



Intel® Omni-Path Software

Release Notes for V10.7

Rev. 1.0

April 2018



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

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

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

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

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

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

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 © 2018, Intel Corporation. All rights reserved.



Contents

- 1.0 Overview of the Release..... 5**
 - 1.1 Audience..... 5
 - 1.2 Software License Agreement..... 5
 - 1.3 If You Need Help..... 5
 - 1.4 Enhancements and Features in this Release..... 5
 - 1.5 Supported Features..... 6
 - 1.6 Deprecated and Removed Features..... 7
 - 1.7 Release Packages..... 7
 - 1.8 Release Compatibility..... 8
 - 1.9 Firmware Files..... 8
 - 1.10 Operating Systems..... 8
 - 1.11 Parallel File Systems..... 9
 - 1.12 MPI Libraries..... 9
 - 1.12.1 Supported MPI Libraries..... 9
 - 1.12.2 Compiler Versions and Distributions..... 10
 - 1.13 Intel Hardware..... 10
 - 1.14 Switch Firmware..... 11
 - 1.15 Document Versions..... 11
 - 1.16 Installation Requirements..... 12
 - 1.16.1 Software and Firmware Requirements..... 12
 - 1.16.2 Installation Instructions..... 12
 - 1.16.3 Installation Path Changes in Release 10.4 (and later)..... 12
 - 1.17 Product Constraints..... 13
 - 1.18 Product Limitations..... 13
 - 1.19 Accelerated RDMA Information..... 14
- 2.0 Issues..... 15**
 - 2.1 Resolved Issues..... 15
 - 2.1.1 Issues Resolved in this Release..... 15
 - 2.1.2 Issues Resolved in Prior Releases..... 16
 - 2.2 Open Issues..... 17
- 3.0 Related Information..... 21**
 - 3.1 Intel® Omni-Path Documentation Library..... 21
 - 3.1.1 How to Search the Intel® Omni-Path Documentation Set..... 23



Tables

1	Firmware Files.....	8
2	Supported Operating Systems.....	8
3	Supported MPI Libraries.....	9
4	Compiler Versions and Distributions.....	10
5	Supported Hardware.....	10
6	Supported Document Versions.....	11
7	Issues Resolved in this Release.....	15
8	Issues Resolved in Prior Releases.....	16
9	Open Issues.....	17



1.0 Overview of the Release

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.

1.1 Audience

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

1.2 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.

1.3 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.

1.4 Enhancements and Features in this Release

The following enhancements and features are new in this release:

- Additional operating systems supported. See [Table 2](#) on page 8 for details.
- Additional hardware. See [Table 5](#) on page 10 for details.
- Topology-aware job scheduling, which is enabled by the opamgt library and allows developers to write code that interfaces to the SA/PA. New features in this release include: Added EM and EA headers (stl_em_eostl.h and stl_ea_eostl.h resp) to opamgt RPM. See the *Intel® Omni-Path Management API Programmer's Guide* for details.
- Support for IntelSM C/C++ Compiler 18.0
- UEFI, TMM, and Firmware Tools are now standalone rpms.
- Product Constraints described in [Product Constraints](#) on page 13.



1.5 Supported Features

- The list of supported operating systems is in [Table 2](#) on page 8.
- The list of supported hardware is in [Table 5](#) on page 10.
- 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>
- 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 [Accelerated RDMA Information](#) on page 14 for details.
- Support for active optical cables (AOC) on server platforms using integrated HFI for OPA (commonly known as "-F").
- 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
- Signed Kernel Modules, as required to support UEFI Secure Boot.
- Support for NVMe over Fabric Protocol
- Support for IBM* Platform MPI and IBM* Spectrum MPI. See [Table 3](#) on page 9 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.
- In Release 10.4, changes were made to the installation path for all Intel® Omni-Path Software files. See [Installation Path Changes in Release 10.4 \(and later\)](#) on page 12 for details and user action.
- Legacy BIOS Boot Mode Enhancements to support boot over fabric, custom board descriptions, and pre-boot platform configuration data for AOC support.
- Multi-endpoint functionality. See the *Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide* for details.
- SNMP MIB support. See the *Intel® Omni-Path Fabric Switches Release Notes* for details.
- Support for Power Class 2 active optical cables (AOC). See [Product Constraints](#) on page 13 for more information.
- Sandia* OpenSHMEM over Open Fabrics Interface (OFI).
- The `openmpi_gcc_hfi-X.X.X` version includes support for Open Fabrics Interface (OFI) libfabric. See the *Intel® Omni-Path Fabric Host Software User Guide* for details.



