



Intel® Omni-Path Fabric Software in Red Hat* Enterprise Linux* 8.3

Release Notes

Rev. 1.0

November 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 product plans and roadmaps are subject to change without notice.

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 may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. 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.....	9
1.8 Intel Hardware.....	9
1.9 Intel® OPA Compatibility Matrix.....	10
1.10 Installation Requirements.....	10
1.10.1 Best Practices.....	10
1.10.2 Installation Instructions.....	11
1.11 Product Constraints.....	12
1.12 Product Limitations.....	12
2.0 Issues	13
2.1 Resolved Issues.....	13
2.2 Open Issues.....	13
2.2.1 Third Party Open Issues.....	15



Tables

1	Supported Document Versions.....	5
2	Supported Hardware.....	9
3	Intel® OPA Compatibility Matrix.....	10
4	Issues Resolved in this Release.....	13
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 V10.10.1 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 V10.10.1 release version.

Table 1. Supported Document Versions

Title	Doc. Number	Revision
<i>Intel® Omni-Path Fabric Quick Start Guide</i>	J57479	8.0
<i>Intel® Omni-Path Fabric Setup Guide</i>	J27600	12.0
<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>	H76456	12.0
<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>	H76466	7.0
<i>Intel® Omni-Path Fabric Software Installation Guide</i>	H76467	16.0
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	H76457	12.0
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	H76458	12.0
<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i>	H76469	16.0
<i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i>	H76468	15.0
<i>Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide</i>	H76471	15.0
<i>continued...</i>		

Title	Doc. Number	Revision
Intel® Omni-Path Fabric Host Software User Guide	H76470	16.0
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	14.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	18.0
Intel® Omni-Path IP and LNet Router Design Guide (Old title: Intel® Omni-Path IP and Storage Router Design Guide)	H99668	10.0
Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note	J57474	9.0
Intel® Omni-Path Management API Programmer's Guide	J68876	8.0
Configuring Non-Volatile Memory Express* (NVMe*) over Fabrics on Intel® Omni-Path Architecture Application Note	J78967	2.0
Intel® Omni-Path Fabric Software Release Notes	K82053	1.0
Intel® Omni-Path Fabric Manager GUI Release Notes	K69636	2.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)	K82050	2.0
Intel® Omni-Path Fabric Unified Extensible Firmware Interface (UEFI) Release Notes	K50782	2.0
Intel® Omni-Path Fabric Thermal Management Microchip (TMM) Release Notes	K38341	2.0
Intel® Omni-Path Fabric Firmware Tools Release Notes	K50784	2.0

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-10.10.1.0.35-1.el8.x86_64
opa-basic-tools-10.10.1.0.35-1.el8.x86_64
opa-fastfabric-10.10.1.0.35-1.el8.x86_64
opa-fm-10.10.1.0.35-1.el8.x86_64
opa-libopamgt-10.10.1.0.35-1.el8.x86_64
libpsm2-11.2.91-1.el8.x86_64
Firmware binaries delivered by Intel
<i>continued...</i>

Intel® Omni-Path Software Packages
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
linux-firmware-20200619-99.git3890db36.el8.noarch
libfabric-1.10.0-1.el8.x86_64
rdma-core-29.0-3.el8.x86_64 (libhfi1)
openmpi-4.0.3-3.el8.x86_64
mpitests-openmpi-5.6.2-1.el8.x86_64
mpitests-mvapich2-5.6.2-1.el8.x86_64
mpitests-mvapich2-psm2-5.6.2-1.el8.x86_64
mvapich2-2.3.3-1.el8.x86_64
mvapich2-psm2-2.3.3-1.el8.x86_64

HFI Programmable Firmware

Click [here](#) and select "Latest release library" to download 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 10 for the firmware versions compatible with this release.

1.6 Supported Features

The following features have been deprecated or removed:

- CongDiscards counters are no longer thresholded in the default configuration file (`/etc/opa/opamon.conf`). As a result, they are no longer reported by `opareport -o errors`, `opaextractstat`, `opaextractstat2`, `opaextractbadlinks`, and `opafabricanalysis` when using the default configuration file.
- Support for SHMEM has been removed in this release.
- Native verbs support in Open MPI (openib BTL) is no longer maintained. It has been removed from the Open MPI build in IFS.

