



Intel[®] Omni-Path Fabric Software in SUSE* Linux* Enterprise Server 12 SP5

Release Notes

Rev. 1.0

December 2019



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.

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 © 2019, 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 Best Practices.....	9
1.10.2 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 URL:

<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 release version.

For documents compatible with this release, refer to the V10.8.0 revisions listed in the table below.

Table 1. Supported Document Versions

Title	Doc. Number	Revision
<i>Intel® Omni-Path Fabric Quick Start Guide</i>	J57479	5.0
<i>Intel® Omni-Path Fabric Setup Guide</i>	J27600	9.0
<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	10.0
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	H76457	10.0
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	H76458	10.0
<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>	H76469	10.0
<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>	H76468	10.0
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>	H76471	10.0
<i>continued...</i>		



Title	Doc. Number	Revision
Intel® Omni-Path Fabric Host Software User Guide	H76470	10.0
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	10.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	12.0
Intel® Omni-Path IP and LNet Router Design Guide (Old title: Intel® Omni-Path IP and Storage Router Design Guide)	H99668	7.0
Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note	J57474	5.0
Intel® Omni-Path Management API Programmer's Guide	J68876	4.0
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	K21143	1.0
Intel® Omni-Path Fabric Manager GUI Release Notes	K21144	1.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally- managed switches)	K21142	1.0
Intel® Omni-Path Fabric Unified Extensible Firmware Interface (UEFI) Release Notes	K21145	1.0
Intel® Omni-Path Fabric Thermal Management Microchip (TMM) Release Notes	K21147	1.0
Intel® Omni-Path Fabric Firmware Tools Release Notes	K21148	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-10.8.0-1.30.x86_64
opa-basic-tools-10.8.0-1.30.x86_64
opa-fastfabric-10.8.0-1.30.x86_64
opa-fm-10.8.0-1.30.x86_64
opa-fmgui-10.1.0.0.115-1.35.noarch
opa-libopamgt-10.8.0-1.30
libpsm2-11.2.78-1.31.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-20190618-5.11.1.noarch
libfabric1-1.7.1-1.31.x86_64
rdma-core-22.3-2.21.x86_64
openmpi-1.10.7-8.12.x86_64
mpitests-3.2-13.12.x86_64
mpitests-openmpi-3.2-13.15.x86_64
mpitests-mvapich2-3.2-13.14.x86_64
mvapich2-psm2-2.3.1-4.8.x86_64
mvapich2-2.3.1-4.9.x86_64
mpitests-mvapich2-psm2-3.2-13.13.x86_64

HFI Programmable Firmware

To download Intel programmable firmware for HFIs, refer to the following:

- [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* 12 SP5 for Intel® Omni-Path Fabric Software. Note that these MPIs are designed to run with PSM2 by default.

- OpenMPI 2.1.2
- MVAPICH2-2.3B

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® Xeon Phi™ x200 Product Family	Knights Landing CPU-based servers
Intel® Xeon Phi™ 72x5 Processor Family	Knights Mill 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
<i>continued...</i>	



Hardware	Description
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

1.9 Intel® OPA Compatibility Matrix

The following component versions are compatible with Intel® Omni-Path software in SLES* 12 SP5.

Table 3. Intel® OPA Compatibility Matrix

UEFI	TMM	Managed Switch	Externally-Managed Switch	FM GUI
1.9.0.1.0	10.9.0.0.208	10.8.0.0.186	10.8.0.0.186	10.9.2.0.7
1.8.1.0.0	10.8.0.0.214	10.8.0.0.186	10.8.0.0.186	10.8.0.0.206
1.7.2.0.0	10.7.0.0.3	10.7.0.0.146	10.7.0.0.144	10.7.0.0.145

1.10 Installation Requirements

This section provides instructions and information on installing the software.

1.10.1 Best Practices

- Intel recommends that users update to the latest versions of Intel® Omni-Path firmware and software to obtain the most recent functional and security updates.
- To improve security, the administrator should log out users and disable multi-user logins prior to performing provisioning and similar tasks.

1.10.2 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: AOC , 5m FINISAR CORP P/N FCBN425QB1C05 Rev A 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 </pre>

continued...



Step	Task/Prompt	Action
		<pre>STARTMODE=auto IPOIB_MODE='connected' MTU=65520</pre> <p>NOTE: To configure datagram mode for AIP, change IPOIB_MODE=datagram and remove (comment out) MTU= of the ifcfg-ib0 file. Further details can be found in the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>.</p>
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>
(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> Type zypper install -y opa-fmgui and press Enter .
End Task		

1.11 Product Constraints

- The minimum firmware version for Intel® Omni-Path Host Fabric Interface Silicon 100 Series Switch ASIC is 10.7.
- 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.
- Before using Intel® MPI provided by Intel® Parallel Studio 2018, please verify its OS compatibility.