- Open Fabrics Interface (OFI) libfabric is provided in the Basic package. Existing libfabric installations will be upgraded when the Basic package is installed. fabtests that support 1.4.x versions of libfabric are not guaranteed to work with the Basic package version.

1.6 Depreciated and Removed Features

- In this release, Host Fabric Manager support for infinite Switch Lifetime Limit (SLL) and HoQ lifetime limit (HoQLife) have been removed.
- In V10.8, the INSTALL script will be modified to remove the following options:
 - rebuild Intel® OPA-specific OFA Delta user modules
 - rebuild gasnet
 - rebuild openshmem
 - rebuild verbs MPIs
 - rebuild uefi

If you need information on how to build Intel® OPA-specific OFA delta user modules, contact Intel Customer Support for details.

1.7 Release Packages

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

- Basic for compute nodes
- IFS for management nodes

The Basic package includes:

- Software that installs the following packages to the distribution OpenFabrics Alliance* (OFA):
 - hfi1-firmware, libpsm2 (for all RHEL* and SLES* 12 SP2) and ibpsm2-2 (for SLES* 12 SP3 and newer), hfi1-diahtools-sw
 - Open MPI and MVAPICH2. See [MPI Libraries](#) on page 9 for details.
 - mpitests
 - mpi-selector
 - Sandia* OpenSHMEM
 - Open Fabrics Interface (OFI) libfabric
 - Firmware files listed in [Table 1](#) on page 8.
- 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.
- Firmware installation tools, including hfi1_eprom and TMM update tools.

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.
- 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 *Intel® Omni-Path Fabric Suite FastFabric User Guide*.

1.8 Release Compatibility

This release is backwards compatible with the most recent minor release version. For example, Release 10.N is backwards compatible with Release 10.N-1 and Release 10.N-1.x.

1.9 Firmware Files

This release of the Intel® Omni-Path Software contains the firmware files listed in the following table.

Table 1. Firmware Files

Description	File Name	Version
Intel® Omni-Path HFI platform file ¹	hfi1_platform.dat	HFI_TYPE1 v1.0.1.0
<i>Note:</i> 1. If you have a non-Intel HFI, contact the manufacturer's support team for details.		

1.10 Operating Systems

This release of the Intel® Omni-Path Software supports the operating systems listed in the following table.

Table 2. Supported Operating Systems

Operating System	Update/ SP	Kernel Version
Red Hat* Enterprise Linux* (RHEL*) 7.3 X86_64	Update 3	3.10.0-514.el7.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.4 X86_64	Update 4	3.10.0-693.el7.x86_64
CentOS* 7.3 X86_64	Update 3	3.10.0-514.el7.x86_64
CentOS* 7.4 X86_64	Update 4	3.10.0-693.el7.x86_64
Scientific Linux* 7.3 X86_64	Update 3	3.10.0-514.el7.x86_64
Scientific Linux* 7.4 X86_64	Update 4	3.10.0-693.el7.x86_64
SUSE* Linux* Enterprise Server (SLES*) 12.2 X86_64	Service Pack 2	4.4.21-69.1-default
SUSE* Linux* Enterprise Server (SLES*) 12.3 X86_64	Service Pack 3	4.4.73-5.1-default



- Notes:
- The Intel® Xeon Phi™ x200 Product Family (Knights Landing CPU-based servers) and Intel® Xeon® Scalable Processors (Skylake CPU-based servers) support the following operating systems: RHEL* 7.3, RHEL* 7.4, CentOS* 7.3, SLES* 12 SP2, and SLES* 12 SP3.
 - The next generation Intel® Xeon Phi™ Processor (codename Knights Mill) supports the following operating systems: RHEL* 7.4 and SLES* 12 SP3.
 - PSM2 GPUDirect* RDMA with CUDA ToolKit 8.x is supported on RHEL* 7.3, RHEL* 7.4, and SLES* 12 SP2. Support for CUDA Toolkit 9.1 is also available with minimal validation for this list of operating systems.

1.11 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 (formerly known as General Parallel File System, GPFS), version 4.2.1.0
 - RHEL* versions supported by Intel® Omni-Path Software.

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

1.12 MPI Libraries


1.12.1 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 3. Supported MPI Libraries

MPI Implementation	Included in Basic package?	Runs Over	Compiled With
Open MPI 2.1.2	Yes	Verbs	N/A
	Yes	PSM2	GCC, Intel
Open MPI 2.1.2-cuda	No	Verbs	N/A
	Yes	PSM2	GCC
MVAPICH2-2.3B	Yes	Verbs	GCC
	Yes	PSM2	GCC, Intel
Intel® MPI Library 2017 Update 3	No	Verbs	N/A
	No	PSM2	N/A
Intel® MPI Library 2018	No	Verbs	N/A

continued...



