yes	PSM2	GCC, Intel
Intel® MPI 5.1.3 or later	No	Verbs	N/A
	No	PSM2	N/A
IBM* Platform* MPI version 9.1.4.3	No	Verbs	N/A
	No	PSM2	N/A
IBM* Spectrum* MPI version 10.1.0	No	Verbs	N/A
	No	PSM2	N/A

1.11.2 Compiler Versions and Distributions

The MPI libraries listed in the preceding table that are included in the release and built with PSM2 support were built with the following compiler versions:

Table 1-4. Compiler Versions and Distributions

Compiler	Linux* Distribution	Compiler Version
(GNU) gcc	RHEL* 7.2	gcc (GCC) 4.8.5 20150623 (Red Hat* 4.8.5-4)
(GNU) gcc	RHEL* 7.3	gcc (GCC) 4.8.5 20150623 (Red Hat* 4.8.5-11)
(GNU) gcc	SLES* 12 SP 1	gcc (SUSE* Linux*) version 4.8.5
(GNU) gcc	SLES* 12 SP 2	gcc (SUSE* Linux*) version 4.8.5
(Intel) icc	RHEL* 7.2	icc (ICC) 15.0.1
(Intel) icc	RHEL* 7.3	icc (ICC) 15.0.1
(Intel) icc	SLES* 12 SP 1	icc (ICC) 15.0.1
(Intel) icc	SLES* 12 SP 2	icc (ICC) 15.0.1

Note: Refer to the *Intel® Omni-Path Fabric Host Software User Guide* for set up information when using Open MPI with the SLURM PMI launcher and PSM2.

1.12 Hardware

Table 1-5 lists the hardware supported in this release.

Note: The Intel® PSM2 implementation has a limit of four (4) HFIs.

Table 1-5.   Supported Hardware

Hardware	Description
Intel® Xeon® Processor E5-2600 v3 product family	Haswell CPU-based servers
Intel® Xeon® Processor E5-2600 v4 product family	Broadwell CPU-based servers
Next generation Intel® Xeon® Processor (codename Skylake)	Skylake CPU-based servers (pre-production samples)
Intel® Xeon Phi™ Processor x200 product family	Knights Landing CPU-based servers
Intel® Omni-Path Host Fabric Interface 100HFA016 (x16)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Host Fabric Interface 100HFA018 (x8)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Switch 100SWE48Q	Managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48U	Externally-managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48UFH	Externally-managed 48-port Edge Switch, hot-swap power and fans
Intel® Omni-Path Switch 100SWE48QFH	Managed 48-port Edge Switch, hot-swap power and fans
Intel® Omni-Path Switch 100SWE24Q	Managed 24-port Edge Switch
Intel® Omni-Path Switch 100SWE24U	Externally-managed 24-port Edge Switch
Intel® Omni-Path Director Class Switch 100SWD24	Director Class Switch 100 Series, up to 768 ports
Intel® Omni-Path Director Class Switch 100SWD06	Director Class Switch 100 Series, up to 192 ports

Note: For RHEL* 6.7 and CentOS* 6.7, only the following processors are supported:
- Intel® Xeon® Processor E5-2600 v3 product family
- Intel® Xeon® Processor E5-2600 v4 product family
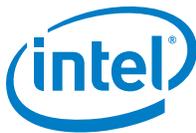
1.13 Documentation Versions

Table 1-6 lists the end-user document versions supported by this release.

Table 1-6. Supported Documentation Versions

Title	Doc. Number	Revision
Key: Shading indicates the URL to use for accessing the particular document.		
• Intel® Omni-Path Switches Installation, User, and Reference Guides: http://www.intel.com/omnipath/SwitchPublications		
• Intel® Omni-Path Software Installation, User, and Reference Guides (includes HFI documents): http://www.intel.com/omnipath/FabricSoftwarePublications		
• Drivers and Software (including Release Notes): http://www.intel.com/omnipath/Downloads		
Intel® Omni-Path Fabric Quick Start Guide	J57479	1.0
New title: *Intel® Omni-Path Fabric Setup Guide* (Old title: *Intel® Omni-Path Fabric Staging Guide*)	J27600	5.0
Intel® Omni-Path Fabric Switches Hardware Installation Guide	H76456	6.0
Intel® Omni-Path Host Fabric Interface Installation Guide	H76466	5.0
Intel® Omni-Path Fabric Software Installation Guide	H76467	6.0
Intel® Omni-Path Fabric Switches GUI User Guide	H76457	6.0
Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide	H76458	6.0
Intel® Omni-Path Fabric Suite FastFabric User Guide	H76469	6.0
Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide	H76472	6.0
Intel® Omni-Path Fabric Suite Fabric Manager User Guide	H76468	6.0
Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide	H76471	6.0
Intel® Omni-Path Fabric Host Software User Guide	H76470	6.0
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	6.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	8.0
Intel® Omni-Path IP and Storage Router Design Guide	H99668	5.0
Building Lustre Servers with Intel® Omni-Path Architecture Application Note*	J10040	1.0
Building Containers for Intel® Omni-Path Fabrics using Docker and Singularity* Application Note*	J57474	2.0
Intel® Omni-Path Fabric Software Release Notes	J66909	1.0
Intel® Omni-Path Fabric Fabric Manager GUI Release Notes	J60016	1.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)	J66908	1.0

For details on which document to use for a particular task, see Table 3-1 on page 24.

