



Intel® Omni-Path Fabric Software in Red Hat* Enterprise Linux* 7.5

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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit <http://www.intel.com/design/literature.htm>.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. 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 © 2018, Intel Corporation. All rights reserved.



Contents

1.0 Overview of the Release	5
1.1 Audience.....	5
1.2 Related Documents.....	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.....	8
1.10 Installation Requirements.....	9
1.10.1 Installation Instructions.....	9
1.11 Product Constraints.....	10
1.12 Product Limitations.....	11
2.0 Issues	12
2.1 Resolved Issues.....	12
2.2 Open Issues.....	14



Tables

1	Related Documents.....	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.....	14



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 Related Documents

The following table lists the V10.5 end user documents supported by this release.

Table 1. Related Documents

Title	Doc. Number	Revision
Key: Shading indicates the URL to use for accessing the particular document.		
<ul style="list-style-type: none"> Intel® Omni-Path Switches Installation, User, and Reference Guides for: http://www.intel.com/omnipath/SwitchPublications 		
<ul style="list-style-type: none"> Intel® Omni-Path Software Installation, User, and Reference Guides (includes HFI documents) for: http://www.intel.com/omnipath/FabricSoftwarePublications 		
<ul style="list-style-type: none"> Drivers and Software (including Release Notes) for: http://www.intel.com/omnipath/Downloads 		
<i>Intel® Omni-Path Fabric Quick Start Guide</i>	J57479	2.0
<i>Intel® Omni-Path Fabric Setup Guide</i> (Old title: <i>Intel® Omni-Path Fabric Staging Guide</i>)	J27600	6.0
<i>Intel® Omni-Path Fabric Switches Hardware Installation Guide</i>	H76456	6.0
<i>Intel® Omni-Path Host Fabric Interface Installation Guide</i>	H76466	5.0
<i>Intel® Omni-Path Fabric Software Installation Guide</i>	H76467	7.0
<i>Intel® Omni-Path Fabric Switches GUI User Guide</i>	H76457	7.0
<i>Intel® Omni-Path Fabric Switches Command Line Interface Reference Guide</i>	H76458	7.0
<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	7.0
<i>continued...</i>		



Title	Doc. Number	Revision
Intel® Omni-Path Fabric Suite Fabric Manager User Guide	H76468	7.0
Intel® Omni-Path Fabric Suite Fabric Manager GUI User Guide	H76471	7.0
Intel® Omni-Path Fabric Host Software User Guide	H76470	7.0
Intel® Performance Scaled Messaging 2 (PSM2) Programmer's Guide	H76473	7.0
Intel® Omni-Path Fabric Performance Tuning User Guide	H93143	9.0
Intel® Omni-Path IP and Storage Router Design Guide	H99668	5.0
Building Lustre* Servers with Intel® Omni-Path Architecture Application Note	J10040	1.0
Building Containers for Intel® Omni-Path Fabrics using Docker* and Singularity* Application Note	J57474	3.0
Intel® Omni-Path Management API Programmer's Guide	J68876	1.0
Intel® Omni-Path Fabric Software Release Notes	J75208	2.0
Intel® Omni-Path Fabric Manager GUI Release Notes	J75209	2.0
Intel® Omni-Path Fabric Switches Release Notes (includes managed and externally- managed switches)	J75207	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-10.5.0.0.140-1.el7.x86_64
opa-basic-tools-10.5.0.0.140-1.el7.x86_64
opa-fastfabric-10.5.0.0.140-1.el7.x86_64
opa-fm-10.5.1.0.1-1.el7.x86_64
opa-libopamgt-10.5.0.0.140-1.el7.x86_64
libpsm2-10.3.8-2.el7.x86_64
Firmware binaries delivered by Intel
8051 firmware version 1.24.0
SBus Master firmware version 0x10130001
<i>continued...</i>