MPI Implementation	Included in Basic package?	Runs Over	Compiled With
	No	PSM2	N/A
IBM* Platform* MPI version 9.1.4.3	No	Verbs	N/A
	No	PSM2	N/A
IBM* Spectrum* MPI version 10.1.1.0	No	Verbs	N/A
	No	PSM2	N/A

1.12.2 Compiler Versions and Distributions

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

Table 4. Compiler Versions and Distributions

Compiler	OS Distribution	Compiler Version
(GNU) gcc	RHEL* 7.3	gcc version 4.8.5 20150623 (Red Hat* 4.8.5-11) (GCC)
(GNU) gcc	RHEL* 7.4	gcc version 4.8.5 20150623 (Red Hat 4.8.5-16) (GCC)
(GNU) gcc	SLES* 12 SP 2	gcc version 4.8.5 (SUSE* Linux*)
(GNU) gcc	SLES* 12 SP 3	gcc version 4.8.5 (SUSE* Linux*)
(Intel) icc	RHEL* 7.3	icc (ICC) 20170811, icc version 18.0.0 (gcc version 4.8.5 compatibility)
(Intel) icc	RHEL* 7.4	icc (ICC) 20170811, icc version 18.0.0 (gcc version 4.8.5 compatibility)
(Intel) icc	SLES* 12 SP 2	icc (ICC) 20170811, icc version 18.0.0 (gcc version 4.8.0 compatibility)
(Intel) icc	SLES* 12 SP 3	icc (ICC) 20170811, icc version 18.0.0 (gcc version 4.8.0 compatibility)

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

1.13 Intel Hardware

The following table lists the Intel hardware supported in this release. The table does not include OEM-specific hardware, such as custom adapters and switches.

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

Table 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
Intel® Xeon® Scalable Processors	Skylake CPU-based servers
Intel® Xeon Phi™ x200 Product Family	Knights Landing CPU-based servers
Next generation Intel® Xeon Phi™ Processor (codename Knights Mill)	Knights Mill CPU-based servers
Intel® Omni-Path Host Fabric Interface 100HFA016 (x16)	Single Port Host Fabric Interface (HFI)
<i>continued...</i>	



Hardware	Description
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

1.14 Switch Firmware

The following firmware is supported for Intel® Omni-Path switches:

- Intel® Omni-Path Switch Firmware 10.7.x revision (managed and externally-managed switches)
- Intel® Omni-Path Switch Firmware 10.6.x revision (managed and externally-managed switches)

Refer to the *Intel® Omni-Path Fabric Switches Release Notes* for more information.

1.15 Document Versions

The following table lists the end user document versions supported by this release.

Table 6. Supported Document Versions

Title	Doc. Number	Revision
<i>Intel® Omni-Path Fabric Quick Start Guide</i>	J57479	4.0
<i>Intel® Omni-Path Fabric Setup Guide</i>	J27600	8.0
<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>	H76456	7.0
<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>	H76466	5.0
<i>Intel® Omni-Path Fabric Software Installation Guide</i>	H76467	9.0
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	H76457	9.0
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	H76458	9.0
<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i> (Merged with: <i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i>)	H76469	9.0
<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>	H76468	9.0
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>	H76471	9.0
<i>Intel® Omni-Path Fabric Host Software User Guide</i>	H76470	9.0
continued...		



Title	Doc. Number	Revision
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	9.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	11.0
Intel® Omni-Path IP and LNet Router Design Guide (Old title: Intel® Omni-Path IP and Storage Router Design Guide)	H99668	6.0
Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note	J57474	4.0
Intel® Omni-Path Management API Programmer's Guide	J68876	3.0
Configuring Non-Volatile Memory Express* (NVMe*) over Fabrics on Intel® Omni-Path Architecture Application Note	J78967	1.0
Intel® Omni-Path Fabric Software Release Notes	J95967	1.0
Intel® Omni-Path Fabric Manager GUI Release Notes	J95968	1.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)	J95964	1.0
Intel® Omni-Path Fabric Unified Extensible Firmware Interface (UEFI) Release Notes	J98868	1.0
Intel® Omni-Path Fabric Thermal Management Microchip (TMM) Release Notes	J98871	1.0
Intel® Omni-Path Fabric Firmware Tools Release Notes	J98870	1.0

Related Links

[Intel Omni-Path Documentation Library](#) on page 21

1.16 Installation Requirements

1.16.1 Software and Firmware Requirements

Table 2 on page 8 lists the operating systems supported by this release. Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for the required packages.

1.16.2 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 this 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 Switches Hardware Installation Guide* for related firmware requirements.