1.14 Installation Requirements

1.14.1 Software and Firmware Requirements

Table 1-2 lists the operating systems supported by this release. Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for the required packages.

1.14.2   Installation Instructions

There are two Intel® Omni-Path Fabric Software packages:

- `IntelOPA-IFS.<distro>-x86_64.<version>.tgz` for the management node.
- `IntelOPA-Basic.<distro>-x86_64.<version>.tgz` for compute nodes.

The packages in the tgz file are RPMs. Installing individual RPMs is not supported in the 10.4.2 release.

Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for related software requirements and complete installation procedures. Refer to the *Intel® Omni-Path Fabric Hardware Installation Guide* for related firmware requirements.

1.14.3   Installation Path Changes

If you are upgrading from a previous Intel® Omni-Path Fabric Software installation, Intel recommends that you remove certain RPMs prior to upgrading, due to changes in installation paths for RPMs and configuration files.

Intel recommends that you run the following command:

```
rpm -e opa-fm opa-fastfabric
```

If `opa-mpi-apps` is installed, run this command also:

```
rpm -e opa-mpi-apps
```

Pre-existing configuration files are automatically saved by the RPM as `.rpmsave` files. (RPM will notify you about these files during removal.) If you want to keep these configuration files, you should move them to their new locations. A mapping of old configuration file locations to new locations is shown in the following table.

Old Location	New Location (Release 10.4)
/etc/sysconfig/opafm.xml	/etc/opa-fm/opafm.xml
/etc/sysconfig/allhosts	/etc/opa/allhosts
/etc/sysconfig/chassis	/etc/opa/chassis
/etc/sysconfig/esm_chasis	/etc/opa/esm_chassis
/etc/sysconfig/hosts	/etc/opa/hosts
/etc/sysconfig/opafastfabric.conf	/etc/opa/opafastfabric.conf
/etc/sysconfig/opaff.xml	/etc/opa/opaff.xml
/etc/sysconfig/opamon.conf	/etc/opa/opamon.conf
/etc/sysconfig/ports	/etc/opa/ports
/etc/sysconfig/switches	/etc/opa/switches

1.14.4   Installation Prerequisites for RHEL* 6.7 and CentOS* 6.7

Install the following packages using yum from the RHEL* or CentOS* distributions:

- libibverbs
- librdmacm
- libibcm
- qperf
- perftest

- rdma
- infinipath-psm
- opensm-devel
- expat
- elfutils-libelf-devel
- libstdc++-devel
- gcc-gfortran
- atlas
- c-ares
- tcl
- expect
- tcsh
- sysfsutils
- pciutils
- bc (command line calculator for floating point math)
- rpm-build
- redhat-rpm-config
- kernel-devel
- opensm-libs

1.15   Product Constraint

Power class 2 and power class 3 Active Optical Cables (AOC) are not supported in this release.

1.16   Product Limitations

This release has the following product limitations:

- The embedded version of the Fabric Manager supports a maximum of 100 nodes within a fabric. This is due to the memory and processing resources available in the embedded environment.

- Performance Administration (PA) Failover should **not** be enabled with FMs running on differing software versions.
  PA Failover is enabled via
  `configuration:<PM>/<ImageUpdateInterval> > 0`

- Enabling UEFI Optimized Boot on some platforms can prevent the HFI UEFI driver from loading during boot. To prevent this, do not enable UEFI Optimized Boot.

1.17   RHEL* 6.7 and CentOS* 6.7 Support

1.17.1   RHEL* 6.7 and CentOS* 6.7 Limitations

RHEL* 6.7 and CentOS* 6.7 are supported in this release with the following limitations:

- Processor support:
  — Intel® Xeon® Processor E5-2600 v3 product family

— Intel® Xeon® Processor E5-2600 v4 product family

- File system support:

  — GPFS

  — NFS

  — Intel® Enterprise Edition for Lustre* software

  Note: For Intel® Enterprise Edition 3.0 Clients (support RHEL* 6.7) and Intel® Enterprise Edition 3.1 Servers (support RHEL* 7.3): You cannot upgrade your Clients beyond version 3.0 until you move to a newer RHEL* version.

- MVAPICH2 and Open MPI have been compiled for PSM2 to support the following versions of the compilers:

Compiler	Linux* Distribution	Compiler Version
(GNU) gcc	RHEL* 6.7 CentOS* 6.7	gcc (GCC) 4.4.7
(Intel) icc	RHEL* 6.7 CentOS* 6.7	icc (ICC) 15.0.1

- Performance is within 2%-5% of RHEL* 7.2 performance for the following features:

  — PSM bandwidth

  — MPI latency

  — Verbs bandwidth

1.17.2 RHEL* 6.7: Building Lustre* Kernel Modules for Intel® Omni-Path Support

The Intel® Enterprise Edition for Lustre* kernel modules are dependent on InfiniBand* core modules. In order for them to load properly in a system where the OPA stack is installed, they must have been compiled against `Module.symvers`. The IFS package installs `Module.symvers`, which is generated while building InfiniBand* core and Intel® Omni-Path kernel modules. This file contains symbol information that can be used to build higher-level kernel modules such as Lustre* that are dependent on InfiniBand* core or Intel® Omni-Path kernel modules.

