

Intel[®] Omni-Path Fabric Software

Release Notes for 10.2

November 2016



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

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

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

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


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

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at <http://www.intel.com/> or from the OEM or retailer.

Intel, Xeon, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2015-2016, Intel Corporation. All rights reserved.



Contents

1.0 Overview of the Release	4
1.1 Introduction	4
1.2 Audience	4
1.3 Software License Agreement	4
1.4 If You Need Help	4
1.5 New Enhancements and Features in this Release	4
1.6 Supported Features	5
1.7 Release Packages	5
1.8 Firmware Files	6
1.9 Operating Systems	6
1.10 Parallel File System	7
1.11 Compilers	7
1.11.1 MPI	7
1.11.2 MVAPICH2 and Open MPI	7
1.12 Hardware	8
1.13 Installation Requirements	8
1.13.1 Software and Firmware Requirements	8
1.13.2 Installation Instructions	8
1.14 Product Constraints	10
1.15 Product Limitations	10
1.16 Documentation	10
2.0 Issues	12
2.1 Introduction	12
2.2 Resolved Issues	12
2.3 Open Issues	14
Tables	
1-1 Firmware Files	6
1-2 Operating Systems Supported	6
1-3 MPI Compilers	7
1-4 MVAPICH2 and Open MPI	7
1-5 Hardware Supported	8
1-6 Related Documentation for this Release	11
2-1 Resolved Issues	12
2-2 Open Issues	14



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 www.intel.com for additional detail.

1.5 New Enhancements and Features in this Release

The following enhancements and features are new for the 10.2 release:

- Added support tools to perform PA queries with time inputs.
- A new API has been added in the Linux* kernel v4.4 for a new memory registration model, so memory registration occurs for each RDMA command.
- Intel® Enterprise Edition for Lustre* software version 3.0
- Added support for RHEL* 6.7. See [Table 1-2](#) for details



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 *Intel® Omni-Path Fabric Host Software User Guide*.
- Supports Dual Rail: Two Intel® Omni-Path Host Fabric Interface (HFI) cards in the same server connected to the same fabric.
- Supports Dual Plane: Two HFI cards in the same server connected to separate fabrics.
- Limited validation testing performed on network storage file systems:
 - NFS over TCP/IP
- Active Optical Cables
- MPI applications are provided in a stand-alone rpm package.
- Intel® Xeon® v4 processor (codename Broadwell) support
- KNL support (B0 technology preview)
- 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.


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 in [Table 1-6](#).

1.8 Firmware Files

This release of the Intel® Omni-Path Software contains the firmware files listed in [Table 1-1](#).

Table 1-1. Firmware Files


Description	File Name	Version
HF11 UEFI Option ROM	HfiPcieGen3_0x29.efi	0x29
UEFI UNDI	HfiUndiStub_0x29.rom	0x29
HF11 SMBus Microcontroller Firmware (Thermal Monitor)	hfi1_smbus.fw	10.2.0.0.154

1.9 Operating Systems

This release of the Intel® Omni-Path Software supports the operating systems listed in [Table 1-2](#).

Table 1-2. Operating Systems Supported

Operating System	Update/SP	Kernel Version
Red Hat* Enterprise Linux* (RHEL*) 6.7 X86_64	Update 7	2.6.32-573.el7.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.1 X86_64	Update 1	3.10.0-229.el7.x86_64
Red Hat* Enterprise Linux* (RHEL*) 7.2 X86_64	Update 2	3.10.0-327.el7.x86_64
CentOS* 7.1 X86_64	N/A	3.10.0-229.el7.x86_64
CentOS* 7.2 X86_64	N/A	3.10.0-327.el7.x86_64
Scientific Linux* 7.1 X86_64	N/A	3.10.0-229.el7.x86_64
Scientific Linux* 7.2 X86_64	N/A	3.10.0-327.el7.x86_64
SUSE* Linux* Enterprise Server (SLES*) 12 X86_64	N/A	3.12.28-4-default
SUSE* Linux* Enterprise Server (SLES*) 12.1 X86_64	Service Pack 1	3.12.49-11.1-default



1.10 Parallel File System

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

- Intel® Enterprise Edition Lustre* (IEEL) 3.X
 - All supported RHEL versions.
- IBM* General Parallel File System (GPFS) version 4.0.1
 - All supported RHEL versions.

