

# **Intel<sup>®</sup> Omni-Path Fabric Software in SUSE\* Linux\* Enterprise Server 12 SP4**

**Release Notes**

---

*Rev. 2.0*

*March 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](http://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 © 2018, Intel Corporation. All rights reserved.



## Contents

---

<b>1.0 Overview of the Release</b> .....	<b>5</b>
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
<b>2.0 Issues</b> .....	<b>12</b>
2.1 Resolved Issues.....	12
2.2 Open Issues.....	15



## Tables

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



## 1.0 Overview of the Release

These Release Notes are intended for Intel® Omni-Path IFS 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 V10.7.

- Intel® Omni-Path Switches Installation, User, Reference Guides, and Release Notes  
<http://www.intel.com/omnipath/SwitchPublications>
- Intel® Omni-Path Software Installation, User, Reference Guides, and Release Notes (includes HFI documents)  
<http://www.intel.com/omnipath/FabricSoftwarePublications>

The following table lists the end user document versions supported by this release.

**Table 1. Supported Document Versions**

Title	Doc. Number	Revision
<i>Intel® Omni-Path Fabric Quick Start Guide</i>	J57479	4.0
<i>Intel® Omni-Path Fabric Setup Guide</i>	J27600	8.0
<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	9.0
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	H76457	9.0
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	H76458	9.0
<i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i> (Merged with: <i>Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide</i> )	H76469	9.0
<b>continued...</b>		



Title	Doc. Number	Revision
Intel® Omni-Path Fabric Suite Fabric Manager User Guide	H76468	9.0
Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide	H76471	9.0
Intel® Omni-Path Fabric Host Software User Guide	H76470	9.0
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	9.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	11.0
Intel® Omni-Path IP and LNet Router Design Guide (Old title: Intel® Omni-Path IP and Storage Router Design Guide)	H99668	6.0
Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note	J57474	4.0
Intel® Omni-Path Management API Programmer's Guide	J68876	3.0
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	J95967	1.0
Intel® Omni-Path Fabric Manager GUI Release Notes	J95968	1.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally-managed switches)	J95964	1.0
Intel® Omni-Path Fabric Unified Extensible Firmware Interface (UEFI) Release Notes	J98868	1.0
Intel® Omni-Path Fabric Thermal Management Microchip (TMM) Release Notes	J98871	1.0
Intel® Omni-Path Fabric Firmware Tools Release Notes	J98870	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
<b>Packages created by Intel</b>
opa-address-resolution-10.7.0-1.27.x86_64
opa-basic-tools-10.7.0-1.27.x86_64
opa-fastfabric-10.7.0-1.27.x86_64
opa-fm-10.7.0-1.27.x86_64
opa-libopamgt-10.7.0-1.27.x86_64
<i>continued...</i>




Intel® Omni-Path Software Packages
libpsm2-2-10.3.37-1.19.x86_64
libfabric-1.5.3-1.31.x86_64
<b>Firmware binaries delivered by Intel</b>
8051 firmware version 1.26.1
SBus Master firmware version 0x10130001
PCIe SerDes firmware version 0x4755
Fabric SerDes firmware version 0x1055
<b>Packages used by Intel</b>
rdma-core-16.5-1.33.x86_64 (libhfi1)
openmpi-1.10.7-3.11.x86_64
mpitests-openmpi-3.2-13.15.x86_64
mvapich2-psm2-2.2-8.8.x86_64
mpitests-mvapich2-psm2-3.2-13.13.x86_64

To download Intel programmable firmware, refer to the following:

- [Unified Extensible Firmware Interface \(UEFI\)](#)
- [Thermal Management Module \(TMM\)](#)
- [Firmware Tools](#)

## 1.6 Supported Features

- The list of supported hardware is in [Table 3](#) on page 8.
- Product Constraints described in [Product Constraints](#) on page 11.
- 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.
- Support for the Enhanced Hypercube Routing Engine is outside the scope of Intel® OPA support. However, Intel partners may offer such support as part of their solutions. In addition there is an open source community who may be able to answer specific questions and provide guidance with respect to the Enhanced Hypercube Routing Engine.