It is installed in:
`/lib/modules/<kernel_ver>/include-ifs-kernel/Module.symvers`

Build Lustre* modules using the following command:
```
# KBUILD_EXTRA_SYMBOLS="/lib/modules/<kernel_ver>/
include-ifs-kernel/Module.symvers" rpmbuild –rebuild –without
servers <pathtomyrpm>
```

1.18 Accelerated RDMA Information

Accelerated RDMA is a Verbs protocol extension to improve the performance of RDMA write and RDMA read operations on Intel® Omni-Path hardware.

This extension improves the efficiency of large message transfers to provide performance benefits for storage protocols and other Verbs-based protocols. The benefits include increased achievable bandwidth with reduced CPU utilization. The Token ID (TID) RDMA protocol accelerates the OpenFabrics Alliance* (OFA) Verbs API with no changes required to API consumers. The acceleration technique is performed by the host driver and the application running over the OFA Verbs API does not need to make any code change.

Accelerated RDMA is off by default.

To enable it, add `cap_mask=0x4c09a01cbba` to the `/etc/modprobe.d/hfi1.conf` file. Instructions on how to do this are in the *Intel® Omni-Path Fabric Performance Tuning User Guide*, "Setting HFI1 Driver Parameters" section.

Note: Accelerated RDMA must be enabled on all nodes to function. Mixing of enabled and disabled nodes will not show performance benefits.
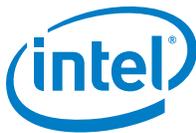
2.0   Issues

2.1   Introduction

This section lists the resolved and open issues in the Intel® Omni-Path Software.

2.2   Resolved Issues

Table 2-1 lists issues that are resolved in this release.

Table 2-1.   Issues resolved in this release (Sheet 1 of 2)

ID	Description	Resolved in Release
135068	Older versions of Grub 2 may not properly boot over Ethernet with the HFI UEFI driver.	10.4.2
138920	Reverted to a prior version of 8051 firmware to resolve a link bring up reliability regression issue while interoperating with older link firmware.	10.4.2
137708	Following a link bounce event, there is a possibility that a link will fail to reach the Armed/Active state. The likelihood of this issue depends largely on the link type: • Compute Nodes: These links are very unlikely to be affected. • FM Nodes: These links are the most exposed. If an FM link is affected and not recovered, there may be downstream effects over time.	10.4.1
130336	`hfilstats` cannot be run at user level due to mount-point privileges.	10.4
131017	Verbs ib_send_bw, ib_read_bw, and ib_write_bw are not working with the -R option to use the RDMA CM API to create QPs and exch data.	10.4
134268	The Option ROM image (e.g. containing a UEFI driver) may not be executed if the BIOS configures the HFI Expansion ROM BAR with an address that is not 16MB aligned.	10.4
134353	Very infrequently, when a link goes down, the logical link state can remain stuck in the 'Init' state.	10.4
134493	When using Mvapich2 with Intel® Omni-Path PSM2, users will notice unexpected behavior when seeding the built-in random number generator with functions like srand or srandom before MPI_Init is called. MPI_Init re-seeds the random number generator with its own value and does not restore the seed set by the user application. This causes different MPI ranks to generate different sequences of random numbers even though they started with the same seed value.	10.4
134821	The UEFI network stack is initialized with a default network address before the driver receives a MAD packet containing an updated and actual subnet prefix. Therefore, in ARP and IP UEFI drivers the old (default) HW address is still used, causing problems with packet receiving and transmitting.	10.4
135040	You can't currently specify portions of an Intel® DCS chassis that is not populated and is not expected to be populated. If CoreFull is 1, all the internal links for that chassis are generated when run against `opaxlattopology`. If CoreFull is 0, none of the links are generated.	10.4
135180	OpenMPI/PSM2 timeouts during MPI stress tests on Haswell and Intel® Xeon Phi™ mixed fabrics.	10.4
135326	Calling `opasmaquery` fails when called from a non-SM node to a node which has not booted to the OS.	10.4
135355	Due to changes in where the IFS packages are installed, customers using the FastFabric tools and upgrading to 10.3 from an earlier release must find each occurrence of `/opt/opa` in the `opafastfabric.conf` file and replace the string with `/usr/lib/opa`.	10.4
135545	A change has been made to several SA record attributes which causes incompatibilities between the Fabric tool suite and the SA.	10.4

Table 2-1.   Issues resolved in this release (Sheet 2 of 2)