1.16.3 Installation Path Changes in Release 10.4 (and later)

If you are upgrading an Intel® Omni-Path Fabric Software installation (Release 10.3 or earlier), Intel recommends that you perform the following steps before upgrading, due to changes in installation paths for RPMs and configuration files.



Run `./INSTALL -u` to uninstall existing packages.

Run `./INSTALL -a` to complete the installation.

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 (Release 10.3 and earlier)	New Location (Release 10.4 and later)
<code>/etc/sysconfig/opafm.xml</code>	<code>/etc/opa-fm/opafm.xml</code>
<code>/etc/sysconfig/allhosts</code>	<code>/etc/opa/allhosts</code>
<code>/etc/sysconfig/chassis</code>	<code>/etc/opa/chassis</code>
<code>/etc/sysconfig/esm_chassis</code>	<code>/etc/opa/esm_chassis</code>
<code>/etc/sysconfig/hosts</code>	<code>/etc/opa/hosts</code>
<code>/etc/sysconfig/opafastfabric.conf</code>	<code>/etc/opa/opafastfabric.conf</code>
<code>/etc/sysconfig/opaff.xml</code>	<code>/etc/opa/opaff.xml</code>
<code>/etc/sysconfig/opamon.conf</code>	<code>/etc/opa/opamon.conf</code>
<code>/etc/sysconfig/ports</code>	<code>/etc/opa/ports</code>
<code>/etc/sysconfig/switches</code>	<code>/etc/opa/switches</code>

1.17 Product Constraints

- Power class 2 AOC are supported and require the same level firmware on both ends of the cable. Specifically, 10.6 host software and 1.6 level UEFI are required for proper operation. Integrated HFI (-F) requires a specific BIOS level to support power class 2 AOC; contact your BIOS vendor for more information.
- On certain systems with more than 1 HFI, the ports can come up in an unexpected way. This can lead to the wrong HFI being associated with `ib0` or `ib1`. To ensure ports come up in the expected order each time, use the module parameter `hfi1 port_reorder=1` and if you are loading the driver during `initramfs` boot, then ensure that you rebuild `initramfs`.

1.18 Product Limitations

This release has the following product limitations:

- The embedded version of the Fabric Manager supports a maximum of 100 HFI ports involving less than 20 switch ASICs. Calculate the number of switch ASICs in your fabric as follows:
 - One ASIC per Intel® Omni-Path Edge Switch 100 Series
 - Two ASICs per Intel® Omni-Path Director Class Switch 100 Series Leaf module
 - Two ASICs per Intel® Omni-Path Director Class Switch 100 Series Spine module
- Performance Administration (PA) Failover should **not** be enabled with FMs running on differing software versions.



To disable PA failover, edit the `/etc/sysconfig/opafm.xml` file and in the `<Pm>` section, change `<ImageUpdateInterval>` to 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.

1.19 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

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

2.1 Resolved Issues

2.1.1 Issues Resolved in this Release

The following table lists issues that are resolved in this release.

Table 7. Issues Resolved in this Release

ID	Description	Resolved in Release
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	10.7
134471	The HFI UEFI driver cannot boot via PXE using Grub 2.	10.7
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.	10.7
135084	In rare circumstances, the HFI may not appear in the PCI config space after a power cycle.	10.7
135975	After performing an OPA software configuration update, switches will show the new settings when queried by opaswitchadmin tools, however, individual ports will continue to operate using the previous settings, including LinkWidth enable.	10.7
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.	10.7
136500	RDMA perftests can hang on start on a client side when RDMA CM (-R option) is used.	10.7
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. In Release 10.6, a new configuration value is present in the FM configuration that determines how much time will be allotted to timeouts before abandoning a sweep. If you are upgrading from a previous version of the FM and retaining a configuration file that does not include this new parameter, the value will be set too low and cause sweeps to abandon after only a single timeout is witnessed.	10.7
137951	In the HFI BIOS screen for Advanced NIC Configuration, a warning message about incorrect custom P_Key value is not completely displayed.	10.7

continued...