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 Compilers

1.11.1 MPI

This release supports the following MPI implementations:

Table 1-3. MPI Compilers

MPI Implementation	Runs Over	Compiled With
Open MPI 1.10.2	Verbs	GCC
	PSM2	GCC, Intel
MVAPICH2-2.1	Verbs	GCC
	PSM2	GCC, Intel
Intel® MPI 5.1.3	Verbs	GCC
	PSM2	GCC, Intel

1.11.2 MVAPICH2 and Open MPI

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

Table 1-4. MVAPICH2 and Open MPI (Sheet 1 of 2)

Compiler	Linux* Distribution	Compiler Version
(GNU) gcc	RHEL* 7	gcc (GCC) 4.8.2 20140120
(GNU) gcc	RHEL* 7.1	gcc (GCC) 4.8.3 20140911
(GNU) gcc	RHEL* 7.2	gcc (GCC) 4.8.5 20150623 (Red Hat* 4.8.5-4)
(GNU) gcc	SLES* 12	gcc (SUSE* Linux*) 4.8.3 20140627
(GNU) gcc	SLES* 12 SP 1	gcc (SUSE* Linux*) version 4.8.5
(Intel) icc	RHEL* 7	icc (ICC) 15.0.1
(Intel) icc	RHEL* 7.1	icc (ICC) 15.0.1



Table 1-4. MVAPICH2 and Open MPI (Sheet 2 of 2)

Compiler	Linux* Distribution	Compiler Version
(Intel) icc	RHEL* 7.2	icc (ICC) 15.0.1
(Intel) icc	SLES* 12	icc (ICC) 15.0.1
(Intel) icc	SLES* 12 SP 1	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.

Table 1-5. Hardware Supported

Hardware	Description
Intel® Xeon® v3 Processor	Haswell CPU-based servers
Intel® Xeon® v4 Processor	Broadwell 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 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 the following hardware is supported:

- Intel® Xeon® v3 Processor
- Intel® Xeon® v4 Processor

1.13 Installation Requirements


1.13.1 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.

1.13.2 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 10.2 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.

1.13.2.1 Installation Prerequisites for RHEL* 6.7

Install the following packages using yum from the RHEL* 6.7 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

1.13.2.2 Disable SRP Daemon Autostart

By default, RHEL* 7.1 is configured to autostart the SRP daemon on every node. This daemon performs a fabric search for SRP target devices every minute, which may impact the performance of HPC applications, especially in larger fabrics.

Intel recommends you disable the SRP daemon on all compute nodes and any service nodes. Typically, the SRP daemon only needs to be enabled on filesystem server nodes that are directly accessing block storage devices.

When installing either the IntelOPA-Basic or IntelOPA-IFS packages, there is a prompt: `Enable SRP initiator autoload?` Intel recommends you use the default option `no` on all compute nodes and any service nodes.



When installing using other means (such as provisioning systems or direct installation of the rpms), you must manually disable the SRP daemon autostart. After installing the rdma-*
noarch.rpm package, edit the /etc/rdma/rdma.conf file. The default setting is: SRP_LOAD=yes. Intel recommends you change this to no.

1.14 Product Constraints

None.

1.15 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 limited memory and processing resources available in the embedded environment.
- PXE boot can be accomplished using UEFI BIOS. PXE boot is not supported in legacy boot mode. Refer to the *Intel® Omni-Path Fabric Software Installation Guide*, "HFI UEFI PXE Installation and Configuration" Appendix for more information.
- 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.
- RHEL6.7 Support:
 - CPU support:
 - Intel® Xeon® v3 Processor - Haswell CPU-based servers, Intel® Xeon® v4 Processor - Broadwell CPU-based servers
 - File system support:
 - GPFS
 - NFS
 - Lustre
 - 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	gcc (GCC) 4.4.7
(Intel) icc	RHEL* 6.7	icc (ICC) 15.0.1

- Performance:
 - PSM bandwidth
 - MPI latency
 - Verbs bandwidth within 2%-5% of RHEL7.2 performance

1.16 Documentation

Table 1-6 lists the end user documentation for the current release.