Intel® Omni-Path Software Packages
PCIe SerDes firmware version 0x4755
Fabric SerDes firmware version 0x1055
Packages used by Intel
rdma-core-15-2.el7.x86_64 (libhfi1)
openmpi-1.10.7-1.el7.x86_64
mpitests-openmpi-5.4-1.el7.x86_64
mvapich2-2.2-psm2-2.2-1.1.el7.x86_64
mpitests-mvapich222-psm2-5.4-1.el7.x86_64

1.6 Supported Features

- The list of supported hardware is in [Table 3](#) on page 8.
- 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 10 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.
- Topology-aware job scheduling, which is enabled by the opamgt library and allows developers to write code that interfaces to the SA/PA. See the *Intel® Omni-Path Management API Programmer's Guide* for details.
- 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
- Support for IBM* Platform MPI and IBM* Spectrum MPI. See [Supported MPI Libraries](#) on page 8 for details.
- 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.
- Product Constraints described in [Product Constraints](#) on page 10.



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 1.10.7	Yes	PSM2
MVAPICH2-2.2	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® 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
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* 7.5.



Table 4. Intel® OPA Compatibility Matrix

UEFI	TMM	Managed Switch	Externally-Managed Switch	FM GUI
10.6.1.0.2	1.6.0.0.0	10.4.0.0.146	10.6.1.0.3	10.6.1.0.1
10.5.1.0.2	1.5.2.0.0	10.4.0.0.146	10.5.1.0.2	10.5.0.0.141
10.4.2.0.7	1.4.2.0.0	10.4.0.0.146	10.4.3.0.1	10.4.3.0.1

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.
- 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® OP Software configuration using RHEL* OS:

Step	Task/Prompt	Action
Install OPA-Basic Software		
1.	At the command prompt, enter the installation command for opa-basic-tools.	Type yum 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 </pre>
<i>continued...</i>		



Step	Task/Prompt	Action
		<pre>FCBN425QB1C05 Rev A 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 .
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 BOOTPROTO=static IPADDR=10.228.200.173 BROADCAST=10.228.203.255 NETWORK=10.228.200.0 NETMASK=255.255.252.0 ONBOOT=yes CONNECTED_MODE=yes MTU=65520</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 and require RHEL* 7.5 on both ends of the cable. Specifically, 1.5 level UEFI or higher are required for proper operation. Integrated HFI (-F) requires a specific BIOS level to support power class 2 AOC; contact your BIOS vendor for more information.
- The minimum firmware version for Intel® Omni-Path Host Fabric Interface Silicon 100 Series Switch ASIC is 10.2.



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/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
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.	RHEL* 7.5
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 opafm.xml file and restart the FM. Note: 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.	RHEL* 7.5
134409	In links exhibiting a high error rate, a rare PortRcvError 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 PortRcvErrors from occurring.	RHEL* 7.5
135259	In links exhibiting a high error rate, a rare PortRcvError 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 PortRcvErrors from occurring.	RHEL* 7.5
135390	Very old HFI adapters may be programmed with an obsolete version of the AOC platform configuration file. In these cases, errors such as the following may be observed: <pre>[26.903186] hfi1 0000:d5:00.0: hfi1_0: parse_platform_config:Bad config file [26.903186] hfi1 0000:d5:00.0: hfi1_0: parse_platform_config:File claims to be larger than read size [27.351555] hfi1 0000:d5:00.0: hfi1_0: tune_serdes: Unknown port type</pre> Update the platform configuration file on the HFI to the current version. For details, see the <i>Intel® Omni-Path Fabric Software Installation Guide</i> , section B.1.	RHEL* 7.5
135929	Intel® Omni-Path Boot nodes occasionally dropped from fabric when switching master SM from one node to another.	RHEL* 7.5
136049	The expected width of a card is not showing up correctly in <code>opaverifyhosts</code> .	RHEL* 7.5
136437	When using previous RHEL* distros, the default generic PXE boot image does not work due to missing driver and firmware files.	RHEL* 7.5
136727	Initialization of PSM2 library fails with the following error message: Error: PSM is in the finalized state	RHEL* 7.5
136901	Occasionally, nodes may be dropped by the Fabric Manager while they are in a pre-boot mode. This can occur when the node has multiple HFIs on a single socket.	RHEL* 7.5
<i>continued...</i>		