ID	Description	Resolved in Release
135648	MPI applications are installed under the `/usr/lib` directory structure, which may be set up to be read-only overall. This causes resulting FastFabric operations to fail since mpi_apps contain source code and run scripts for sample MPI applications, test programs and benchmarks.	10.4
135711	After generating the `opafm.xml` file from the `config_generate` script, the FE is not enabled.	10.4
135873	hostverify.sh fails with RHEL* 6.7 due to the Intel P-State driver not being the default cpufreq driver.	10.4
136137	The hfi1_eprom tool man page contains incorrect information in the -d device option.	10.4
136733	Slow memory deregistration has been observed.	10.4
136902	A snapshot file with a multicast group with rate 10g will not be read properly. The following error is returned: opafabricanalysis: Port 0:0 Error: Unable to analyze fabric snapshot. See /var/usr/lib/opa/analysis/latest/fabric.0:0.links.stderr opafabricanalysis: Possible fabric errors or changes found	10.4
136945	When using the TID RDMA feature (Accelerated RDMA), certain Mvapich over Verbs tests may cause error messages.	10.4
136985	opahfirev has output errors when the HFI driver is not installed.	10.4
136995	The opahfirev tool output uses the term "HWRev" to indicate the revision of the silicon on the card.	10.4
137015	The state and configuration of ipoib interfaces are controlled by the NetworkManager service. The NetworkManager in RHEL* 7.2 mistakenly assumes the ipoib interface is type 'ethernet' and fails to initialize it, due to a mismatch against its actual type which is 'infiniband'.	10.4
137096	The IFS package does not install all the RPMs that it contains. In particular, infiniband-diags and libibmad are not automatically installed. The absence of infiniband-diags may result in failure of node descriptions to be populated, such that all hosts have the same `hfi1_0` description.	10.4
137108	When using the TID RDMA feature (Accelerated RDMA), virtual machines, and other cases where the IOMMU is enabled, do not operate correctly. This can lead to stability issues, and possibly data corruption, because the address used to receive data into will be incorrect.	10.4
137142	When using the TID RDMA feature (Accelerated RDMA), certain MPI benchmark tests may cause Kernel panic.	10.4
137221	Querying for switch info with opasmaquery while using the -g option will print incorrect IPv4 addresses.	10.4

Table 2-2 lists issues that are resolved in prior releases.

Table 2-2.   Issues resolved in prior releases (Sheet 1 of 2)

ID	Description	Resolved in Release
132219	Server platforms running IFS 10.3.0 release (or Intel® OPA software delivered in certain Linux* OS distributions) and using integrated HFI for OPA (commonly known as "-F") may not support Active Optical Cables (AOC) after boot up.	10.3.1
133377	irqbalance settings are not being honored correctly after a reboot.	10.3
133707	Updating to the RHEL* 7.2 kernel for the CVE-2016-0728 update in OSes prior to 7.2 causes the Intel® Omni-Path installation to fail.	10.3
134111	On some older HFI and HFI-like cards, running hfi1_eprom -V -c to inquire the version of the AOC configuration file on the card may return an invalid version of "etnIRFWI".	10.3
134124	HFI port stuck in INIT state due to SM failure to set pkeys.	10.3
134135 134429	When running communication-intensive workloads with 10KB MTU, it is possible to encounter node and/or job failures.	10.3

Table 2-2.   Issues resolved in prior releases (Sheet 2 of 2)

ID	Description	Resolved in Release

134283	When downgrading on a SLES* 12.X system from Intel® OPA version 10.2.X to a previous version, the following install errors occur:
	`ERROR – Failed to install`
	and
	`error: Failed dependencies: libibmad5 is needed by opa-basic-tools...`
		10.3

134772	`opatmmtool` will fail if provided with a filename (full path) that is longer than 63 characters.	10.3
134866	`hostverify.sh` cannot properly detect if SRP is enabled on target node.	10.3.1
134956	ib0 fails to become ready on warm reboots.	10.3.1
135000	Fabric Manager configuration files that specify IncludeGroup fields with undefined or nonexistent device groups could cause Fabric Manager failure.	10.3
135649	The XPPSL kernel changes conflict with items in the SLES* 12 SP1 kernel RPM. This causes the recompile of the SLES* 12 SP1 compat-rdma package to have an error.	10.3.1
135729 135870	KNL-F/SKL-F ports are offline in pre-boot setting when connected with AOC.	10.3.1
135812	FM may crash and restart in the event of a failure during topology assignments. This may result in mismatched port physical states on a link. While unlikely, this event may occur when there are integrity issues on a link.	10.3.1
135958	Spurious segmentation faults with greater than 2MB PSM2 transfers on Intel® Xeon Phi™ platforms.	10.3.1
136027	IFS `hostverify.sh` script does not provide reliable results for pstates_on and governor tests on RHEL* 7.3 and SLES* 12 SP2.	10.3.1
136028	Two versions of the UEFI firmware are contained in the hfi-uefi RPM in the 10.3.0 IFS and BASIC packages. The files are functionally identical except the unsigned files (`HfiPcieGen3Loader_<version number>.unsigned.rom` and `HfiPcieGen3_<version number>.unsigned.efi`) are not signed for secure boot.	10.3.1
136152	Server platforms using integrated HFI for OPA (commonly known as "-F") require BIOS that provides UEFI version 1.3.1.0.0 and a configuration data file for pre-boot support of Active Optical Cables (AOC). Some servers may not have these files available in BIOS and will therefore not support AOC in pre-boot.	10.3.1
136215	For RHEL6.7, the `opaconfig` command will not change the autostart settings for OPA service.	10.3.1
136318	SM crashes showing segfault errors in logs and high CPU usage. These crashes were caused by a mismatch of pahistory file versions.	10.3
136621	PCIe Fatal Errors during reboot cycles on server platforms using integrated HFI for OPA (commonly known as "-F").	10.3.1
136628	A bug in the Linux* kernel (CVE-2016-5195, also called Dirty COW) requires you to update the kernel for your operating system.	10.3.1
136723	Upgrading your OPA installation from version 10.2 to 10.3 may not install the correct host driver.	10.3.1

2.3   Open Issues

Table 2-3 lists the open issues for this release.