ID	Description	Resolved in Release
139550	Infrequently, an AOC may exhibit an unexpectedly high local link integrity error rate after the link comes up, relative to the error rate on previous link up occasions. This can be determined by observing a link quality of <5. These links may eventually experience a link width downgrade.	10.7
139981	When Accelerated RDMA (TID RDMA) is enabled, GPFS file system can hang.	10.7
140691	When running <code>opaswitchadmin</code> against multiple externally-managed switches simultaneously, it sends schedule requests in parallel to those hosts. It is possible that some hosts may intermittently fail due to timeouts at high levels of parallelism.	10.7
140707	When using the Upgrade option in the INSTALL script TUI, a message is displayed stating that "Up To Date" items will be reinstalled. However, the <code>opa.log</code> file does not indicate these items were reinstalled.	10.7
140881	In rare cases when an LNI failure occurs, the link will not come up after manually disabling and re-enabling the link.	10.7
140911	In Release 10.6, the OFI verbs provider does not support FI_EP_RDM End Point type. This End Point type is needed for Open MPI OFI support. Therefore, Open MPI OFI support will not run over the verbs provider.	10.7
141219	When adaptive routing is disabled, the output for <code>opasmaquery</code> for <code>portgroup</code> appears as shown below: <pre># opasmaquery -l 1 -o portgroup PG: 0x0000 Egress:None</pre>	10.7
141420	When Accelerated RDMA is enabled, the <code>kdeth_qp</code> module parameter cannot be changed.	10.7
141558 141852	In IFS versions 10.3.1 and earlier, the <code>rpm</code> contains the Epoch tag. When upgrading to a newer version without the Epoch tag (i.e., 10.4 through 10.6.x), the <code>rpm</code> tools act as though the old <code>rpm</code> is the newer version. This issue causes a failure in the IFS upgrade.	10.7
141782	In Release 10.7, the <code>opapmaquery -n</code> argument no longer accepts decimal input. In the user documentation and the man pages, the <code>-n</code> argument is described as the port in hexadecimal. Until this release, the tool also accepted decimal input.	10.7
141845	Resolved FM process out of memory condition	10.7
141909	Resolved multiple FM synchronization issue that can lead to FM failure.	10.7

2.1.2 Issues Resolved in Prior Releases

The following table lists issues that were resolved in prior releases.

Table 8. Issues Resolved in Prior Releases

ID	Description	Resolved in Release
135390	In Release 10.6.1, the driver can parse older versions of the platform configuration file.	10.6.1
141142	In Release 10.6.1, issues loading the HFI driver with GPUDirect* RDMA CUDA support have been resolved.	10.6.1
133633	OpenMPI and MVAPICH2 compiles fail to link properly when using the Intel compilers. This issue is caused by a bug in Intel Compiler 2015.	10.6
134904	Legacy PXE boot using iPXE while the HFI UEFI driver is loaded may cause a hang.	10.6
136419 137106	When running SLES* 12.2 with inbox OPA drivers installed, the state may not change from "Offline" to "Physical Linkup (Init)" as expected. When running SLES* 12.3 with inbox OPA drivers, this issue does not occur.	10.6
<i>continued...</i>		



ID	Description	Resolved in Release
137409	When using DHCP from an Intel® OPA HFI, the DHCP client-identifier field (option 61) must be used to set up DHCP static leases. However, what is sent in the DHCP discover packet is different for PXE boot attempts versus a DHCP boot attempt. The DHCP client-identifier can be explicitly specified in DHCP client configuration or NetworkManager scripts.	10.6
138909	Installation fails due to a build error in IntelOPA-Tests, which is caused by a version mismatch between the ICC libraries in the IFS package and the ICC libraries available in the system.	10.6
139397	IPoIB traffic stalls during reboot testing.	10.6
139660	Following a boot it is possible, although rare, that the IPoIB interface will fail to come up. Hosts attempting to ping this host will get no response.	10.6
139692	On SLES* 12 SP2, when installing the Release 10.5 software with this command: <code>rpm -i hfidiaqs-0.8-66.x86_64.rpm</code> The following error is returned: <code>error: Failed dependencies: libreadline.so.7() (64bit) is needed by hfidiaqs-0.8-66.x86_64</code>	10.6
139834	When using the FastFabric TUI to run "Perform Single Host Verification", the test hangs during operation.	10.6
140073	Error message when running MPI tests: <code>Unable to initialize PSM2 CUDA support</code> In Release 10.5, if IFS is installed with CUDA* support, then the NVIDIA* CUDA* Toolkit is also required to be installed in order to run any MPI application using PSM2 transport. Release 10.5 IFS installations without CUDA* support do not require NVIDIA* CUDA* Toolkit to be installed.	10.6
140151	On RHEL* 6.7, during IFS 10.5 installation, warning messages were displayed about the <code>/etc/sysconfig/irqbalance</code> file. This issue has been resolved.	10.6
140208	The node description of a node may change after rebooting. This issue has been seen on RHEL* and SLES* distributions and requires restart of the <code>rdma-ndd</code> service.	10.6
140229	In Release 10.6, the <code>opaswitchadmin</code> tool was updated to address a condition that was seen during firmware upgrade of a large number of switches.	10.6
140281	On SLES* 12 SP2, during IFS 10.5 installation, a warning message about <code>IRQBALANCE_ARGS</code> was displayed. This issue has been resolved.	10.6
140527	In Release 10.6, an installation issue caused by a dependency on <code>libuuid-devel</code> has been resolved.	10.6
140909	In Release 10.6, the filepath for <code>bios_images</code> was updated from: <code>/opt/opa/bios_images</code> to <code>/usr/share/opa/bios_images</code>	10.6


2.2 Open Issues

The following table lists the open issues for this release.

