



Intel[®] Omni-Path Fabric Software in SUSE* Linux* Enterprise Server 15 SP2

Release Notes

Rev. 2.0

August 2020



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

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

All 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.

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](https://www.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 © 2020, Intel Corporation. All rights reserved.



Contents

- 1.0 Overview of the Release..... 5**
 - 1.1 Audience..... 5
 - 1.2 Document Versions..... 5
 - 1.3 Software License Agreement..... 6
 - 1.4 If You Need Help..... 6
 - 1.5 Packages in This Release..... 6
 - 1.6 Supported Features..... 7
 - 1.7 Supported MPI Libraries..... 8
 - 1.8 Intel Hardware..... 8
 - 1.9 Intel® OPA Compatibility Matrix..... 9
 - 1.10 Installation Requirements..... 9
 - 1.10.1 Installation Instructions..... 9
 - 1.11 Product Constraints..... 11
 - 1.12 Product Limitations..... 11
- 2.0 Issues..... 12**
 - 2.1 Resolved Issues..... 12
 - 2.2 Open Issues..... 13
 - 2.2.1 Third Party Open Issues..... 15



Tables

1	Supported Document Versions.....	5
2	Supported Hardware.....	8
3	Intel® OPA Compatibility Matrix.....	9
4	Issues Resolved in this Release.....	12
5	Open Issues.....	13
6	Third Party Open Issues.....	15



1.0 Overview of the Release

These Release Notes are intended for Intel® Omni-Path Fabric software provided in box with the OS 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 Document Versions

Intel® Omni-Path publications are available at the following URLs. For documents compatible with this release, refer to the V10.9.3 documents listed in the table below.

<https://www.intel.com/content/www/us/en/design/products-and-solutions/networking-and-io/fabric-products/omni-path/downloads.html>

Click *Latest Release Library*. To view previous release versions, click **Filter and Group**, select the **Topics** filter, and choose the corresponding V10.9.3 release version.

Table 1. Supported Document Versions

Title	Doc. Number	Revision
Intel® Omni-Path Fabric Quick Start Guide	J57479	6.0
Intel® Omni-Path Fabric Setup Guide	J27600	10.0
Intel® Omni-Path Fabric Switches Hardware Installation Guide	H76456	10.0
Intel® Omni-Path Host Fabric Interface Installation Guide	H76466	5.0
Intel® Omni-Path Fabric Software Installation Guide	H76467	14.0
Intel® Omni-Path Fabric Switches GUI User Guide	H76457	10.0
Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide	H76458	10.0
Intel® Omni-Path Fabric Suite FastFabric User Guide	H76469	14.0
Intel® Omni-Path Fabric Suite Fabric Manager User Guide	H76468	13.0
Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide	H76471	13.0
Intel® Omni-Path Fabric Host Software User Guide	H76470	14.0
<i>continued...</i>		



Title	Doc. Number	Revision
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	12.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	16.0
Intel® Omni-Path IP and LNet Router Design Guide (Old title: Intel® Omni-Path IP and Storage Router Design Guide)	H99668	8.0
Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note	J57474	7.0
Intel® Omni-Path Management API Programmer's Guide	J68876	6.0
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	K59645	1.0
Intel® Omni-Path Fabric Manager GUI Release Notes	K61052	1.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)	K38337	1.0
Intel® Omni-Path Fabric Unified Extensible Firmware Interface (UEFI) Release Notes	K50782	1.0
Intel® Omni-Path Fabric Thermal Management Microchip (TMM) Release Notes	K38341	1.0
Intel® Omni-Path Fabric Firmware Tools Release Notes	K38342	1.0

1.3 Software License Agreement

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

1.4 If You Need Help

Technical support for Intel® Omni-Path products is available 24 hours a day, 365 days a year. Please contact Intel Customer Support or visit <http://www.intel.com/omnipath/support> for additional detail.

