

Intel[®] Omni-Path Fabric Software

Release Notes for 10.1

May 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 | 5 |
| 1.1 Introduction | 5 |
| 1.2 Audience | 5 |
| 1.3 Software License Agreement | 5 |
| 1.4 If You Need Help | 5 |
| 1.5 New Features in this Release | 5 |
| 1.6 Supported Features | 5 |
| 1.7 Release Packages | 6 |
| 1.8 Firmware Files | 7 |
| 1.9 Operating Systems | 7 |
| 1.10 Parallel File System | 8 |
| 1.11 Compilers | 8 |
| 1.11.1 MPI | 8 |
| 1.11.2 MVAPICH2 and Open MPI | 8 |
| 1.12 Hardware | 9 |
| 1.13 Installation Requirements | 10 |
| 1.13.1 Software and Firmware Requirements | 10 |
| 1.13.2 Installation Instructions | 10 |
| 1.14 Product Constraints | 10 |
| 1.15 Product Limitations | 11 |
| 1.16 Other Information | 11 |
| 1.17 Documentation | 11 |
| 2.0 Issues | 13 |
| 2.1 Introduction | 13 |
| 2.2 Resolved Issues | 13 |
| 2.3 Open Issues | 15 |
| Tables | |
| 1-1 Firmware Files | 7 |
| 1-2 Operating Systems Supported | 7 |
| 1-3 MPI Compilers | 8 |
| 1-4 MVAPICH2 and Open MPI | 9 |
| 1-5 Hardware Supported | 9 |
| 1-6 Related Documentation for this Release | 12 |
| 2-1 Resolved Issues | 13 |
| 2-2 Open Issues | 15 |



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 Features in this Release

- The following features are new for the 10.1 release: Additional operating systems supported. See [Table 1-2](#) for details
- Additional hardware supported. See [Table 1-5](#) for details
- Boot over Intel® Omni-Path using secure UEFI. This replaces the previously released Boot over Fabric.

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, Intel, and PGI compilers.
 - MVAPICH2 built for verbs and PSM2 using gcc, Intel, and PGI 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 *Intel® Omni-Path Fabric Software Installation Guide* 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 |
|---|----------------------|--------------|
| HFI1 UEFI Option ROM | HfiPcieGen3_0x1b.efi | 0x1b |
| UEFI UNDI | HfiUndiStub_0x1b.rom | 0x1b |
| HFI1 SMBus Microcontroller Firmware (Thermal Monitor) | hfi1_smbus.fw | 10.1.0.0.101 |

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*) 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:

- IBM* General Parallel File System (GPFS) version 4.0.1
 - RHEL* 7.1 and 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 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

| 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 |
| (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 |

See ID [132345](#) for expected behavior that may be observed 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 |

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 10.1 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 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 autoloader? 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 UNDI 0x1b firmware resident in this release is not compatible in server configurations that include dual HFI PCIe cards.
- 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.
- PA Failover should **not** be enabled with FMs running on differing software versions. PA Failover is enabled via configuration:<PM>/<ImageUpdateInterval> > 0

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
<http://www.intel.com/SUPPORT/NETWORK/OMNI-SWTCH100/SB/CS-035856.HTM>
- Intel® Omni-Path Host Fabric Interface Installation, User, and Reference Guides (includes software documents)
<http://www.intel.com/support/network/omni-adptr/sb/CS-035857.htm>
- Drivers and Software (including Release Notes)
<https://downloadcenter.intel.com/>