## 1.7 Supported MPI Libraries

The table below lists the different MPI libraries supported by Intel® Omni-Path Fabric Software. Note that the second column indicates whether the MPI library is included in the distribution or not.

**Table 2. Supported MPI Libraries**

MPI Implementation	Included in Distribution?	Runs Over
Open MPI 2.1.2	Yes	PSM2
MVAPICH2-2.3B	Yes	PSM2

## 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 3. 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
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
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
<i>continued...</i>	





Hardware	Description
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 SP4.

**Table 4. Intel® OPA Compatibility Matrix**

UEFI	TMM	Managed Switch	Externally-Managed Switch	FM GUI
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
1.6.0.0.0	10.4.0.0.146	10.6.1.0.3	10.6.1.0.1	10.6.0.0.136

## 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® OP Software configuration using SLES\* OS:



Step	Task/Prompt	Action
<b>Install OPA-Basic Software</b>		
1.	At the command prompt, enter the installation command for opa-basic-tools.	Type <b>zypper install -y opa-basic-tools</b> and press <b>Enter</b> .
2.	At the command prompt, reboot the server.	Type <b>reboot</b> and press <b>Enter</b> .
3.	Check your link using opainfo.	Type <b>opainfo</b> and press <b>Enter</b> . 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 <b>zypper install -y rdma-core</b> and press <b>Enter</b> .
5.	On all compute nodes: install the PSM2 library.	Type <b>zypper install -y libpsm2-2</b> and press <b>Enter</b> .
<b>Install Intel® Omni-Path Fabric Suite Components on the Management Node</b>		
6.	Install FastFabric.	Type <b>zypper install -y opa-fastfabric</b> and press <b>Enter</b> .
8.	Install Fabric Manager.	Type <b>zypper install -y opa-fm</b> and press <b>Enter</b> .
9.	Start the Fabric Manager.	Type <b>systemctl start opafm</b> and press <b>Enter</b> .
<b>Set up IPoIB IPV4 Configuration</b>		
10.	Manually edit or create the ifcfg-ibX file.	<i>Note:</i> Use the OS distribution-supplied instructions for setting up network interfaces. Type <b>cat /etc/network/ifcfg-ib0</b> and press <b>Enter</b> . 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 IPOIB MODE='connected' MTU=65520</pre>
11.	Bring up the ib0 interface.	Type <b>ifup ib0</b> and press <b>Enter</b> .
12.	Perform a test ping.	Type <b>ping &lt;remote IPoIB address&gt;</b> and press <b>Enter</b> .

*continued...*



Step	Task/Prompt	Action
		For example: <pre data-bbox="870 373 1396 464">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>
<b>(Optional) Install the Fabric Manager GUI</b>		
13.	On one node in the fabric: install the Fabric Manager GUI.	<i>Note:</i> Intel recommends not to install the Fabric Manager GUI on the Management Node where the Fabric Manager is being used. Type <b>zypper install -y opa-fmgui</b> and press <b>Enter</b> .
<b>End Task</b>		

### 1.11 Product Constraints

- The minimum firmware version for Intel® Omni-Path Host Fabric Interface Silicon 100 Series Switch ASIC is 10.6.
- Power class 2 AOC are supported. You must use 1.5 (or newer) UEFI for proper operation. 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/sysconfig/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 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 5. Issues Resolved in this Release**