Documents 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

Table 1-6. Related Documentation for this Release

Document Title
Hardware Documents
<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>
<i>Intel® Omni-Path Edge Switch Platform Configuration Reference Guide</i>
<i>Intel® Omni-Path Fabric Managed Switches Release Notes</i>
<i>Intel® Omni-Path Fabric Externally-Managed Switches Release Notes</i>
<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>
Fabric Software Documents
<i>Intel® Omni-Path Fabric Software Installation Guide</i>
<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>
<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>
<i>Intel® Omni-Path Fabric Host Software User Guide</i>
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI Online Help</i>
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>
<i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i>
<i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>
<i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>
<i>Intel® Omni-Path Host Fabric Interface Platform Configuration Reference Guide</i>
<i>Intel® Omni-Path Fabric Software Release Notes</i>
<i>Intel® Omni-Path Fabric Manager GUI Release Notes</i>
<i>Intel® Omni-Path Storage Router Design Guide</i>
<i>Intel® Omni-Path Fabric Staging Guide</i>
<i>Building Lustre* Servers with Intel® Omni-Path Architecture Application Note</i>



2.0 Issues

2.1 Introduction

This section provides a list of 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. Resolved Issues (Sheet 1 of 2)

ID	Component	Description	Resolved in Release
130233	HFI Host Pre-boot Software	Intel® Omni-Path DHCP fails to ID system based on hardware ID is no longer an issue due to the boot solution being used has been deprecated for Omni-Path.	10.2
130435	HFI Host Driver	Packet traffic stops after reconnecting a Fabric Manager HFI port to a management allowed port that was previously marked "management not allowed".	10.2
133269	Host Software/Verbs	GFP_NOIO error message using GFP_KERNEL when using IPoIB on OPA in connected mode.	10.2
132345	Host Fabric Software/MPI	When using Open MPI with the SLURM PMI launcher and PSM2, you may see the following error: Error obtaining unique transport key from ORTE (orte_precondition_transports not present in the environment). This is expected behavior. For details, see the <i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i> , in the description of psm2_ep_open() function, unique_job_key parameter.	10.2
132714	Host Fabric Software/MPI	The latest mvapich2 requires autoconf 2.67/automake 1.15/ libtool 2.4.3 or newer to rebuild.	10.2
132718	Software Installation	The Intel® Omni-Path SLES* 12.1 package contains infiniband-diags version 1.6.4 which does not include the rdma-ndd service.	10.2
133038	HFI Host Driver	PSM fails to acquire a context despite the kernel having contexts available. The following error message is returned: assign_context command failed: Device or resource busy This cannot be duplicated for resolution.	10.2
133607	Fabric Management Tools/FastFabric	opaverifyhosts fails to run on RHEL* 7.2	10.2
133676	Fabric Management Software/Fabric Manager	Creating a virtual fabric with standby enabled causes opapaqueries to fail.	10.2
133709	HFI Host Fabric Software/Tools	Transmit Discard (xmit discard) counts are captured at the port level through the opapmaquery command. opapmaquery additionally shows per VL counts for xmit discard (based on tool options), but currently shows zero counts.	10.2
133820	HFI Host Pre-boot Software	The UEFI driver may not work on some systems. The system freezes when the UEFI driver starts.	10.2



Table 2-1. Resolved Issues (Sheet 2 of 2)

ID	Component	Description	Resolved in Release
133918	HFI Host Pre-boot Software	Booting over Omni-Path HFI occasionally fails with "buffer too small, need netbuf_size = 16777216".	10.2
133987	Fabric Management Software/Fabric Manager	SM crash with general protection fault during compounding of PM sweep image	10.2
134118	HFI Host Fabric Software/ PSM2	PSM2_MQ_RNDV_HFI_WINDOW stalls HPL on bin1 KNL nodes	10.2
134131	HFI Host Pre-boot Software	Servers with dual HFI PCIe cards become unbootable after flashing the UNDI 0x1c firmware to one or both of the HFI PCIe cards. Firmware has be upgraded to 0x29.	10.2
134274 134301	HFI Host Fabric Software/ PSM2	Some MPI jobs are hanging including MPI_alltoall benchmarks that stall at 4K/8K messages sizes with higher PPN on OPA	10.2
134456	HFI Host Pre-boot Software	A page fault is detected by debug enabled BIOS versions when running UEFI driver 0x1C	10.2
134466	HFI Host Pre-boot Software	UEFI driver 0x1C causes a Red Screen on certain next gen servers running debug enabled BIOS versions.	10.2
134571	HFI Host Pre-boot Software	Version 10.1.1 0x1c signed UEFI option ROM fails secure boot	10.2
134598 134700	OPA Switch Firmware	Incorporate v0.45 and v0.47 of DC firmware into the Intel® Omni-Path Switch firmware build.	10.2