1.5 Packages in This Release

Intel® Omni-Path Software Packages
Packages created by Intel
opa-address-resolution-10.9.3-1.63.x86_64
opa-basic-tools-10.9.3-1.63.x86_64
opa-fastfabric-10.9.3-1.63.x86_64
opa-fm-10.9.3-1.63.x86_64
opa-fmgui-10.1.0.0.115-8.32.noarch
libopamgt-devel-10.9.3-1.63.x86_64
libpsm2-2-11.2.86-1.39.x86_64
<i>continued...</i>



Intel® Omni-Path Software Packages
Firmware binaries delivered by Intel
8051 firmware version 1.27.0
SBus Master firmware version 0x10130001
PCIe SerDes firmware version 0x4755
Fabric SerDes firmware version 0x1055
Packages used by Intel
kernel-firmware-20200107-3.12.1.noarch.rpm
libfabric1-1.9.0-1.61.x86_64
rdma-core-27.1-1.12.x86_64
openmpi2-2.1.6-10.19.x86_64
openmpi3-3.1.6-1.19.x86_64
mpitests-3.2-8.27.x86_64
mpitests-openmpi2-3.2-8.54.x86_64
mpitests-openmpi3-3.2-8.54.x86_64
mpitests-mvapich2-3.2-8.52.x86_64
mvapich2-psm2-2.3.3-1.34.x86_64
mvapich2-2.3.3-1.42.x86_64
mpitests-mvapich2-psm-3.2-8.35.x86_64

HFI Programmable Firmware

Click [here](#) and select *Latest release library* then expand *Software and Drivers* to download the compatible Intel programmable firmware for HFIs:

- Unified Extensible Firmware Interface (UEFI)
- Thermal Management Module (TMM)
- Firmware Tools

NOTE

Refer to the [Intel® OPA Compatibility Matrix](#) on page 9 for the firmware versions compatible with this release.

1.6 Supported Features

- The list of supported hardware is in [Table 2](#) on page 8.
- Product constraints are described in [Product Constraints](#) on page 11.
- Support for multiple virtual fabric security.
- UEFI, TMM, and Firmware Tools are now standalone rpms.
- Active Optical Cables. For details, see the Cable Matrix at: <https://www.intel.com/content/www/us/en/products/network-io/high-performance-fabrics/omni-path-cables.html>



- Support for active optical cables (AOC) on server platforms using integrated HFI for OPA (commonly known as "-F").
- Support for Power Class 2 active optical cables (AOC). See [Product Constraints](#) on page 11 for more information.
- 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.
- Support for OpenFabrics Interfaces (OFI), a framework that includes libraries (including libfabric) and applications used to export fabric communication services to applications.
- Support for NVMe over Fabric Protocol
- Virtual Fabric creation has been enhanced to better support advanced topologies, including the ability to place multicast traffic on a separate SL from unicast traffic. For details, see the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide*, section 2.

1.7 Supported MPI Libraries

The list below shows the different MPI libraries tested with SLES* 15 SP2 for Intel® Omni-Path Fabric Software.

- OpenMPI3-3.1.6
- MVAPICH2-2.3.3

1.8 Intel Hardware

The following table lists the Intel hardware supported in this release.

NOTE

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

Table 2. 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
2nd Generation Intel® Xeon® Scalable Processors	Cascade Lake CPU-based servers
Intel® Omni-Path Host Fabric Interface 100HFA016 (x16)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Host Fabric Interface 100HFA018 (x8)	Single Port Host Fabric Interface (HFI)
Intel® Omni-Path Switch 100SWE48Q	Managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48U	Externally-managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48UFH	Externally-managed 48-port Edge Switch, hot-swap power and fans
<i>continued...</i>	



Hardware	Description
Intel® Omni-Path Switch 100SWE48QH	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

1.9 Intel® OPA Compatibility Matrix

The following component versions are compatible with Intel® Omni-Path software in SLES* 15 SP2.

Table 3. Intel® OPA Compatibility Matrix

UEFI	TMM	Managed Switch	Externally-Managed Switch	FM GUI
1.9.2.0.3	10.9.0.0.208	10.8.2.0.6	10.8.2.0.6	10.10.0.0.437
1.9.0.1.0	10.8.0.0.214	10.8.0.0.186	10.8.0.0.186	10.9.3.0.22
1.8.1.0.0				10.9.2.1.1

1.10 Installation Requirements

This section provides instructions and information on installing the software.

1.10.1 Installation Instructions

Perform the steps in this section to install the default Intel® Omni-Path Software configuration.

Assumptions

- You are logged in as root or with root privileges.
- You have a list of IPv4 addresses and netmasks for each IPoIB interface you are going to set up.

References

- Refer to the *Intel® Omni-Path Fabric Software Installation Guide* for related software requirements and next steps.
- Refer to the *Intel® Omni-Path Fabric Switches Hardware Installation Guide* for related firmware requirements.