Table 9. Open Issues

ID	Description	Workaround
129563	Memory allocation errors with <code>MVAPICH2-2.1/Verbs</code> .	<i>Note:</i> To avoid this issue, use MPIs over PSM. If you are using MPIs over verbs, the following workaround is required:

continued...



ID	Description	Workaround
		<ul style="list-style-type: none"> When running MVAPICH2 jobs with a large number of ranks (for example, > 36 ranks but ≤ 72 ranks), you must set the following parameters in <code>/etc/security/limits.conf</code>: <ul style="list-style-type: none"> — hard memlock unlimited — soft memlock unlimited Also, you must increase the <code>lkey_table_size:LKEY</code> table size in bits (2^n, where $1 \leq n \leq 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>, HFI1 Driver Module Parameters chapter.
132207	Kernel crash caused by the <code>ib_srpt</code> module.	Install this kernel patch: https://git.kernel.org/cgit/linux/kernel/git/torvalds/linux.git/commit/?id=51093254bf879bc9ce96590400a87897c7498463
133604	Bonding driver shows incorrect hardware address of IPoIB interfaces.	This workaround applies to RHEL* 7.3 (and earlier) and SLES* 12 SP3 (and earlier) distributions: Use the <code>opainfo</code> command to retrieve the PortGUID and <code>ip addr show ib0</code> to get the correct 20-byte hardware address of OPA network interface.
134819	In KNL-F EFI shell, the command <code>ifconfig -l</code> does not correctly display the IP address after being assigned via DHCP.	Launch a newer version of the EFI shell from the embedded shell.
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.
135830	On Intel® Xeon Phi™ systems, failure observed during software upgrade when rebuilding the boot image. Error message contains: Rebuilding boot image with <code>"/usr/bin/dracut -f"</code>	Due to the extended processing time of the <code>dracut</code> command on the Intel® Xeon Phi™ platform, Intel recommends the following: <ul style="list-style-type: none"> Install and configure Intel® Xeon Phi™ systems separately. Change the <code>FF_TIMEOUT_MULT</code> value in <code>opafastfabric.conf</code> from 2 to 6 for Intel® Xeon Phi™ systems.
135981 143631	In RHEL 7.3, <code>irqbalance</code> spams messages every ten seconds with snippet below. The HFI driver sets affinity and the Omni Fabric install sets <code>/etc/sysconfig/irqbalance</code> to use <code>exact</code> . <pre>IRQBALANCE_ARGS= --hintpolicy=exact</pre> Irq balancer scans the <code>isolcpus</code> and <code>nohz_full</code> kernel masks and adds the corresponding CPUs to the <code>banned_cpus</code> mask. This works fine for valid masks, but not for the default, empty masks. In this case, when they read from the <code>sysfs</code> they return empty strings, <code>"\n"</code> or <code>"0x0, \n"</code>	Uncomment the following from <code>/etc/sysconfig/irqbalance</code> : <pre>#IRQBALANCE_BANNED_CPUS=</pre>
136432	Certain <code>perftest</code> tools such as <code>ib_write_bw</code> do not work on RHEL* 7.3 when using the RDMA CM with UD QPs.	Roll back the <code>perftest</code> package to the level found in RHEL* 7.2, which is <code>perftest-2.4</code> . Then install this package on RHEL* 7.3.
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.

continued...