Product Improvements

- Top 10 Class Storage Performance Tuning

Intel® Omni-Path drivers have been further hardened and tuned to meet the needs of high-end HPC filesystems. Refer to *Intel® Omni-Path Fabric Performance Tuning User Guide*.

- Installation Enhancements

- Meta-packages are available to facilitate yum/zypper-based IFS installation. Each meta-package corresponds to one component in the IFS INSTALL script. Install, upgrade, or downgrade using a meta-package will have the same effect as using the INSTALL script.

- Alias packages are introduced to further facilitate IFS installation. Each alias package represents a typical HPC node (management node, compute node, and service node) IFS installation. Install using an alias package will install a full set of rpms for the corresponding HPC node.

- Performance Administration (PA) History Scalability

The performance of the PA when monitoring large clusters and/or clusters with longer retention of performance history has been improved. This benefit can be seen when using `opareport` (especially with `--begin` and `--end` options), `opatop`, the FM GUI, the OPA Management API, and many other tools. In addition, the following PA history analysis capabilities have been enhanced:

- `opatop` now permits direct access to historic PM images by inputting a specific time. When a long history is retained, this capability is more effective than simply stepping forward or backward among the images. Refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide*, Navigating PM Sweeps section.
- FastFabric now includes a more effective way to extract all of the counters for the fabric and report the specific link and neighbor associated with them.

The new `opaextractperf2` script permits easy analysis of information, in a spreadsheet or through other customer scripts, for the fabric. `opareport` focus options and `--begin` and `--end` options may be used to obtain the subset of needed information.

The new `opamergeperf2` script combines two previous `opaextractperf2` runs with running counters into a csv file showing the delta between the two points in time.

Refer to the respective sections in the *Intel® Omni-Path Fabric Suite FastFabric User Guide*.

- Improved Fabric Debug/Diagnostics

- The FM will now log the *LinkDownReason* in Node Disappearance log messages.
- `opareport` now includes the link down reason logs in `-o comps -d5` reports and snapshots. `opareport` also allows focus on link down reason through the `-F ldr:` option.
- `opareport` now includes the *LinkCRCMode* in `-o comps`.
- `opareport -o links` now includes port counters.
- PM logging for fabric error events is now improved to include the error information provided by the hardware. On security, routing, and some signal integrity error events, the PM will get and clear hardware *ErrorInfo* and log the



information. FM configuration allows a threshold for the number of such log events per sweep. This log can provide additional information such as which PKey was attempting to be used or which DLID was unable to be routed. This information can aid sysadmins to debug faulty applications or configurations.

- The `opaextractbadlinks`, `opaextractlids`, and `opaextractmissinglinks` tools have been enhanced to permit use of `opareport -F` focus options to limit the scope of the report

Refer to the *Intel® Omni-Path Fabric Suite FastFabric User Guide* for more information.

- The *Intel® Omni-Path Management API Programmer’s Guide* documents the performance characteristics of PA query functions. This can permit developers of PA applications to design their PA usage for improved scalability and performance.
- The Intel® Omni-Path Fabric Manager can now be more tolerant of non-responsive SMAs on hosts whose port remains active. This can permit applications to soak the CPU or memory of a server without the host being removed from the fabric by the FM. This is controlled via the `NonRespMaxCount` FM configuration option. Refer to the *Intel® Omni-Path Fabric Suite Fabric Manager User Guide* for more information.
- A new `PSM2_AVX512` environment variable has been added that allows the user to control whether PSM2 uses AVX512 code for performing larger memory copies. Disabling AVX512 in PSM2 may improve performance for some applications and traffic patterns. Refer to the *Intel® Performance Scaled Messaging 2 (PSM2) Programmer’s Guide*.

1.7 Supported MPI Libraries

The list below shows the MPI libraries supported for Intel® Omni-Path Fabric Software.