Table 2-3.   Open Issues (Sheet 1 of 6)

ID	Description	Workaround
129563	Memory allocation errors with Mvapich2-2.1/Verbs.	When running mvapich2 jobs with a large number of ranks (for example, between 36 and 72 ranks), you must set the following parameters in `/etc/security/limits.conf`: * hard memlock unlimited * soft memlock unlimited Also, you must increase the lkey_table_size:LKEY table size in bits (2^n, $1 <= n <= 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to Appendix A in the *Intel® Omni-Path Fabric Performance Tuning User Guide*.
131745	When running OpenMPI 1.10.0 on SLES* 12 with large number of ranks per node (over 40), it may happen that the ORTE daemon (orted) "hangs" during the finalization of job. This is an issue in Open MPI with the version of glibc used in SLES* 12. It is being researched by the Open MPI community in issue: https://github.com/open-mpi/ompi/issues/1136	Stopping and resuming the "hung" orted process allows the job to finish normally. To find the hung process, run the ps and find a node with several job zombie processes. In that same node, identify the orted process ID and send a stop signal (kill -19 <PID>) and a continue signal (kill -18 <PID>).
132207	Kernel crash caused by the ib_srpt module.	Install this kernel patch: https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=51093254bf879bc9ce96590400a87897c7498463
133596	When running the install script and installing all available packages, the installer may return a `prereq not installed` error.	Refer to the *Intel® Omni-Path Fabric Software Installation Guide*, in the OS RPMs Installation Prerequisites section for the complete list.
133604	Bonding driver shows incorrect hardware address of IPoIB interfaces.	Use the `opainfo` command to retrieve the PortGUID and `ip addr show ib0` to get the correct 20-byte hardware address of OPA network interface.
133633	OpenMPI and Mvapich2 compiles fail to link properly when using the Intel compilers.	No workaround available.
134409	In links exhibiting a high error rate, a rare PortRcvError is possible, resulting in a link down event. Such links should retrain and return to operation without user interaction. In cases where the Link Quality is less than or equal to 3, the interconnect in the link should be evaluated for possible replacement to prevent future PortRcvErrors from occurring.	Reboot or bounce the link.
134471	The HFI UEFI driver cannot boot via PXE using Grub 2.	Contact Intel Customer Support for assistance.
134494	Open MPI uses srand() family functions at MPI_Init() time. Therefore, if the user sets srand() before calling MPI_Init(), the values will be altered.	a) Fixed in Open MPI 2.0.1. b) Call srand() functions family after calling MPI_Init().
134819	In KNL-F EFI shell, the command `ifconfig -l` does not correctly display the IP address after being assigned via DHCP.	Launch a newer version of the EFI shell from the embedded shell.
134904	Legacy PXE boot using iPXE while the HFI UEFI driver is loaded may cause a hang.	Configure PXE operation to boot using UEFI boot mode.
135028	NVMe over Fabric Protocol is only supported on Intel® OPA with Linux* kernel 4.5 and later versions.	To use NVMe functionality on Intel® OPA, you must patch the kernel.
135084	In rare circumstances, the HFI may not appear in the PCI config space after a power cycle.	Reboot or power cycle the platform.

Table 2-3.   Open Issues (Sheet 2 of 6)

ID	Description	Workaround
135259	In links exhibiting a high error rate, a rare PortRcvError is possible, resulting in a link down event. Such links should retrain and return to operation without user interaction. In cases where the Link Quality is less than or equal to 3, the interconnect in the link should be evaluated for possible replacement to prevent future PortRcvErrors from occurring.	Reboot or bounce the link.
135360	On a system running RHEL* 7.2 or 7.3, if two `kmem_cache_creates()` occur with the same name, a kernel panic may result when the caches are deleted. The kernel panic is caused by `hfi1_user_sdma_free_queues`.	Install this patch: https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/commit/?id=26ea12dec0c84133add937455be76d44fe253d85

135390	Very old HFI adapters may be programmed with an obsolete version of the AOC platform configuration file. In these cases, errors such as the following may be observed:
	`[   26.903186] hfi1 0000:d5:00.0: hfi1_0:`
	`parse_platform_config:Bad config file`
	`[   26.903186] hfi1 0000:d5:00.0: hfi1_0:`
	`parse_platform_config:File claims to be`
	`larger than read size`
	`[   27.351555] hfi1 0000:d5:00.0: hfi1_0:`
	`tune_serdes: Unknown port type`
		Update the platform configuration file on the HFI to the current version. For details, see the *Intel® Omni-Path Fabric Software Installation Guide,* section B.1.

135929	Intel® Omni-Path Boot nodes occasionally dropped from fabric when switching master SM from one node to another.	Reboot PXE client node.
135963	Cannot install IFS software on RHEL* 7.3 using the command: `./INSTALL -vv -a`	Use the `-v` option instead.
135975	After performing an OPA software configuration update, some unmanaged switches do not update the settings for LinkWidth and LinkWidthDnGrade enables.	A reboot is required for configuration changes made to an externally managed switch to become active.
136049	The expected width of a card is not showing up correctly in `opaverifyhosts`.	For a cluster with mixed server or HFI configurations, the correct edited `hostverify.sh` script should be pushed to each group of servers. If using the TUI: • Create a /etc/opa/*hosts file for each type of server configuration. For example: `computehosts`, `storagehosts`, `mgmthosts`, etc. • Pick the desired hosts file in menu item 0 of the "Host Verification/Admin" menu, then run the "Perform Single Host Verification" function. • Edit the sample `hostverify.sh` script, putting in the proper settings for the server config (HFI PCIe bus, server memory size, expected single node HPL performance for server, etc). • When prompted, run the hostverify function on the given subset. • Repeat for each of the `*hosts` files.