2.3 Open Issues

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

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

ID	Component	Description	Workaround
129563	HFI Host Fabric Software/ MPI	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 , $1 \leq n \leq 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to Appendix A in the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i> .
130336	HFI Host Fabric Software/ Tools	hfi1stats cannot be run at user level due to mount-point privileges	The administrator can provide "sudo access" to "hfi1stats" or provide root access to users.
131017	Open Software	Verbs ib_send_bw, ib_read_bw & ib_write_bw are not working with the -R option to use the RDMA CM API to create QP's and exch data.	The following combinations of client and server DO NOT allow RDMA CM connections: Client Server SLES* 12.1 SLES* 12.1 or SLES* 12.0 or (intermittent) RHEL* 7.2 SLES* 12.0 SLES* 12.1 or SLES* 12.0 Using SLES* clients results in "Unexpected CM event bl blka 0" errors. Additionally, there are long (5-10 sec) initial delays when using these combinations: Client Server SLES* 12.1 RHEL* 7.2 or 7.1 SLES* 12.0 RHEL* 7.2 or 7.1
131745	HFI Host Fabric Software/ MPI	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.	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	Open Software	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
133377	HFI Host Driver	irqbalance settings are not being honored correctly after a reboot.	Restart irqbalance after the hfi1 driver has loaded.
133408	HFI Hardware/PCIe card	In rare scenarios, loading the hfi1 host driver causes syslog messages including text such as "kernel:Uhhuh. NMI received for unknown reason 39 on CPU 0." On some systems, this may also be reported as "PCIe Fat" or as "CPU ERR2".	The error reporting is a transitory condition that arises on driver load. You can ignore the error and continue. Alternatively, ensure that the UEFI driver is being executed. This will perform the initialization before the OS boots and will avoid this error reporting. (For details, see the <i>Intel® Omni-Path Fabric Host Software User Guide</i> .)