Procedures

Perform the following steps to install the default Intel® Omni-Path Software configuration using SLES* OS:



Step	Task/Prompt	Action
Install OPA-Basic Software		
1.	At the command prompt, enter the installation command for opa-basic-tools.	Type zypper install -y opa-basic-tools and press Enter .
2.	At the command prompt, reboot the server.	Type reboot and press Enter .
3.	Check your link using opainfo.	Type opainfo and press Enter . Example output: <pre>hfil_0:1 PortGID: 0xfe80000000000000:001175010163f931 PortState: Active LinkSpeed Act: 25Gb En: 25Gb LinkWidth Act: 4 En: 4 LinkWidthDnGrd ActTx: 4 Rx: 4 En: 3,4 LCRC Act: 14-bit En: 14-bit,16-bit, 48-bit Mgmt: True LID: 0x00000010-0x00000010 SM LID: 0x0000000c SL: 0 QSFP: Copper , 2m Hitachi Metals P/N IQSFP26C-20 Rev 03 Xmit Data: 0 MB Pkts: 251 Recv Data: 0 MB Pkts: 251 Link Quality: 5 (Excellent)</pre>
4.	Install the rdma-core rpm.	Type zypper install -y rdma-core and press Enter .
5.	On all compute nodes: install the PSM2 library.	Type zypper install -y libpsm2-2 and press Enter .
Install Intel® Omni-Path Fabric Suite Components on the Management Node		
6.	Install FastFabric.	Type zypper install -y opa-fastfabric and press Enter .
8.	Install Fabric Manager.	Type zypper install -y opa-fm and press Enter .
9.	Start the Fabric Manager.	Type systemctl start opafm and press Enter .
Set up IPoIB IPV4 Configuration		
10.	Manually edit or create the ifcfg-ibX file.	<i>Note:</i> Use the OS distribution-supplied instructions for setting up network interfaces. Type cat /etc/network/ifcfg-ib0 and press Enter . Example output: <pre>BOOTPROTO=static IPADDR=192.168.0.1 BROADCAST=192.168.0.255 NETWORK=192.168.0.0 NETMASK=255.255.255.0 STARTMODE=auto</pre>
11.	Bring up the ib0 interface.	Type ifup ib0 and press Enter .
12.	Perform a test ping.	Type ping <remote IPoIB address> and press Enter . For example: <pre>ping 10.228.200.161 PING 10.228.200.161 (10.228.200.161) 56(84) bytes of data: 64 bytes from 10.228.200.161: icmp_seq=1 ttl=64 time=0.863 ms</pre>
<i>continued...</i>		



Step	Task/Prompt	Action
(Optional) Install the Fabric Manager GUI		
13.	On one node in the fabric: install the Fabric Manager GUI.	<p><i>Note:</i> Intel recommends not to install the Fabric Manager GUI on the Management Node where the Fabric Manager is being used.</p> <p>Type <code>zypper install -y opa-fmgui</code> and press Enter.</p>
End Task		

1.11 Product Constraints

- Power class 2 AOC are supported. You must use UEFI version 1.5 or newer for proper operation. Servers using integrated HFI (-F) requires a specific BIOS level to support power class 2 AOC; contact your BIOS vendor for more information.
- The Secure boot feature in this release requires a minimum UEFI version of 1.9.2.

1.12 Product Limitations

This release has the following product limitations:

- Performance Administration (PA) Failover should not be enabled with FMs running on differing software versions.
To disable PA failover, edit the `/etc/opa-fm/opa_fm.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.