ID	Description	Workaround
139368	<p>Some applications compiled with older compilers may use a personality bit that signifies that READ should imply EXECUTE permissions.</p> <p>To improve system security, the hfi1 driver does not allow execute permissions on PSM memory maps. Therefore, applications that use READ implies EXECUTE will fail to run.</p>	<p>As root, run the execstack tool to clear the executable bit on the binary:</p> <pre>execstack -c <binary></pre> <p>Alternatively, recompile the binary to not set this personality bit.</p>
139613	<p>The Subsystem Vendor and Subsystem Device ID in the PCI configuration space of Intel® Omni-Path discrete HFI cards may not indicate the correct OEM vendor and device. As a result, the <code>lspci</code> command may show incorrect Subsystem Vendor and Device ID information. This issue affects Intel server boards for Intel® Xeon® Processor v3 and v4 Product Family configured in Legacy OS boot mode.</p>	<p>Reconfigure the system from Legacy OS boot mode to UEFI boot mode.</p>
139743 143031 143115	<p>Under a very heavy load through the IPoIB interface, the kernel warning <code>NETDEV WATCHDOG: ib0 (hfi1): transmit queue 0 timed out</code>, followed by the messages <code>queue stopped 1</code>, <code>tx_head xxx</code>, <code>tx_tail xxx</code> and <code>transmit timeout: latency xxxx msecs</code> may be seen.</p>	<p>Reduce traffic pressure on ipoib interface to resume transmission.</p>
139924	<p>For RHEL* and SLES*, the <code>ibacmp</code> provided in the OS distribution uses incorrect address information when joining multicast groups. This causes name resolution to fail.</p> <p>The <code>dsap</code> provided in the OS distribution works correctly.</p>	<p>The fix for this issue is available in the library <code>rdma-core-15-2.el7</code>.</p> <ul style="list-style-type: none"> Do not use the OS distribution <code>ibacmp</code>. Install <code>rdma-core-15-2.el7</code> manually.
139995	<p>When installing Intel® Omni-Path Software on a clean system (with no prior versions installed), the following error message may be displayed:</p> <pre>cat: //etc/opa/version_delta: No such file or directory</pre>	<p>This message can be safely ignored. The installer is looking for an IFS version file before it has been created on the system. The installation is not impacted.</p>
140310	<p>On RHEL* 7.3, if an ipofabric interface name is non-standard, the <code>ifcfg</code> file is not read correctly.</p> <p>For example, if you are using the ipofabric interface name <code>opa_ib0</code>, and the <code>connected</code> mode and MTU size is specified in the <code>ifcfg-opa_ib0</code> file, when you bring up the interface with <code>ifup</code>, the settings do not take effect.</p>	<p>You must manually enter:</p> <pre>echo connected > /sys/class/net/opa_ib0/mode</pre> <p>for the settings to take effect.</p>
140797 141558 141852 142476	<p>When installing an IFS tar package on a server where the OS-resident OPA software has been installed, the installation may fail to properly update all the RPMs. This is due to the use of Epoch tagging on the OS-resident OPA software which makes the OS-resident RPMs appear to be newer than RPMs in the IFS tar package.</p>	<p>Uninstall the OS-resident OPA software before attempting to install from the IFS tar package.</p> <p>If you are manually updating individual RPMs, use the RPM argument <code>--oldpackage</code> to force the installation of the RPM.</p> <p>If you are using <code>yum update</code> or <code>yum-cron/auto-update</code>, add <code>exclude=opa-*</code> to the <code>/etc/yum.conf</code> file. Alternately you may modify the existing exclude list to contain the pattern, <code>opa-*</code>.</p>
141263	<p>The GUPS application that is included as part of Sandia Open Shmem apps may fail with more than 1% of error when executed with more than four PEs.</p> <p>Note: GUPS does not use atomic XOR operations given they are not supported in the OpenShmem specification, version 1.3. This is addressed in version 1.4.</p>	<p>Update the version of Sandia Open Shmem to 1.4.0.</p>
141273	<p>The in-distro version of <code>perftests</code> has bugs.</p>	<p>Use the upstream version of <code>perftest</code> from http://git.openfabrics.org/?p=~grockah/perftest.git;a=summary.</p>

continued...



ID	Description	Workaround
141586	Using SLES* 12 SP2 and SP3, INSTALL -a fails due to dependency issues.	Prior to installing Intel OPA, install the SLES SDK iso from SUSE, which contains additional, required development packages.
141636	During the Intel OPA installation, the ramdisk may be rebuilt multiple times; however, only the last build is necessary.	None. Ignore the messages.
141740	Due to the race condition during boot up ipoib driver can miss PKey change event that will leave ipoib interface in disconnected state.	The workaround is to reload ipoib driver on the affected host: #> modprobe -r ib_ipoib && modprobe ib_ipoib
141793	Use of a static buffer could produce an incorrect device name (hfi1_x) in dmesg logging.	None.
142330	MPI applications that leverage the PSM2 library's access to the HFI ASICs Memory Mapped IO and that access the MMIO directly (not via PSM2) can potentially cause an "unsupported opcode" error which some servers handle as a critical error.	Disable upstream error reporting using the AER mask register. <ul style="list-style-type: none"> For discrete HFI ASICs (e.g., CHF PCIe card), use <pre>setpci -d 8086:24f0 ECAP_AER +8.l=00100000:00100000</pre> For integrated HFIs (e.g., KNL-F and SKX-F), use <pre>setpci -d 8086:24f1 ECAP_AER +8.l=00100000:00100000</pre>
143296	When irqbalance uses the argument --hintpolicy=exact, it applies the policy of setting the hardware interrupts to CPU core mappings according to device drivers preferences. For the HFI1 driver, it is strongly recommended to preserve interrupt locality for low latency and high bandwidth by having a dedicated CPU core per interrupt.	Always start the user-space process irqbalance using the argument --hintpolicy=exact.
143311	During UEFI pre-boot with the connected switch running 10.7.0.0.134, the OPA link may not complete initialization if the link is bounced or restarted. This behavior is limited to the pre-boot period. There is no exposure once Linux boot has completed.	Avoid link bounce or switch reboots during server reboot periods. If the OPA link fails to come up during UEFI pre-boot, a host reboot is required to recover.
143449	PM will scroll LQI=0 and Integrity Exceeded Threshold logs when an additional VF with QoS enabled and a device group that is not "All". <i>Note:</i> This issue does not occur when running against the default opafm.xml configuration file.	Set the <ProcessVLCounters> field in the opaxml.fm configuration to 0 to stop scrolling of logs related to LQI.
143915	When specific messages that are of non-double word length (not 4 B) and less than 1 MTU in size (10,240 B in most configurations) are sent, they may never complete, causing an application to wait indefinitely. <i>Note:</i> This issue occurs on machines with more than 28 CPU cores (not including hyper-threads) or more than 28 HFI contexts enabled per Intel® Omni-Path HFI.	Add this command line option for OpenMPI: -x PSM2_CUDA_THRESH_RNDV=8192