ID	Description	Resolved in Release
131017	Verbs <code>ib_send_bw</code> , <code>ib_read_bw</code> , and <code>ib_write_bw</code> are not working with the <code>-R</code> option to use the RDMA CM API to create QPs and exch data.	SLES* 12 SP4
133380	The PM has been updated allowing you to change the weight and threshold of PA categories. This enables you to recalculate values using already stored port data. Also, PA query time, memory usage, and disk space usage will decrease with the new PM History version. To change thresholds and weights, edit the <code>opafm.xml</code> file and restart the FM. <i>Note:</i> The FM no longer supports the previous short term history (STH) file after this change. The old files do not need to be removed as they will age out normally.	SLES* 12 SP4
134353	Very infrequently, when a link goes down, the logical link state can remain stuck in the 'Init' state.	SLES* 12 SP4
134409 135259	In links exhibiting a high error rate, a rare <code>PortRcvError</code> is possible, resulting in a link down event. Such links should retrain and return to operation without user interaction. In cases where the Link Quality is less than or equal to 3, the interconnect in the link should be evaluated for possible replacement to prevent future <code>PortRcvErrors</code> from occurring.	SLES* 12 SP4
134493	When using <code>MVAPICH2</code> with Intel® Omni-Path PSM2, users will notice unexpected behavior when seeding the built-in random number generator with functions like <code>srand</code> or <code>srandom</code> before <code>MPI_Init</code> is called. <code>MPI_Init</code> re-seeds the random number generator with its own value and does not restore the seed set by the user application. This causes different MPI ranks to generate different sequences of random numbers even though they started with the same seed value.	SLES* 12 SP4
134494	Open MPI uses <code>srand()</code> family functions at <code>MPI_Init()</code> time. Therefore, if the user sets <code>srand()</code> before calling <code>MPI_Init()</code> , the values will be altered.	SLES* 12 SP4
135040	You cannot currently specify portions of an Intel® Omni-Path Director Class Switch chassis that are not populated and are not expected to be populated. If <code>CoreFull</code> is 1, all the internal links for that chassis are generated when run against <code>opaxlattopology</code> . If <code>CoreFull</code> is 0, none of the links are generated.	SLES* 12 SP4
135180	OpenMPI/PSM2 timeouts during MPI stress tests on Haswell and Intel® Xeon Phi™ Processor mixed fabrics.	SLES* 12 SP4
135326	Calling <code>opasmaquery</code> fails when called from a non-SM node to a node which has not booted to the OS.	SLES* 12 SP4
135390	The driver can parse older versions of the platform configuration file.	SLES* 12 SP4
135545	A change has been made to several SA record attributes which causes incompatibilities between the Fabric tool suite and the SA.	SLES* 12 SP4
		<i>continued...</i>