Table 2-3.   Open Issues (Sheet 3 of 6)

ID	Description	Workaround
136160	On some Intel® Xeon Phi™ with integrated Intel® Omni-Path fabric platforms, the second integrated HFI is discovered first and is subsequently identified as the first HFI device. As a result, when issuing Intel® Omni-Path commands, the second HFI appears first in the results. In Linux* and various Intel® Omni-Path tools, the HFI reporting order may be the opposite of the order appearing on the Intel® Xeon Phi™ with integrated Intel® Omni-Path fabric cable/faceplate.	You can identify the second integrated HFI by inspecting the Node GUID or Port GUID/Port GID reported by `opainfo` or other commands such as `hfi1_control -i`. Note that bit 39 of the PortGUID, the most significant bit, is set for the second HFI, and is clear for the first HFI. Keep in mind that when issuing various Intel® Omni-Path CLI commands targeted at a specific HFI using the `-h` option, `-h 1` correlates to the device that is listed as `hfi1_0`. As a result, the issued command affects the second HFI instance in cases where the second HFI port instance appears first. By default, ports are ordered as enumerated by the kernel. There is a new module parameter called `port_reorder`. When set, the HFI1 ports on the same ASIC will be enumerated in increasing order. To enable this feature, use the command: `modprobe hfi1 port_reorder=1`
136419	When running SLES* 12.2 with inbox OPA drivers installed, the state may not change from "Offline" to "Physical Linkup (Init)" as expected.	Add a `platform.dat` file in `/lib/firmware/updates`, then restart.
136432	Certain perftest tools such as ib_write_bw do not work on RHEL* 7.3 when using the RDMA CM with UD QPs.	Roll back the perftest package to the level found in RHEL* 7.2, which is perftest-2.4. Then install this package on RHEL* 7.3.

Table 2-3.   Open Issues (Sheet 4 of 6)

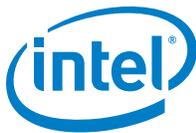| ID	Description	Workaround

136436	On SLES* 12.2, node_desc is not populated with the host name when system is booted up.	Install and run the rdma-ndd daemon on each node.
		
		1. Unpack IFS:
		`# tar xzf IntelOPA-IFS.SLES122-x86_64.10.3.0.0.81.tgz`
		`# ls`
		`IntelOPA-IFS.SLES122-x86_64.10.3.0.0.81`
		`IntelOPA-IFS.SLES122-x86_64.10.3.0.0.81.tgz`
		2. Uninstall infiniband-diags and libibnetdisc5 libraries. (SLES* splits out the libibnetdisc library but it is included in the IFS infiniband-diags version.)
		`# rpm -e infiniband-diags`
		`# rpm -e libibnetdisc5`
		3. Install the older version of infiniband-diags from the IFS package.
		`# cd IntelOPA-OFED_DELTA.SLES122-x86_64.10.3.0.0.82/`
		`# rpm -Uvh ./infiniband-diags-1.6.7-2.x86_64.rpm`
		4. Enable rdma-ndd:
		`# systemctl daemon-reload`
		`# systemctl status rdma-ndd`
		`  rdma-ndd.service - RDMA Node Description Daemon`
		`   Loaded: loaded (/usr/lib/systemd/system/rdma-ndd.service; disabled; vendor preset: disabled)`
		`   Active: inactive (dead)`
		`# systemctl enable rdma-ndd`
		`Created symlink from /etc/systemd/system/multi-user.target.wants/rdma-ndd.service to /usr/lib/systemd/system/rdma-ndd.service.`
		5. Start rdma-ndd and check the status:
		`# systemctl start rdma-ndd`
		`# systemctl status rdma-ndd`
		6. Test that it is working:
		`# cat /sys/class/infiniband/hfi1_0/node_desc`
		`phs1fnive08u26 hfi1_0`
		`# hostname foo`
		`# cat /sys/class/infiniband/hfi1_0/node_desc`
		`foo hfi1_0`
		`# hostname phs1fnive08u26`
		`# cat /sys/class/infiniband/hfi1_0/node_desc`
		`phs1fnive08u26 hfi1_0`
		`# reboot`
		`...`
		`# cat /sys/class/infiniband/hfi1_0/node_desc`
		`phs1fnive08u26 hfi1_0`

136437	When using RHEL* 7.2, the default generic PXE boot image does not work due to missing driver and firmware files.	Contact Intel Customer Support for assistance.
136500	RDMA perftests can hang on start on a client side when RDMA CM (-R option) is used.	Intel recommends that you use the same version of perftests across your fabric. Obtain the latest perftests version from the upstream repository.
136727	Initialization of PSM2 library fails with the following error message: Error: PSM is in the finalized state	The openmpi-mca-params.conf file is automatically edited during IFS install. If you have edited the file after installation, then you may see this error. To correct the error, edit the file `/usr/mpi/gcc/openmpi-1.10.4-hfi/etc/openmpi-mca-params.conf` and set the default parameters for mpirun to: `-mca pml cm -mca mtl psm2`

Table 2-3.   Open Issues (Sheet 5 of 6)