- OpenMPI 4.0.3
- MVAPICH2-2.3

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)
<i>continued...</i>	

Hardware	Description
Intel® Omni-Path Switch 100SWE48Q	Managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48U	Externally-managed 48-port Edge Switch
Intel® Omni-Path Switch 100SWE48UFH	Externally-managed 48-port Edge Switch, hot-swap power and fans
Intel® Omni-Path Switch 100SWE48QFH	Managed 48-port Edge Switch, hot-swap power and fans
Intel® Omni-Path Switch 100SWE24Q	Managed 24-port Edge Switch
Intel® Omni-Path Switch 100SWE24U	Externally-managed 24-port Edge Switch
Intel® Omni-Path Director Class Switch 100SWD24	Director Class Switch 100 Series, up to 768 ports
Intel® Omni-Path Director Class Switch 100SWD06	Director Class Switch 100 Series, up to 192 ports

1.9 Intel® OPA Compatibility Matrix

The following component versions are compatible with Intel® Omni-Path software in RHEL* 8.3.

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	10.10.3.1.1
1.9.0.1.0				

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.
- To improve security, Intel recommends updating the default HTTPS certificate. Refer to the *Intel® Omni-Path Fabric Switches GUI User Guide*, Updating the Certificate for details.
- To improve security, Intel recommends configuring the `MgmtAllowed` setting and consider limiting access to port configuration changes by limiting access to Userspace Management Datagrams (UMADs). Refer to the *Intel® Omni-Path Fabric Software Installation Guide*, About User Queries Settings for more information.



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.
- RHEL* packages are available in a yum repository.

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 RHEL* OS:

Step	Task/Prompt	Action
Install OPA-Basic Software		
1.	At the command prompt, enter the installation command for <code>opa-basic-tools</code> .	Type <code>yum install -y opa-basic-tools</code> and press Enter .
2.	At the command prompt, reboot the server.	Type <code>reboot</code> and press Enter .
3.	Check your link using <code>opainfo</code> .	Type <code>opainfo</code> and press Enter . Example output: <pre>hfi1_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 <code>rdma-core</code> rpm.	Type <code>yum install -y rdma-core</code> and press Enter .
5.	On all compute nodes: install the PSM2 library.	Type <code>yum install -y libpsm2</code> and press Enter .
Install Intel® Omni-Path Fabric Suite Components on the Management Node		
6.	Install FastFabric.	Type <code>yum install -y opa-fastfabric</code> and press Enter .
7.	Install the <code>opa-address-resolution</code> rpm on all nodes.	Type <code>yum install -y opa-address-resolution</code> and press Enter .
<i>continued...</i>		

Step	Task/Prompt	Action
8.	Install Fabric Manager.	Type <code>yum install -y opa-fm</code> and press Enter .
9.	Start the Fabric Manager.	Type <code>systemctl start opafm</code> and press Enter .
Set up IPoIB IPV4 Configuration		
10.	Manually edit or create the <code>ifcfg-ibX</code> file.	<p><i>Note:</i> Use the OS distribution-supplied instructions for setting up network interfaces.</p> <p>Type <code>cat /etc/network-scripts/ifcfg-ib0</code> and press Enter.</p> <p>Example output:</p> <pre>DEVICE=ib0 TYPE=infiniband BOOTPROTO=static IPADDR=10.228.200.173 BROADCAST=10.228.203.255 NETWORK=10.228.200.0 NETMASK=255.255.252.0 ONBOOT=yes</pre>
11.	Bring up the <code>ib0</code> interface.	Type <code>ifup ib0</code> and press Enter .
12.	Perform a test ping.	<p>Type <code>ping <remote IPoIB address></code> and press Enter.</p> <p>For example:</p> <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>
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.

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/opafm.xml` file and in the `<Pm>` section, change `<ImageUpdateInterval>` to 0.
- Enabling UEFI Optimized Boot on some platforms can prevent the HFI UEFI driver from loading during boot. To prevent this, do not enable UEFI Optimized Boot.