2.0 Issues

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

2.1 Resolved Issues

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

Table 4. Issues Resolved in this Release

ID	Description	Resolved in Release
139368	Some applications compiled with older compilers may use a personality bit that signifies that READ should imply EXECUTE permissions. 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.	SLES* 15 SP2
139743 143031 143115	Under a very heavy load through the IPoIB interface, the kernel warning NETDEV WATCHDOG: ib0 (hfi1): transmit queue 0 timed out, followed by the messages queue stopped 1, tx_head xxx, tx_tail xxx and transmit timeout: latency xxxx msecs may be seen. As of V10.9.0, this issue is no longer valid.	SLES* 15 SP2
145855	If the Admin VF is not running on VL0, the HSM may get into a state where it is unable to talk to the fabric. The sweep will log the following errors: <pre>opamgt ERROR: [<pid>] omgt_send_mad2: send failed; Invalid argument, agent id 2 MClass 0x81 method 0x1 attrId 0X11 attrM 0x0 WARN [topology]: SM: sm_send_stl_request_impl: Error Sending to Path:[1] Lid:[0xffffffff] [Can't find node in topology!]. AID:[NODEINFO] TID:[0x0000000000000031] Status:[OK (0x00000000)] WARN [topology]: SM: topology_main: TT: too many errors during sweep, will re-sweep in a few seconds rc: 108: unrecoverable error</pre>	SLES* 15 SP2
146456	In a fabric with only one Edge switch using the fat tree routing algorithm, a port can get stuck in the Init (LinkUp) state after the port is bounced.	SLES* 15 SP2
STL-46606 STL-47956 STL-48661	Bouncing a link or rebooting a device under certain fabric conditions may cause a switch in the fabric to be removed from the Fabric Manager's internal view of the topology leading to fabric disruptions and instability.	SLES* 15 SP2
STL-46790	In cases where GSI services are active and the FM is receiving capability change traps (common after node reboots), FM responsiveness may be impacted. This could result in data traffic disruption or unexpected FM failovers. GSI traffic would include the PM, SA, and DBSync (FM failover).	SLES* 15 SP2
STL-47546	When an ISL goes down in the middle of an FM sweep (due to a disruption in the fabric such as a reboot), the SA copy of topology becomes invalid when the Fattree routing algorithm is used. SA queries that use this topology (e.g., path record query) fail. <i>Note:</i> A path record query failure can be seen in FM log as "INVALID TOPOLOGY" messages. The issue will resolve after the FM's next successful sweep.	SLES* 15 SP2

continued...



ID	Description	Resolved in Release
STL-49210	For SLES* 15 and newer, when Accelerated IP (AIP) is disabled and IPoFabric is configured in datagram mode, no bulk traffic will occur if the configured ib0 MTU size is larger than what default (non-AIP) IPoFabric can support.	SLES* 15 SP2
STL-49734	Some applications compiled with older compilers may use a personality bit that signifies that READ should imply EXECUTE permissions. 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.	SLES* 15 SP2
STL-58309	It is possible for a memory region not to be freed due to the lack of rcu synchronization for the memory region's refcount. This has been observed when executing the <code>mmvdisk</code> while running GPFS version RAID version 5.0.2-3.0.1.	SLES* 15 SP2

2.2 Open Issues

The following table lists the open issues for this release.

Table 5. Open Issues

ID	Description	Workaround
139613	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.	Reconfigure the system from Legacy OS boot mode to UEFI boot mode.
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, use <pre>setpci -d 8086:24f0 ECAP_AER +8.1=00100000:00100000</pre> For integrated HFIs, use <pre>setpci -d 8086:24f1 ECAP_AER +8.1=00100000:00100000</pre>
143296	When <code>irqbalance</code> uses the argument <code>--hintpolicy=exact</code> , 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 <code>irqbalance</code> using the argument <code>--hintpolicy=exact</code> .
145623	For systems running on SLES* 15, there is a known issue with <code>irqbalance</code> .	Contact Intel Customer Support for more information.
STL-47003	Due to a kernel setting, <code>hfiDiags</code> cannot work while the HFI driver is loaded. The tool and driver are mutually exclusive.	Boot the kernel with <code>iomem=relaxed</code> . This turns off the resource exclusivity check.
STL-48921	If you attempt to unload or restart the <code>hfi1</code> driver while the <code>ib_ipoib</code> driver is loaded, it may fail with a message similar to this: <code>modprobe: FATAL: Module hfi1 is in use</code>	To avoid this issue, unload the <code>ib_ipoib</code> driver before unloading or restarting the <code>hfi1</code> driver.

continued...