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
132207	Kernel crash caused by the <code>ib_srpt</code> module.	SLES* 12 SP5
141793	Use of a static buffer could produce an incorrect device name (<code>hfi1_x</code>) in <code>dmesg</code> logging.	SLES* 12 SP5
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.	SLES* 12 SP5
143449	PM will scroll <code>LQI=0</code> and <code>Integrity Exceeded Threshold</code> 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 <code>opafm.xml</code> configuration file.	SLES* 12 SP5
144165	Nodes unable to ping on IPoIB. <i>Note:</i> This issue occurs when a host port disappears and reappears from the FM's topology (usually due discovery timeout or major fabric disruption), while the port remains ACTIVE the entire time. This results in the host port not being a member of the IP multicast groups. The primary symptom is the inability to resolve IP addresses via ARP.	SLES* 12 SP5
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* 12 SP5
STL-46935 (141273)	The in-distro version of <code>perftests</code> has bugs.	SLES* 12 SP5
STL-47901	The output of the module parameter <code>num_user_contexts</code> has been changed. The value of the module parameter <code>num_user_context</code> used at driver init time determines the number of receive contexts reserved for PSM. The default value of -1 caused the driver init sequence to determine how many CPUs are available, and assigned the context count to that value. <i>It would then update the module parameter <code>num_user_contexts</code> with that value.</i> This incorrect behavior has been fixed. <code>num_user_contexts</code> is no longer updated.	SLES* 12 SP5
STL-48734	Some applications compiled with older compilers may use a personality bit that signifies that READ should imply EXECUTE permissions. To improve system security, the <code>hfi1</code> driver does not allow execute permissions on PSM memory maps. Therefore, applications that use READ implies EXECUTE will fail to run.	SLES* 12 SP5



2.2 Open Issues

The following table lists the open issues for this release.

Table 5. Open Issues

ID	Description	Workaround
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.
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.
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 data-bbox="922 821 1211 863">setpci -d 8086:24f0 ECAP_AER +8.1=00100000:00100000</pre> For integrated HFIs, use <pre data-bbox="922 940 1211 982">setpci -d 8086:24f1 ECAP_AER +8.1=00100000:00100000</pre>
143031	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, tx_head xxx, tx_tail xxx</code> and <code>transmit timeout: latency xxxx msecs</code> may be seen.	The default <code>ib_ipoib</code> module option <code>send_queue_size</code> is 256. To avoid TX timeouts, Intel recommends that you modify <code>send_queue_size</code> to 2048 or higher, up to a maximum of 8192.
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> .
146456	In a fabric with only one Edge switch using the fat tree routing algorithm, a port can get stuck in the <code>Init (LinkUp)</code> state after the port is bounced.	Perform one of the following workarounds: <ul style="list-style-type: none"> Use the shortest path routing algorithm. Restart the Fabric Manager.
STL-46193	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>	Install parallel gzip (<code>pigz</code>) and make it the default compression tool prior to upgrading the software.
STL-46550	Running workloads with more than 78 ranks with the Open MPI OFI MTL over OFI Verbs;OFI_RXM provider may result in a hang with message sizes larger than 65 KB.	To avoid this issue, increase the TX/RX size by setting the following environment variables on the <code>mpirun</code> command line: <ul style="list-style-type: none"> <code>-x FI_VERBS_RX_SIZE=2000</code> <code>-x FI_VERBS_TX_SIZE=2000</code>

continued...



ID	Description	Workaround
STL-46790	<p>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.</p> <p>GSI traffic would include the PM, SA, and DBSync (FM failover).</p>	<p>Apply the Github patch: https://github.com/intel/opa-fm/commit/86d2a898652529234a368706e711aabee438b9456</p> <p>Or, apply the following manual workaround:</p> <ul style="list-style-type: none"> Disable trap handling via the <IgnoreTraps> FM XML config option. Note that this prevents event-driven sweeps, so the FM will only handle fabric changes based on the configured regular sweep interval (default is every 5 minutes). Disable the PM. This may mitigate the most likely cause of failure, given the higher volume of PM traffic. Note that this disables performance counter collection, and may not completely mitigate the issue.
STL-47546	<p>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.</p> <p><i>Note:</i> A path record query failure can be seen in FM log as "INVALID TOPOLOGY" messages.</p> <p>The issue will resolve after the FM's next successful sweep.</p>	<p>To avoid this issue, use the Shortestpath routing algorithm instead of Fattree.</p>
STL-49732	<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>
STL-47578 STL-55021	<p>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.</p>	<p>Perform the following workaround:</p> <ol style="list-style-type: none"> Unload <code>hfi1</code> kernel driver: <code>rmmmod hfi1</code> Run <code>hfi1_eprom</code> commands. Reload <code>hfi1</code> kernel driver: <code>modprobe hfi1</code>



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
STL-47095 (129563)	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 (not recommended with OPA), 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>, HFI1 Driver Module Parameters chapter.