Table 1-6. Related Documentation for this Release (Sheet 1 of 2)

| 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> |





Table 1-6. Related Documentation for this Release (Sheet 2 of 2)

| Document Title |
|--|
| <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>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 3)

| ID | Component | Description | Resolved in Release |
|--------|---|--|---------------------|
| 127520 | Host Fabric Software/ MPI/PSM2 | A segmentation fault may occur when running MVAPICH2 benchmark jobs utilizing greater than 72 mpi ranks per node. | 10.1 |
| 130791 | Host Fabric Software | Using opapacketcapture with large amounts of traffic can cause performance issues on the given host. | 10.1 |
| 131058 | HFI Host Fabric Software | Running a job with Intel® MPI v5.0.2.044 using the OFA fabric (indicated by -IB or I_MPI_FABRICS=ofa) fails to start any test when the number of process ranks requested is greater than 36. The following error is generated: [10] Abort: Got completion with error 12, vendor code=0, dest rank= at line 858 in file ../../../../src/mpid/ch3/channels/nemesis/netmod/ofa/ofa_poll.c | 10.1 |
| 131126 | Fabric Management Software | The Intel® Omni Path Fabric Manager service supports automatic restarting of failed FM instances. When a component of the FM fails prematurely, the service will automatically re-launch the component that has failed to preserve fabric stability and reliability. | |
| 131934 | Fabric Management Software | Redundant FMs may fail to synchronize and may log dbsync errors related to PM history redundancy. | 10.1 |
| 131947 | HFI Host Driver | There is no support for HFI1 counters via: /sys/class/infiniband/hfi1_0/ports/1/counters | 10.1 |
| 132127 | HFI Host Driver | Certain server BIOS does not recognize HFI and fails to appear as a "UEFI Network" in the standard UEFI boot option. | 10.1 |
| 132197 | HFI Host Fabric Software/mvapich2/ Verbs | The following mpi apps hang when run on mvapich2/verbs: <ul style="list-style-type: none"> • mpi_stress • osu_alltoall • osu_multi_lat The app osu_mbw_mr does not hang, but performance decreases as message size increases. | 10.1 |
| 132570 | Fabric Management Tools/FastFabric | Opareport -F should have a way to query for copper vs optical cables | 10.1 |
| 132579 | Fabric Management Tools/FastFabric | Need to improve tools for controlling LED beaconing on devices. CLI command opaLEDports has been developed to control LED beaconing on devices | 10.1 |



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

| ID | Component | Description | Resolved in Release |
|--------|---|--|---------------------|
| 132592 | Fabric Management Tools/FastFabric | The opaverifyhosts command does not include tests to determine if Intel P-State driver is disabled. | 10.1 |
| 132639 | Fabric Management Tools/FastFabric | opatop port counter details screen does not fit in standard 24-line display | 10.1 |
| 132606 | Fabric Management Tools/FastFabric | opareports does not support -m/--smadirect option | 10.1 |
| 132613 | Fabric Management Tools/FastFabric | opaextractstat2 does not allow extra arguments. Fix applies also for opaextractstat and opaextractperf. | 10.1 |
| 132617 | Fabric Management Software/Fabric Manager | Fabric Manager SwitchCascadeActivateEnable configuration setting can leave IPOIB in a non-working state for HFIs. This problem occurs when the Sm.SwitchCascadeActivateEnable setting in opafm.xml is enabled for all devices, and the FM is restarted. The restarted SM never issues a client reregistration of SA subscriptions request to endnodes. | 10.1 |
| 132623 | Fabric Management/PM | Due to how the Performance Agent (PA) builds PA responses, a port may exist in one PA query and not in another. This can depend upon several factors, for example: <ul style="list-style-type: none"> • PmPortGroup • Internal/External membership • A possible Performance Manager (PM) Query Failure and/or a possible PM topology failure) | 10.1 |
| 132643 | Host Software/Tools | hfi1_control -v Parameter does not work. | 10.1 |
| 132647 | Fabric Management Tools/FastFabric | opatop now shows congestion and BW on same screens | 10.1 |
| 132649 | Fabric Management Tools/FastFabric | Fast Fabric will copy the sample hostverify script to \$FF_HOSTVERIFY_DIR before asking user to edit the file. | 10.1 |
| 132650 | Fabric Management Tools/FastFabric | FastFabric utility opaverifyhosts has sample cpu tests which combines many different tests that are not available separately. | 10.1 |
| 132662 | HFI Host Driver | During EFI boot, an EFI driver on the adapter is launched that returns error messages related to invalid temperature sensor readings. | 10.1 |
| 132730 | Fabric Management Tools/FastFabric | opalinkanalysis command does not output identified port when identifying an incorrectly placed cable. | 10.1 |
| 132760 | DC Link Software | Link width downgrade to 0x width on TX [tx 0x0 (0), rx 0xf (4)] (AOC Link goes down.) | 10.1 |
| 132812 | Switch Fabric Management | A GUI-less Linux* install (default option) does not install the bc package which is needed for some Fabric Manager scripts. The issue manifests itself by failing the hfi_pkt test and the log shows: bc command was not found. | 10.1 |
| 132875 | Software Architecture | Added new error reporting fields/features that are available in certain PA queries and have included updates to the FastFabric, Fabric Manager, and Fabric Manager GUI to use these features. Queries Affected: > FocusPorts & VfFocusPorts Added a Flags field for each port in the link record indicating a port's status (OK, Ignored, PMA Failure, topology Failure) in the relevant Image. > GroupInfo, & Vfinfo Added 2 failed ports Counters to indicate the number of ports within each subgroup (Internal, Send, & Receive) that encountered either a PMA Failure or a topology Failure. | 10.1 |
| 133035 | HFI Host Pre-boot Software | HFI UEFI boot asserts. | 10.1 |
| 133080 | HFI Host Driver | Serial number reported using the hfi1_control -iv command is different than expected. | 10.1 |