ID	Description	Workaround
136728	If hundreds of links are bouncing while the FM is sweeping, the FM sweep time may be significantly extended. This can result in unexpected delays in FM responsiveness to fabric changes or host reboots. (The issue is that active links bounce between the time FM discovers one side of the link versus the other side of the link.) In Release 10.3.1 a change was made to improve the FM responsiveness in large fabrics of >1000 nodes when numerous links bounce (or hosts are rebooted) at once.	The following workarounds are recommended: • When rebooting nodes on a production cluster, perform reboots in batches of 300 nodes or less. • During cluster deployment, carefully follow the procedures in the *Intel® Omni-Path Fabric Setup Guide* and use FastFabric to check signal integrity and placement of all cables. Correct or disable any problematic links before starting production use of the cluster. • When replacing or expanding a production cluster, repeat the procedures in the *Intel® Omni-Path Fabric Setup Guide* to verify the new hardware. Correct or disable any problematic links before resuming production use of the cluster. • Use the PM, FM logs, FM GUI, FastFabric, and other tools to monitor signal integrity and link stability. Correct or disable any problematic links when discovered.
136821	When performing boot over fabric, links may take up to 6 minutes to become active.	None.
136822	The Intel UEFI driver contained in the server BIOS must be executed for proper support of Active Optical Cables (AOC) in an integrated HFI environment. Some BIOS do not execute the UEFI in Legacy BIOS Boot mode, and there are BIOS configuration settings that may prevent the UEFI from executing in any mode.	Avoid the use of Legacy BIOS boot mode if your platform does not execute the HFI driver in that mode. Avoid BIOS settings or other configuration settings that do not execute the HFI driver during boot.
136901	Occasionally, nodes may be dropped by the Fabric Manager while they are in a pre-boot mode. This can occur when the node has multiple HFIs on a single socket.	Bounce the link of the dropped pre-boot port.
136971	When using the Accelerated RDMA feature (TID RDMA), certain Verbs Multi-PPN tests may cause error messages.	To avoid errors when running Verbs Multi-PPN tests, limit the parallel processes to 16 or fewer.
137054	Pinging an Intel® OPA UEFI permanent IP address from a DHCP server fails on subsequent reboots unless the corresponding network interface has first been initialized in the UEFI network stack.	Before pinging a UEFI permanent IP address, first initialize the corresponding network interface in the UEFI network stack.
137106	When running SLES* 12.2 with inbox OPA drivers installed, the state may not change from "Offline" to "Physical Linkup (Init)" as expected.	Add a `platform.dat` file in `/lib/firmware/updates`, then restart.
137212	The RHEL* 6.7 base version of the perftest package includes a ib_send_lat utility that may cause a segmentation fault when run with the -z option.	Run the utility without using the -z option.
137364	The node description of a node may change after rebooting. This issue has been seen on RHEL* and SLES*.	You must manually enable and start the rdma-ndd service on Intel® OPA software installations.
137372	Packets may be stuck in kernel when attempting writes to file system via IPoIB interface.	Underlying software should re-establish connection and retry sending the data.
137499	HFI links may occasionally take several minutes to reach link up.	None.
137577	`opatmmtool` does not provide a correct error message if it is run on a system that does not have a TMM.	Do not run `opatmmtool` on a system that does not have a TMM.
137616	When booting in legacy BIOS boot mode on RHEL* 7.x and SLES* 12.x, the following message is present in the kernel: Request for unknown module key 'Intel Corporation: Intel(R) Omni-Path HFI UEFI: 719ebaa125172ba69ad01b850b7458f85c89bb07' err -11	No workaround required. This does not prevent hfi1 driver functionality.

Table 2-3.   Open Issues (Sheet 6 of 6)

ID	Description	Workaround

137869	When using GPUDirect* RDMA, you must run .INSTALL using the -G option.
	In this release, the INSTALL script is not passing the -G argument to remote hosts.
		Copy the Basic tarball to remote hosts, untar, and launch INSTALL with -G manually on the remote host. This will install the correct GPUDirect* enabled driver and PSM.
		For clusters with many nodes, using `opafastfabric` to perform the installation of hosts that use GPUDirect* may be preferred.
		To do this, create two different FastFabric hosts files, one with hosts that use GPUDirect* (for example, `hosts_cuda`) and another with all the other hosts (for example, `hosts_other`).
		Run `opafastfabric` against the GPUDirect* hosts using the following command:
		`HOSTS_FILE=hosts_cuda FF_INSTALL_OPTIONS="-G" FF_UPGRADE_OPTIONS="-G" opafastfabric`
		When this is complete, exit `opafastfabric` and run `opafastfabric` against the other hosts using the following command:
		`HOSTS_FILE=hosts_other opafastfabric`
		
		Note: For other `opafastfabric` operations, the `HOSTS_FILE` selection can be made using the TUI menu option 0 as needed.

138171	During execution of an `opacapture` command, the core file may not be properly copied. This can occur on systems running the SLES* operating system when the following message appears during the capture: `/usr/lib/opa-fm/bin/fm_capture: line 148: [: too many arguments` If a core file exists in the default directory, that core file will be copied.	Replace line 148 in `/usrlib/opa-fm/bin/fm_capture` with the following: `if [ -n "$core_path" -a "${core_path:0:1}" = "/" ]`
138183	In this release, additional fields were added to the `opareport -o snapshot -r` XML output format that are not present in Release 10.3. Therefore, Release 10.3 snapshot files will report a "Mandatory Tag Not Found" parser error using Release 10.4 Fabric Manager tools.	Regenerate any such snapshot files using the `opareport` tool in Release 10.4.
138188	Coexistence feature with Intel® True Scale HCA cards does not function on servers running RHEL* 7.3 OS.	Customers requiring this feature must use RHEL* 7.2 OS.