2.0 Issues

This section lists the 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
STL-58420	Back-to-back runs of opacapture will corrupt the module use count, preventing an unload of the hfi1 modules.	RHEL* 8.3
STL-60132	The following panic can be seen during shutdown: [188.570075] BUG: unable to handle kernel NULL pointer dereference at 0000000000000102 [188.570114] IP: [ffffffff810a65f2] __queue_work+0x32/0x3c0	RHEL* 8.3
STL-60473	While running PSM2 MPI jobs, the following message can be seen in dmesg: WARNING: CPU: 25 PID: 32006 at lib/list_debug.c:29 __list_add+0x65/0xc0	RHEL* 8.3
STL-60536	When TID RDMA is enabled with 256 nodes or greater, it is possible that a crash or failure may occur under stressed conditions.	RHEL* 8.3
STL-60722	The completion ring for a kernel completion queue is leaked when the CQ is destroyed.	RHEL* 8.3

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.

continued...

ID	Description	Workaround
		<ul style="list-style-type: none"> For discrete HFI ASICs (e.g., CHF PCIe card), use <pre>setpci -d 8086:24f0 ECAP_AER +8.1=00100000:00100000</pre> For integrated HFIs (e.g., KNL-F and SKX-F), use <pre>setpci -d 8086:24f1 ECAP_AER +8.1=00100000:00100000</pre>
STL-56557	<p>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.</p> <p>Note that the IPv4 broadcast group is not affected as the IPoIB network device handles joining that multicast group directly.</p> <p>This problem does not affect IPoIB child network devices.</p>	<p>Perform one of the following workarounds:</p> <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-57127	<p>Restarting the irqbalance service while the HFI1 driver is loaded may cause improper or inconsistent distribution of interrupts, which may result in low network transfer performance.</p>	<p>Irqbalance must be stopped before starting the HFI1 module. Perform the following command sequence:</p> <pre>systemctl stop irqbalance modprobe -r ib_ipoib modprobe -r hfi1 modprobe hfi1 modprobe ib_ipoib systemctl start irqbalance</pre>
STL-59663	<p>Certain conditions result in hfi1 driver interrupts not being distributed correctly, even with irqbalance running. This may result in a large run to run performance variation or low overall performance.</p>	<p>Refer to the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i>, Identifying to Which CPU Core an Interrupt is Bound section for help with determining which cores are running the hfi1 interrupts. If interrupts are arbitrarily distributed or significantly reusing select cpu cores, refer to Manually Changing IRQ Affinity section.</p> <p>Contact Intel Customer Support if you need further guidance.</p>
STL-59934	<p>Customers experience errors similar to hfi1_0: Send Context 8(151) Error: WriteOverflow in dmesg or console logs.</p>	<p>Contact Intel Customer Support for more information.</p> <p>Setting pio_threshold=0 may prevent the error.</p>
STL-59955	<p>irqbalance cannot start user/admin-supplied hint script when SELinux is enabled and enforcing.</p>	<p>Disable SELinux.</p> <p>Refer to the Red Hat Bugzilla https://bugzilla.redhat.com/show_bug.cgi?id=1784080</p>
STL-60805	<p>Mixing IB_WR_RDMA_WRITE_WITH_IMM requests with IB_WR_RDMA_WRITE requests can cause a timeout when TID RDMA is enabled:</p> <pre>hfi1 0000:7f:00.0: hfi1_0: [QP74] hfi1_tid_timeout 4084</pre> <p><i>Note:</i> This issue is ULP-dependent and does not appear to be typical.</p>	<p>Disable TID RDMA.</p> <p>Refer to the Red Hat Bugzilla https://bugzilla.redhat.com/show_bug.cgi?id=1870281</p> <p>A fix is scheduled for 8.3 z-stream kernel-4.18.0-240.8.el8.</p>
STL-60921	<p>When Adaptive Routing is enabled, the topology thread responsible for handling the sweep functionality of the SM sometimes hangs. This prevents the SM from responding to changes within the fabric.</p>	<p>None.</p> <p>The FM must be restarted in order to resolve the problem.</p>

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/Verbs.	<p>Note: 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.
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="394 976 902 1066">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>