ID	Description	Resolved in Release
135711	After generating the <code>opafm.xml</code> file from the <code>config_generate</code> script, the FE is not enabled.	SLES* 12 SP4
135929	Intel® Omni-Path Boot nodes occasionally dropped from fabric when switching master SM from one node to another.	SLES* 12 SP4
136049	The expected width of a card is not showing up correctly in <code>opaverifyhosts</code> .	SLES* 12 SP4
136567	The Intel® OPA software does not support SSH passphrases. In earlier releases, if a passphrase was created, an error message was displayed to remove the passphrase and proceed. When SSH keys are created, the passphrase is automatically left empty.	SLES* 12 SP4
136727	Initialization of PSM2 library fails with the following error message: <code>Error: PSM is in the finalized state</code>	SLES* 12 SP4
136728	If hundreds of links are bouncing while the FM is sweeping, the FM sweep time may be significantly extended. This can result in unexpected delays in FM responsiveness to fabric changes or host reboots. (The issue is that active links bounce between the time FM discovers one side of the link versus the other side of the link.) A new configuration value is present in the FM configuration that determines how much time will be allotted to timeouts before abandoning a sweep. If you are upgrading from a previous version of the FM and retaining a configuration file that does not include this new parameter, the value will be set too low and cause sweeps to abandon after only a single timeout is witnessed.	SLES* 12 SP4
136733	Slow memory deregistration has been observed.	SLES* 12 SP4
136902	A snapshot file with a multicast group with rate 10g will not be read properly. The following error is returned: <code>opafabricanalysis: Port 0:0 Error: Unable to analyze fabric snapshot. See /var/usr/lib/opa/analysis/latest/fabric.0:0.links.stderr opafabricanalysis: Possible fabric errors or changes found</code>	SLES* 12 SP4
136985	<code>opahfirev</code> has output errors when the HFI driver is not installed.	SLES* 12 SP4
136995	The <code>opahfirev</code> tool output uses the term "HWRev" to indicate the revision of the silicon on the card.	SLES* 12 SP4
137123	The Fabric Manager is not compatible with older versions of the FM GUI. You must use the same version of both Fabric Manager and FM GUI.	SLES* 12 SP4
137221	Querying for switch info with <code>opasmaquery</code> while using the <code>-g</code> option will print incorrect IPv4 addresses.	SLES* 12 SP4
137364	The node description of a node may change after rebooting. This issue requires manual enabling and starting of the <code>rdma-ndd</code> service.	SLES* 12 SP4
137499	HFI links may occasionally take several minutes to reach link up.	SLES* 12 SP4
137708	Following a link bounce event, there is a possibility that a link will fail to reach the Armed/Active state. The likelihood of this issue depends largely on the link type: <ul style="list-style-type: none"> <li>• Compute Nodes: These links are very unlikely to be affected.</li> <li>• FM Nodes: These links are the most exposed. If an FM link is affected and not recovered, there may be downstream effects over time.</li> </ul>	SLES* 12 SP4
137744	The values for <code>MinInitial</code> and <code>MinTail</code> were reported in flits by the <code>opareport</code> , <code>opasmaquery</code> , and <code>opasaquery</code> tools. This output is now converted to bytes, and is displayed in decimal. See the <i>Intel® Omni-Path Fabric Suite FastFabric User Guide</i> for details.	SLES* 12 SP4

*continued...*



ID	Description	Resolved in Release
137791	Changes to settings in the Preemption section of the FM configuration file are only updated on HFI or switch ports when the port is bounced. See the <i>Intel® Omni-Path Fabric Suite Fabric Manager User Guide</i> for details.	SLES* 12 SP4
138047	The Open MPI implementation for MPI_Wtime() may change when using different CPU frequency drivers (intel_pstate vs acpi_freq) and turbo status of the CPU.	SLES* 12 SP4
138108	PKey handling for Active VFs was changed. To avoid disruptions when upgrading from a prior release, you must ensure that all Active VFs have explicit PKeys defined in the opafm.xml configuration file. To find PKeys that are currently assigned to each Active VF, type <code>opareport -o vfinfo</code> and press <b>Enter</b> . Using this information, manually edit the <VirtualFabric> section of the opafm.xml file for each VF in the list to insert the following: <Pkey>pkey_number</Pkey>	SLES* 12 SP4
138183	Additional fields were added to the <code>opareport -o snapshot -r</code> XML output format that were not present in the previous release. Therefore, the previous release snapshot files reported a "Mandatory Tag Not Found" parser error using the newer release Fabric Manager tools. The resolution was to regenerate any such snapshot files using the <code>opareport</code> tool in the newer release.	SLES* 12 SP4
138460	When upgrading from the previous release, the allhost path include statement was not updated with the new path which is: <code>/etc/opa/allhosts</code>	SLES* 12 SP4
139407	Shell history overflow caused by Intel® OPA commands run by scripts. Commands run by FastFabric <code>opahostadmin</code> will now be omitted from the shell history.	SLES* 12 SP4
139550	Infrequently, an AOC may exhibit an unexpectedly high local link integrity error rate after the link comes up, relative to the error rate on previous link up occasions. This can be determined by observing a link quality of <5. These links may eventually experience a link width downgrade.	SLES* 12 SP4
139797	Switch port connected to HFI stuck in LinkTearDown state.	SLES* 12 SP4
139834	When using the FastFabric TUI to run "Perform Single Host Verification", the test hangs during operation.	SLES* 12 SP4
140199	In some scenarios involving device reboot, down links, fabrics with spare ports or fabrics with DCS; that the SM may skip setting some important SMA attributes. This can result in ports which are Active but unable to pass data, resulting in errors and failures from assorted applications depending on which nodes are communicating with each other.	SLES* 12 SP4
140229	The opaswitchadmin tool was updated to address a condition that was seen during firmware upgrade of a large number of switches.	SLES* 12 SP4
140881	In rare cases when an LNI failure occurs, the link will not come up after manually disabling and re-enabling the link.	SLES* 12 SP4
140911	The OFI verbs provider does not support FI_EP_RDM End Point type. This End Point type is needed for Open MPI OFI support. Therefore, Open MPI OFI support will not run over the verbs provider.	SLES* 12 SP4
141219	When adaptive routing is disabled, the output for <code>opasmaquery</code> for <code>portgroup</code> appears as shown below:  <pre># opasmaquery -l 1 -o portgroup PG: 0x0000 Egress:None</pre>	SLES* 12 SP4
141845	Resolved FM process out of memory condition	SLES* 12 SP4
141909	Resolved multiple FM synchronization issue that can lead to FM failure.	SLES* 12 SP4