3.0 Related Information

3.1 Documentation

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

- Intel® Omni-Path Switches Installation, User, and Reference Guides
  www.intel.com/omnipath/SwitchPublications

- Intel® Omni-Path Fabric Software Installation, User, and Reference Guides
  www.intel.com/omnipath/FabricSoftwarePublications

- Drivers and Software (including Release Notes)
  www.intel.com/omnipath/downloads

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

Table 3-1. Intel® Omni-Path Documentation Library (Sheet 1 of 3)

Task	Document Title	Description
Key: Shading indicates the URL to use for accessing the particular document.		
• Intel® Omni-Path Switches Installation, User, and Reference Guides: http://www.intel.com/omnipath/SwitchPublications		
• Intel® Omni-Path Software Installation, User, and Reference Guides (includes HFI documents): http://www.intel.com/omnipath/FabricSoftwarePublications		
• Drivers and Software (including Release Notes): http://www.intel.com/omnipath/Downloads		
Using the Intel® OPA documentation set	*Intel® Omni-Path Fabric Quick Start Guide*	A roadmap to Intel's comprehensive library of publications describing all aspects of the product family. It outlines the most basic steps for getting your Intel® Omni-Path Architecture (Intel® OPA) cluster installed and operational.
Setting up an Intel® OPA cluster	**New title:** *Intel® Omni-Path Fabric Setup Guide* (Old title: *Intel® Omni-Path Fabric Staging Guide*)	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 BKMs are defined to facilitate the installation process and troubleshooting.
Installing hardware	*Intel® Omni-Path Fabric Switches Hardware Installation Guide*	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.
	Intel® Omni-Path Host Fabric Interface Installation Guide	Contains instructions for installing the HFI in an Intel® OPA cluster. A cluster is defined as a collection of nodes, each attached to a fabric through the Intel interconnect. The Intel® HFI utilizes Intel® Omni-Path switches and cabling.

Table 3-1.   Intel® Omni-Path Documentation Library (Sheet 2 of 3)

Task	Document Title	Description
Installing host software Installing HFI firmware Installing switch firmware (externally-managed switches)	*Intel® Omni-Path Fabric Software Installation Guide*	Describes using a Text 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 rpms individually.
Managing a switch using Chassis Viewer GUI Installing switch firmware (managed switches)	*Intel® Omni-Path Fabric Switches GUI User Guide*	Describes the Intel® Omni-Path Fabric Chassis Viewer graphical user interface (GUI). It provides task-oriented procedures for configuring and managing the Intel® Omni-Path Switch family. Help: GUI Online Help.
Managing a switch using the CLI Installing switch firmware (managed switches)	*Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide*	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	*Intel® Omni-Path Fabric Suite FastFabric User Guide*	Provides instructions for using the set of fabric management tools designed to simplify and optimize common fabric management tasks. The management tools consist of TUI menus and command line interface (CLI) commands.
	Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide	Describes the command line interface (CLI) for the Intel® Omni-Path Fabric Suite FastFabric. 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	*Intel® Omni-Path Fabric Suite Fabric Manager User Guide*	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.
	Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide	Provides an intuitive, scalable dashboard and set of analysis tools for graphically monitoring fabric status and configuration. It is a user-friendly alternative to traditional command-line tools for day-to-day monitoring of fabric health. Help: Fabric Manager GUI Online Help.
Configuring and administering Intel® HFI and IPoIB driver Running MPI applications on Intel® OPA	*Intel® Omni-Path Fabric Host Software User Guide*	Describes how to set up and administer the Host Fabric Interface (HFI) after the software has been installed. The audience for this document includes both cluster administrators and Message-Passing Interface (MPI) application programmers, who have different but overlapping interests in the details of the technology.
Writing and running middleware that uses Intel® OPA	*Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide*	Provides a reference for programmers working with the Intel® PSM2 Application Programming Interface (API). The Performance Scaled Messaging 2 API (PSM2 API) is a low-level user-level communications interface.
Optimizing system performance	*Intel® Omni-Path Fabric Performance Tuning User Guide*	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 a storage router on Intel® OPA	*Intel® Omni-Path IP and Storage Router Design Guide*	Describes how to install, configure, and administer an IPoIB router solution (Linux* IP or LNet) for inter-operating between Intel® Omni-Path and a legacy InfiniBand* fabric.
Building a Lustre* Server using Intel® OPA	*Building Lustre* Servers with Intel® Omni-Path Architecture Application Note*	Describes the steps to build and test a Lustre* system (MGS, MDT, MDS, OSS, OST, client) from the HPDD master branch on a x86_64, RHEL*/CentOS* 7.1 machine.

Table 3-1.   Intel® Omni-Path Documentation Library (Sheet 3 of 3)

Task	Document Title	Description
Building Containers for Intel® OPA Fabrics	*Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note*	Provides basic information for building and running Docker* and Singularity* containers on Linux*-based computer platforms that incorporate Intel® Omni-Path networking technology.
Learning about new release features, open issues, and resolved issues for a particular release	*Intel® Omni-Path Fabric Software Release Notes*	
	Intel® Omni-Path Fabric Fabric Manager GUI Release Notes	
	Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)	