ID	Description	Workaround
STL-49313	On OS versions (SLES* 15 and newer) where it is available, configuring IPoIB module parameter <code>ipoib_enhanced=0</code> disables AIP on same node.	To run with AIP, ensure enhanced IPoIB is enabled (<code>ipoib_enhanced=1</code>).
STL-49732	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.	Reconfigure the system from Legacy OS boot mode to UEFI boot mode.
STL-55021	Due to changes in the kernel, a user space application cannot access a hardware resource if that resource is being used by a kernel driver. The result is that the <code>hfi1_eprom</code> cannot access the EEPROM on an HFI when <code>hfi1</code> kernel driver is using the device.	Perform the following workaround: <ol style="list-style-type: none"> 1. Unload <code>hfi1</code> kernel driver: <code>rmmod hfi1</code> 2. Run <code>hfi1_eprom</code> commands. 3. Reload <code>hfi1</code> kernel driver: <code>modprobe hfi1</code>
STL-56557	The primary IPoIB network device associated with any RDMA device may fail to join certain multicast groups. This can prevent IPv6 neighbor discovery and possibly other network ULPs from working correctly. Note that the IPv4 broadcast group is not affected as the IPoIB network device handles joining that multicast group directly. This problem does not affect IPoIB child network devices.	Perform one of the following workarounds: <ul style="list-style-type: none"> • Do not use IPv6 over a parent IPoIB device. • Only load the IPoIB module after the HFI link is ACTIVE with a valid pkey assigned. • Bounce (unload, then load) the IPoIB module anytime a pkey change occurs.
STL-58420	Back-to-back runs of <code>opatmtool</code> or <code>opacapture</code> will corrupt the module use count, preventing an unload of the <code>hfi1</code> modules.	No workaround but the impact occurs only when the TMM firmware needs to be updated or when a capture needs to be done. If a firmware update or <code>opacapture</code> is required, power-cycle the node afterwards. Refer to https://bugzilla.suse.com/show_bug.cgi?id=1174407 .
STL-60132	The following panic can be seen during shutdown: <pre>[188.570075] BUG: unable to handle kernel NULL pointer dereference at 0000000000000102 [188.570114] IP: [ffffffffff810a65f2] __queue_work +0x32/0x3c0</pre>	None. Refer to https://bugzilla.suse.com/show_bug.cgi?id=1174409 .
STL-60473	While running PSM2 MPI jobs, the following message can be seen in <code>dmesg</code> : <pre>WARNING: CPU: 25 PID: 32006 at lib/list_debug.c:29 __list_add+0x65/0xc0</pre>	None. Refer to https://bugzilla.suse.com/show_bug.cgi?id=1174411 .
STL-60722	The completion ring for a kernel completion queue is leaked when the CQ is destroyed.	None. Refer to https://bugzilla.suse.com/show_bug.cgi?id=1173625 .
STL-60776	While running a verbs upper-layer protocol (ULP), incorrect reporting of queue pair (QP) errors associated with receiving messages may cause a QP to enter the error state. The observed symptom may be reported as an error by the ULP or the system may hang if the ULP does not report the error associated with the QP state transition to error. Other symptoms are possible and ULP-dependent.	None. Refer to https://bugzilla.suse.com/show_bug.cgi?id=1174770 .



2.2.1 Third Party Open Issues

The following table lists the third party open issues for this release.

Table 6. Third Party Open Issues

ID	Description	Workaround
129563 (STL-47095)	Memory allocation errors with MVAPICH2-2.1/Verbs.	<p><i>Note:</i> To avoid this issue, use MPIs over PSM.</p> <p>If you are using MPIs over verbs, the following workaround is required:</p> <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>, HF11 Driver Module Parameters chapter.
STL-47571	<p>When trying to run an MPI/PSM job with more MPI ranks than CPU cores (oversubscribing), the job may fail with the following error message:</p> <pre data-bbox="386 953 894 1031">hfi_userinit: assign_context command failed: Device or resource busy PSM2 can't open hfi unit: -1 (err=23)</pre>	<p>Set <code>PSM2_MULTI_EP=0</code> (user environment variable) before or during job launch.</p> <p>For details, see the <i>Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>.</p>
STL-60633	<p>When you install <code>gpfs.callhome-ecc-client-5.0.5-0</code> under SLES using zypper, you may receive an error code if there are no existing credentials.</p>	<p>Ignore the error code.</p> <p><i>Note:</i> To avoid this error, you can touch an empty file <code>/usr/lpp/mmfs/bin/ecc/ECCBase/com.ibm.ws.webservices.thinclient_7.0.0.jar</code> to force zypper to exit with a success code. For example, run the following command before you install <code>gpfs.callhome-ecc-client-5.0.5-0</code>:</p> <pre data-bbox="917 1272 1417 1308">touch /usr/lpp/mmfs/bin/ecc/ECCBase/ com.ibm.ws.webservices.thinclient_7.0.0.jar</pre>
STL-60746	<p>When running verbs on openmpi with the ob1 PML, a crash occurs during MPI Finalize.</p>	<p>Verbs is not a recommended path for MPI traffic. Use PSM2.</p>