ID	Description	Resolved in Release
137123	In past RHEL versions, the Fabric Manager was not compatible with older versions of the FM GUI. You must use the same version of both Fabric Manager and FM GUI.	RHEL* 7.5
137372	Packets may be stuck in kernel when attempting writes to file system via IPoIB interface.	RHEL* 7.5
137499	HFI links may occasionally take several minutes to reach link up.	RHEL* 7.5
137577	opatmmtool does not provide a correct error message if it is run on a system that does not have a TMM.	RHEL* 7.5
137744	In previous releases, the values for MinInitial and MinTail were reported in flits by the opareport, opasmaquery, and opasaquery tools. In the current RHEL* 7.5 release, this output is now converted to bytes, and is displayed in decimal. See the Intel® Omni-Path Fabric Suite FastFabric Command Line Interface Reference Guide for details.	RHEL* 7.5
137791	In the current RHEL* 7.5, 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 Intel® Omni-Path Fabric Suite Fabric Manager User Guide for details.	RHEL* 7.5
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. RHEL 7.5 includes patches for Open MPI 1.10.x that resolve this issue.	RHEL* 7.5
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 opareport -o vfinfo and press Enter . 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>	RHEL* 7.5
138183	In RHEL* 7.5, additional fields were added to the opareport -o snapshot -r XML output format that were not present in RHEL* 7.4. Therefore, RHEL* 7.4 snapshot files reported a "Mandatory Tag Not Found" parser error using RHEL* 7.5 Fabric Manager tools. The resolution was to regenerate any such snapshot files using the opareport tool in RHEL* 7.5.	RHEL* 7.5
139407	Shell history overflow was caused by Intel® OPA commands run by scripts. In RHEL* 7.5, commands run by FastFabric opahostadmin will be omitted from the shell history. This allows users to more easily repeat or review past commands.	RHEL* 7.5
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"> • Compute Nodes: These links are very unlikely to be affected. • FM Nodes: These links are the most exposed. If an FM link is affected and not recovered, there may be downstream effects over time. 	RHEL* 7.5
131017	Verbs ib_send_bw, ib_read_bw, and ib_write_bw are not working with the -R option to use the RDMA CM API to create QPs and exch data.	RHEL* 7.5
134353	Very infrequently, when a link goes down, the logical link state can remain stuck in the 'Init' state.	RHEL* 7.5
134493	When using MVAPICH2 with Intel® Omni-Path PSM2, users will notice unexpected behavior when seeding the built-in random number generator with functions like srand or srandom before MPI_Init is called. MPI_Init 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.	RHEL* 7.5
135040	You can't currently specify portions of an Intel® DCS chassis that is not populated and is not expected to be populated. If CoreFull is 1, all the internal links for that chassis are generated when run against opaxlattopology. If CoreFull is 0, none of the links are generated.	RHEL* 7.5

continued...



ID	Description	Resolved in Release
135180	OpenMPI/PSM2 timeouts during MPI stress tests on Haswell and Intel® Xeon Phi™ mixed fabrics.	RHEL* 7.5
135545	A change has been made to several SA record attributes which causes incompatibilities between the Fabric tool suite and the SA.	RHEL* 7.5
135711	After generating the opafm.xml file from the config_generate script, the FE is not enabled.	RHEL* 7.5
136733	Slow memory deregistration has been observed.	RHEL* 7.5
136902	A snapshot file with a multicast group with rate 10g will not be read properly. The following error is returned: 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	RHEL* 7.5
136985	opahfirev has output errors when the HFI driver is not installed.	RHEL* 7.5
136995	The opahfirev tool output uses the term "HWRev" to indicate the revision of the silicon on the card.	RHEL* 7.5
137221	Querying for switch info with opasmaquery while using the -g option will print incorrect IPv4 addresses.	RHEL* 7.5