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

| ID | Component | Description | Resolved in Release |
|--------|----------------------------|---|---------------------|
| 133114 | HFI Host Driver | In some cases while the link is down, the HFI port may report the following in /var/log/messages: 'handle_qsf_error_conditions: Cable RX channel 1/2 power too low' (AOC link is down.) | 10.1 |
| 133133 | HFI Host Pre-boot Software | PXE boot timeout results in system hang. | 10.1 |
| 133189 | Host Software/FastFabric | The help related to the opainfo "-o" switch is improved | 10.1 |
| 133238 | Host Software/Verbs | Kernel panic in RHEL 7.2 running IPoSTL traffic. | 10.1 |
| 133436 | HFI Host Pre-boot Software | CR has Bad Signature UEFI version 0x15 | 10.1 |

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 Intel® Omni-Path Fabric Performance Tuning User Guide. |
| 130233 | HFI Host Pre-boot Software | Intel® Omni-Path DHCP fails to ID system based on hardware ID. | Workaround not available. |
| 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. |
| 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". | Reboot the Fabric Manager node. |
| 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 |



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

| ID | Component | Description | Workaround |
|--------|---|---|---|
| 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 Intel® PSM2 Programmer's Guide, in the description of psm2_ep_open() function, unique_job_key parameter. | To resolve this error, set the Open MPI environment variable OMPI_MCA_orte_precondition_transports in every node running the job. The format is 16 digit hexadecimal characters separated by a dash. For example, OMPI_MCA_orte_precondition_transports=13241234acffedeb-abcdefabcdef1233 This key is used by the PSM2 library to uniquely identify each different job end point used on the fabric. If two MPI jobs are running on the same node sharing the same HFI and using PSM2, each one should have a different key. |
| 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. | Install infiniband-diags version 1.6.5. |
| 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 | Reload hfi1 driver or reboot node. |
| 133145 | HFI Host Driver | hfi1_eprom has options to write, read, and erase information from fabric EPROM. However, it does not have an option to display current versions of the OptionROM version, Bootcode version, and Configuration file version. It is important for users to have this option to verify successful update. | Refer to Technical Advisory TA0008 for details (IBL #566152). |
| 133241 | HFI Host Driver | Under some circumstances, an HFI adapter returns QSFP memory read errors and is unable to bring up the physical link. (AOC link does not come up.) | Unplug and replug the QSFP module. If failure persists, recovery requires node reboot. Some cases may also require removal of the AUX power from management circuit on HFI adapter. |
| 133269 | Host Software/Verbs | GFP_NOIO error message using GFP_KERNEL when using IPoIB on OPA in connected mode. | Workaround not available. |
| 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 Intel® Omni-Path Fabric Host Software User Guide.) |
| 133607 | Fabric Management Tools/FastFabric | opaverifyhosts fails to run on RHEL7.2 | The user can modify the /opt/opa/samples/hostverify.sh before running the pcicfg, pcispeed and hfi_pkt tests. In line 263 where it says: pci_id=\$(grep -o -m1 "24f[01]" pcicfg.stdout) Change grep to egrep. Same for lines 311 and 498. |
| 133676 | Fabric Management Software/Fabric Manager | Creating a virtual fabric with standby enabled causes opapaqueries to fail. | Do not enable standby on virtual fabrics. |