## 2.2 Open Issues

The following table lists the open issues for this release.

**Table 6. Open Issues**

ID	Description	Workaround
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, the following workaround is required:</p> <ul style="list-style-type: none"> <li>When running MVAPICH2 jobs with a large number of ranks (for example, &gt; 36 ranks but ≤ 72 ranks), you must set the following parameters in <code>/etc/security/limits.conf</code>:                             <ul style="list-style-type: none"> <li>hard memlock unlimited</li> <li>soft memlock unlimited</li> </ul> </li> <li>Also, you must increase the <code>lkey_table_size:LKEY</code> table size in bits (<math>2^n</math>, where <math>1 \leq n \leq 23</math>) 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.</li> </ul>
135830	<p>On Intel® Xeon Phi™ systems, failure observed during software upgrade when rebuilding the boot image. Error message contains:</p> <p>Rebuilding boot image with <code>"/usr/bin/dracut -f"</code></p>	<p>Due to the extended processing time of the <code>dracut</code> command on the Intel® Xeon Phi™ platform, Intel recommends the following:</p> <ul style="list-style-type: none"> <li>Install and configure Intel® Xeon Phi™ systems separately.</li> <li>Change the <code>FF_TIMEOUT_MULT</code> value in <code>opafastfabric.conf</code> from 2 to 6 for Intel® Xeon Phi™ systems.</li> </ul>
139368	<p>Some applications compiled with older compilers may use a personality bit that signifies that READ should imply EXECUTE permissions.</p> <p>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.</p>	<p>As root, run the <code>execstack</code> tool to clear the executable bit on the binary:</p> <pre>execstack -c &lt;binary&gt;</pre> <p>Alternatively, recompile the binary to not set this personality bit.</p>
139613	<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>
142330	<p>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.</p>	<p>Disable upstream error reporting using the AER mask register.</p> <ul style="list-style-type: none"> <li>For discrete HFI ASICs, use</li> </ul> <pre>setpci -d 8086:24f0 ECAP_AER +8.1=00100000:00100000</pre>

**continued...**



ID	Description	Workaround
		<ul style="list-style-type: none"><li>For integrated HFIs, use</li></ul> <pre>setpci -d 8086:24f1 ECAP_AER +8.l=00100000:00100000</pre>
143296	<p>When irqbalance 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.</p> <p>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.</p>	<p>Always start the user-space process irqbalance using the argument <code>--hintpolicy=exact</code>.</p>
143449	<p>PM will scroll LQI=0 and Integrity Exceeded Threshold logs when an additional VF with QoS enabled and a device group that is not "All".</p> <p><i>Note:</i> This issue does not occur when running against the default opafm.xml configuration file.</p>	<p>Set the <code>&lt;ProcessVLCounters&gt;</code> field in the opafm.xml configuration to 0 to stop scrolling of logs related to LQI.</p>