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.	When running MVAPICH2 jobs with a large number of ranks (for example, between 36 and 72 ranks), you must set the following parameters in /etc/security/limits.conf: <ul style="list-style-type: none"> • hard memlock unlimited • soft memlock unlimited Also, you must increase the lkey_table_size:LKEY table size in bits (2^n , $1 \leq n \leq 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to Appendix A in the <i>Intel® Omni-Path Fabric Performance Tuning User Guide</i> .
134494	Open MPI uses srand() family functions at MPI_Init() time. Therefore, if the user sets srand() before calling MPI_Init(), the values will be altered.	<ol style="list-style-type: none"> 1. Fixed in Open MPI 2.0.1. 2. Call srand() functions family after calling MPI_Init().
135390	Very old HFI adapters may be programmed with an obsolete version of the AOC platform configuration file. In these cases, errors such as the following may be observed: <pre>[26.903186] hfi1 0000:d5:00.0: hfi1_0: parse_platform_config:Bad config file [26.903186] hfi1 0000:d5:00.0: hfi1_0: parse_platform_config:File claims to be larger than read size [27.351555] hfi1 0000:d5:00.0: hfi1_0: tune_serdes: Unknown port type</pre>	Update the platform configuration file on the HFI to the current version. For details, see the <i>Intel® Omni-Path Fabric Software Installation Guide</i> , section B.1.

continued...



ID	Description	Workaround
135975	After performing an OPA software configuration update, switches will show the new settings when queried by opaswitchadmin tools, however, individual ports will continue to operate using the previous settings, including LinkWidth enable.	Perform one of the following: <ul style="list-style-type: none"> • Reboot the switch to apply the new settings. • Bounce the affected port(s).
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.) In the previous RHEL version, a change was made to improve the FM responsiveness in large fabrics of >1000 nodes when numerous links bounce (or hosts are rebooted) at once.	Ensure the following value is present in the <Sm> section of the opafm.xml file: <pre><CumulativeTimeoutLimit>300</CumulativeTimeoutLimit></pre>
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.
137409	When using DHCP from an Intel OPA HFI, the DHCP client-identifier field (option 61) must be used to set up DHCP static leases. However, what is sent in the DHCP discover packet is different for PXE boot attempts versus a DHCP boot attempt.	The DHCP client-identifier can be explicitly specified in DHCP client configuration or NetworkManager scripts.
137951	In the HFI BIOS screen for Advanced NIC Configuration, a warning message about incorrect custom P_Key value is not completely displayed.	The valid range for custom P_Key value is 0x8001 to 0xFFFFE.
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.	Bounce the link.
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 lspci 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.
139650	In rare cases, AOCs with mismatched firmware on each end of the link may experience longer than expected link-up times.	If link-up times are unacceptably long, ensure both ends of links are running the same firmware version.
139660	Following a boot, it is possible, although rare, that the IPoIB interface will fail to come up. Hosts attempting to ping this host will get no response.	Reboot or reload the driver.
139834	When using the FastFabric TUI to run "Perform Single Host Verification", the test hangs during operation.	Run the hostverify.sh script manually using the opaverifyhosts command.
140691	When running opaswitchadmin against multiple externally-managed switches simultaneously, it sends schedule requests in parallel to those hosts. It is possible that some hosts may intermittently fail due to timeouts at high levels of parallelism.	The number of hosts that are queried in parallel can be limited by setting the FF_MAX_PARALLEL value in /etc/opa/opafastfabric.conf. Results may vary by fabric. Intel recommends reducing this to 2.