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

ID	Component	Description	Workaround
133604	Open Software	Bonding driver incorrectly shows HW Address of IPoIB interfaces.	Use the <code>opainfo</code> command to retrieve the PortGUID and <code>ip addr show ib0</code> to get the correct 20-byte HW address of OPA network interface.
133633	Open Software	OpenMPI and Mvapich2 compiles fail to link properly when using the Intel Compilers	No work around available.
133707	Software Installation	Updating to the RHEL* 7.2 kernel for the CVE-2016-0728 update in OSs prior to 7.2 causes the Omni-Path installation to fail.	If you wish to apply the CVE-2016-0728 security patch, you must first upgrade your systems to RHEL* 7.2 OS and then apply the patch, otherwise you will end up with only the kernel update for 7.2 and the rest of the OS will think its 7.1.
134111	Software Configuration Management	On some older HFI and HFI-like cards, running <code>hfi1_eprom -V -c</code> to inquire the version of the AOC configuration file on the card may return an invalid version of "etnIRFWI".	Update cards to the latest version.
134124	Fabric Management Software/SM	HFI port stuck in INIT state due to SM failure to set pkeys.	Bounce the link.
134135 134429	HFI Host Fabric Software/ HFI Host Driver	When running communication-intensive workloads with 10KB MTU, it is possible to encounter node and/or job failures.	Change MTU default size in configuration file from 10KB MTU to 8KB MTU (<code>hfi1</code> driver module parameter <code>max_mtu=8192</code>). Refer to Section A.3, "Setting HFI1 Driver Parameters" in the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i> .
134268	HFI Host Pre-boot Software	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.	Use a BIOS that configures the HFI Expansion ROM BAR with an address that is 16MB aligned. Note that the main memory BAR of the device is 64MB in size and therefore requires 64MB alignment. A BIOS implementation that places the Expansion ROM BAR immediately after the main memory BAR will automatically provide this workaround. Many BIOS implementations have this property and automatically meet the workaround criteria.
134283	Software Installation/ Packaging	When downgrading on a SLES* 12.X system from IntelOPA version 10.2.X to a previous version, the install errors occur: "ERROR - Failed to install" and "error: Failed dependencies: libibmad5 is needed by opa-basic-tools..."	After you have uninstalled OPA version 10.2.X, install "libibmad5" from the distribution before you install the 10.1 or earlier version of OPA.
134471	HFI Host Pre-boot Software	The HFI UEFI driver cannot boot via PXE using Grub 2.	Use Elilo instead.
134772	Fabric Management Tools/Basic	<code>opatmmtool</code> will fail if provided a with a filename (full path) that is longer than 63 characters.	Rename the file to have a shorter name, or move the file so it has a shorter filepath.
134866	Fabric Management Software	<code>hostverify.sh</code> cannot properly detect if SRP is enabled on target node.	Modify <code>hostverify.sh</code> with the following: Replace <code>autostart=\$(systemctl is-enabled srpd 2>&1 grep enabled)</code> with <code>autostart=\$(lsmod grep ib_srp)</code>
134956	Fabric Management Tools/FastFabric	<code>ib0</code> fails to become ready on warm reboots.	If the HFI port transition to ACTIVE state takes an unusually long time and <code>ipoib</code> interface remains down, verify that the "NM_CONTROLLED=no" option is not present in the <code>ifcfg-ib*</code> files.



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

ID	Component	Description	Workaround
135000	Fabric Management Software	Fabric Manager configuration files which specify IncludeGroup fields with undefined or nonexistent device groups could cause Fabric Manager failure.	Remove undefined or nonexistent device groups from IncludeGroup fields in configuration file.
135068	HFI Host Pre-boot Software	When PXE booting using older versions of Grub 2 over Ethernet while the HFI UEFI driver is loaded, some servers will crash with an RSOD (Red Screen of Death).	Upgrade to the latest version of Grub 2.
135084	HFI Host Hardware	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 GUID reported by <code>opainfo</code> or other commands such as <code>hfi1_control -i</code> . Note that bit 39 of the PortGUID, the most significant bit will be set for the second HFI, and clear for the first HFI. Keep in mind that when issuing various Intel® Omni-Path CLI commands targeted at a specific HFI using the <code>-h</code> option, <code>-h 1</code> correlates to the device that is listed as <code>hfi1_0</code> . As a result, the issued command will affect the second HFI instance in cases where the second HFI port instance appears first.
135259	DC Link Software	On rare occasions, the HFI links will not come up after a reboot.	Reboot or bounce the link.
135873	Fabric Management Tools/FastFabric	<code>hostverify.sh</code> fails with RHEL6.7 due to the driver for Intel P-State driver not being the default <code>cpufreq</code> driver	Use the below steps to enable Intel P-State driver. 1. Check if any other <code>cpufreq</code> kernel driver is active. # <code>cpupower frequency-info -d</code> 2. Unload another <code>cpufreq</code> kernel driver (if any). # <code>rmmmod acpi_cpufreq</code> 3. Load <code>intel_pstate</code> driver # <code>modprobe intel_pstate</code> 4. Set <code>cpufreq</code> governor to 'performance' # <code>cpupower -c all frequency-set -g performance</code>
136215	Software Installation/ Packaging	For RHEL6.7 <code>opaconfig</code> command will not change the autostart settings for OPA service.	The autostart for <code>opa-scripts</code> is ON by default. You will not be able to turn OFF the autostart for this service using <code>opaconfig</code> command. Use the following to turn OPA service "on" and "off": Turn ON autostart for OPA service <code>chkconfig -level 2345 opa-scripts on</code> Turn OFF autostart for OPA service <code>chkconfig -level 2345 opa-scripts off</code>

§ §