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

| ID | Component | Description | Workaround |
|--------|---|--|--|
| 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. | The VL level counts can be found using the hfi1stats command. |
| 133820 | HFI Host Pre-boot Software | The UEFI driver may not work on some systems. The system freezes when the UEFI driver starts. | There are no known workarounds. |
| 133918 | HFI Host Pre-boot Software | Booting over Omni-Path HFI occasionally fails with "buffer too small, need netbuf_size = 16777216". | Retry boot |
| 133987 | Fabric Management Software/Fabric Manager | SM crash with general protection fault during compounding of PM sweep image | Set ImagesPerComposite to 1 in opafm.xml |
| 134012 | HFI Host Pre-boot Software | The UEFI driver may not work on systems with the HFI card plugged into a PCIe slot that is not directly connected to the PCIe root complex. | <p>There are several possible workarounds:</p> <ul style="list-style-type: none"> • The best solution. <ul style="list-style-type: none"> — Move the card to a slot that is connected directly to the root complex. • Normal UEFI features, but at a reduced speed. <ul style="list-style-type: none"> — Limit the speed of the slot the card is plugged into to PCIe Gen2 or Gen1. • Same limitations as using a legacy BIOS boot. <ul style="list-style-type: none"> — Do not install the UEFI option ROM driver. |
| 134039 | Software Man Pages | Man pages contain extra Unicode and other characters that do not belong in the text. | Disregard Unicode information. |
| 134131 | HFI Host Pre-boot Software | Servers with dual HFI PCIe cards become unbootable after flashing the UNDI 0x1b firmware to one or both of the HFI PCIe cards. | Do not use the UNDI 0x1b firmware on servers with dual HFI PCIe cards |

§ §



- 1.0 Overview of the Release** 5
 - 1.1 Introduction..... 5
 - 1.2 Audience 5
 - 1.3 Software License Agreement 5
 - 1.4 If You Need Help 5
 - 1.5 New Features in this Release..... 5
 - 1.6 Supported Features 5
 - 1.7 Release Packages 6
 - 1.8 Firmware Files 7
 - 1.9 Operating Systems 7
 - 1.10 Parallel File System 8
 - 1.11 Compilers..... 8
 - 1.11.1 MPI 8
 - 1.11.2 MVAPICH2 and Open MPI 8
 - 1.12 Hardware 9
 - 1.13 Installation Requirements 10
 - 1.13.1 Software and Firmware Requirements..... 10
 - 1.13.2 Installation Instructions 10
 - 1.14 Product Constraints 10
 - 1.15 Product Limitations 11
 - 1.16 Other Information 11
 - 1.17 Documentation..... 11
- 2.0 Issues** 13
 - 2.1 Introduction..... 13
 - 2.2 Resolved Issues 13
 - 2.3 Open Issues 15



| | | |
|-----|---|----|
| 1-1 | Firmware Files | 7 |
| 1-2 | Operating Systems Supported..... | 7 |
| 1-3 | MPI Compilers | 8 |
| 1-4 | MVAPICH2 and Open MPI | 9 |
| 1-5 | Hardware Supported..... | 9 |
| 1-6 | Related Documentation for this Release..... | 12 |
| 2-1 | Resolved Issues | 13 |
| 2-2 | Open Issues | 15 |