3.0 Related Information

3.1 Intel® Omni-Path Documentation Library

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

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

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

Task	Document Title	Description
Key: Shading indicates the URL to use for accessing the particular document.		
• Intel® Omni-Path Switches Installation, User, and Reference Guides: http://www.intel.com/omnipath/SwitchPublications		
• Intel® Omni-Path Software Installation, User, and Reference Guides (includes HFI documents): http://www.intel.com/omnipath/FabricSoftwarePublications (no shading)		
• Drivers and Software (including Release Notes): http://www.intel.com/omnipath/Downloads		
Using the Intel® OPA documentation set	<i>Intel® Omni-Path Fabric Quick Start Guide</i>	A roadmap to Intel's comprehensive library of publications describing all aspects of the product family. It outlines the most basic steps for getting your Intel® Omni-Path Architecture (Intel® OPA) cluster installed and operational.
Setting up an Intel® OPA cluster	<i>Intel® Omni-Path Fabric Setup Guide</i>	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.
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. The Intel® HFI utilizes Intel® Omni-Path switches and cabling.
continued...		



Task	Document Title	Description
Installing host software Installing HFI firmware Installing switch firmware (externally-managed switches)	<i>Intel® Omni-Path Fabric Software Installation Guide</i>	Describes using a Text-based User Interface (TUI) to guide you through the installation process. You have the option of using command line interface (CLI) commands to perform the installation or install using the Linux* distribution software.
Managing a switch using Chassis Viewer GUI Installing switch firmware (managed switches)	<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	Describes the Intel® Omni-Path Fabric Chassis Viewer graphical user interface (GUI). It provides task-oriented procedures for configuring and managing the Intel® Omni-Path Switch family. Help: GUI 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> (Merged with: <i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i>)	Provides instructions for using the set of fabric management tools designed to simplify and optimize common fabric management tasks. The management tools consist of 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. It 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 both cluster administrators and Message-Passing Interface (MPI) application programmers, who have different but overlapping interests in the details of the technology.
Writing and running middleware that uses Intel® OPA	<i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>	Provides a reference for programmers working with the Intel® PSM2 Application Programming Interface (API). The Performance Scaled Messaging 2 API (PSM2 API) is a low-level user-level communications interface.
Optimizing system performance	<i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>	Describes BIOS settings and parameters that have been shown to ensure best performance, or make performance more consistent, on Intel® Omni-Path Architecture. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.
Designing an IP, LNet or storage router on Intel® OPA	<i>Intel® Omni-Path IP and LNet Router Design Guide</i> (Old title: <i>Intel® Omni-Path IP and Storage Router Design Guide</i>)	Describes how to install, configure, and administer an IPoIB router solution (Linux* IP or LNet) for inter-operating between Intel® Omni-Path and a legacy InfiniBand* fabric.
continued...		



Task	Document Title	Description
Building a Lustre* Server using Intel® OPA	(OBSOLETE) <i>Building Lustre* Servers with Intel® Omni-Path Architecture Application Note</i>	This document has been removed from the Intel® OPA Documentation Library. For information on how to build and configure a Lustre* server with Intel® OPA, see the Lustre* wiki: http://wiki.lustre.org .
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 (includes managed and externally-managed switches)</i>	
	<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>	

3.1.1 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.