

Intel[®] True Scale Fabric OFED+ Host Software

Release Notes

January 2015



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

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications.

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.

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

Intel and the Intel logo 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 © 2015, Intel Corporation. All rights reserved.



Contents

1.0	Overview of the Release	7
1.1	Introduction	7
1.2	Audience	7
1.3	If You Need Help	7
1.4	New Features and Enhancements	7
1.4.1	Release 7.3.1.0.12 Enhancements	7
1.4.2	Release 7.3.0.0.26 Enhancements	7
1.4.3	Release 7.2.2.0.8 Enhancements	8
1.5	Operating Environments Supported	9
1.6	Qualified Parallel File Systems	10
1.7	Intel Interface for NVIDIA GPUs	10
1.8	Compilers Supported	10
1.8.1	MPI	10
1.8.2	MVAPICH, MVAPICH2 and Open MPI	10
1.9	Hardware Supported	12
1.10	Installation Requirements	12
1.10.1	Software and Firmware Requirements	12
1.10.2	Installation Instructions	12
1.11	Changes for this Release	13
1.11.1	Changes to Software Components	13
1.11.2	Changes to Industry Standards Compliance	13
1.12	Product Constraints	13
1.13	Product Limitations	14
1.14	Other Information	14
1.14.1	Verbs Performance	16
1.15	Documentation	16
1.16		16
2.0	System Issues for Release 7.3.1.0.12	17
2.1	Introduction	17
2.2	Resolved Issues in this Release	17
2.3	Known Issues	18
2.3.1	Open Issues Table	18






Tables

1-1 Operating Environments Supported	9
1-2 CPU Model of Linux Kernel	9
1-3 NVIDIA's CUDA Tested with OFED+	10
1-4 MPI Compilers	10
1-5 MVAPICH, MVAPICH2 and Open MPI	11
1-6 Hardware Supported	12
1-7 Changes to Software Component Support	13
1-8 Changes to Industry Standards Compliance	13
1-9 Related Documentation for this Release	16
2-1 Resolved Issues	17
2-2 Open Issues	18





1.0 Overview of the Release

1.1 Introduction

These Release Notes provide a brief overview of the changes introduced into the Intel® True Scale Fabric OFED+ by this release. References to more detailed information are provided where necessary. The information contained in this document is intended for supplemental use only; it should be used in conjunction with the documentation provided for each component.

These Release Notes list the new features of the release, as well as the system issues that were closed in the development of Release 7.3.1.0.12.

1.2 Audience

The information provided in this document is intended for installers, software support engineers, and service personnel.

1.3 If You Need Help

If you need assistance while working with the OFED+ Host Software, contact your Intel® approved reseller or Intel® True Scale Technical Support:

- By E-mail:
ibsupport@intel.com
- On the Support tab at web site:
<http://www.intel.com/truescale>

For OEM-specific server platforms supported by this release, contact your OEM.

1.4 New Features and Enhancements

This section list the new features and enhancements for Release 7.3.1.0.12, as well as the two previous releases.

1.4.1 Release 7.3.1.0.12 Enhancements

- Added support for
 - RHEL 6.6
 - CentOS 6.6 and 7.0
 - SLES 6.6 and 7.0

1.4.2 Release 7.3.0.0.26 Enhancements

- Added support for



- RedHat EL 7
- openmpi 1.8.1
- mvapich2 1.8.1
- OFED-3.5-2

1.4.3 Release 7.2.2.0.8 Enhancements

- Added support for
 - RedHat EL 5.10 and 6.5
 - CentOS 5.10 and 6.5
 - Scientific Linux 5.10 and 6.5



1.5 Operating Environments Supported

The Release 7.3.0.0.26 version of OFED+ Host Software allows for the Operating Systems listed in [Table 1-1](#).

Table 1-1. Operating Environments Supported

Operating System	Update/SP	Version
RHEL 6 X86_64 (AMD Opteron and Intel EM64T)	Update 4	2.6.32-358.el6.x86_64
	Update 5	2.6.32-431.el6.x86_64
	Update 6	2.6.32-504.el6.x86_64
RHEL 7 X86_64 (AMD Opteron and Intel EM64T)		3.10.0-123.el7.x86_64
SLES 11 X86_64 (AMD Opteron and Intel EM64T)	SP2	3.0.13-0.27-default
	SP3	3.0.76-0.11-default
Community Enterprise Operating System (CentOS) X86_64 (AMD Opteron and Intel EM64T) (6.x)	Update 6.4	2.6.32-358.el6.x86_64
	Update 6.5	2.6.32-431.el6.x86_64
	Update 6.6	2.6.32-504.el6.x86_64
Community Enterprise Operating System (CentOS) X86_64 (AMD Opteron and Intel EM64T) (7.x)		3.10.0-123.el7.x86_64
Scientific Linux X86_64 (6.x)	Update 6.4	2.6.32-358.el6.x86_64
	Update 6.5	2.6.32-431.el6.x86_64
	Update 6.6	2.6.32-504.el6.x86_64
Scientific Linux X86_64 (7.x)		3.10.0-123.el7.x86_64
StackIQ Cluster Manager (Rocks+) HPC 3.1	RHEL 6.4	2.6.32-358.el6.x86_64
	CentOS 6.4	2.6.32-358.el6.x86_64
StackIQ Cluster Manager (Rocks+) HPC 3.2.1	RHEL 6.5	2.6.32-431.el6.x86_64
	CentOS 6.5	2.6.32-431.el6.x86_64
Platform HPC-4.1.1.1	RHEL 6.4	2.6.32-358.el6.x86_64
	RHEL 6.5	2.6.32-431.el6.x86_64

CPU model of Linux kernel can be identified by `uname -m` and `/proc/cpuinfo` shown in [Table 1-2](#)

Table 1-2. CPU Model of Linux Kernel

Model	uname	/proc/cpuinfo
EM64T	x86_64	Intel CPUs
Opteron*	x86_64	AMD CPUs

Note: Other combinations (such as i586 uname) are not currently supported.



1.6 Qualified Parallel File Systems

Lustre and IBM General Parallel File System (GPFS) listed below have been tested for use with this release of the Intel® OFED+ host software using the operating systems listed below:

- Lustre 2.5.3
 - RHEL 6.5
- IBM GPFS 3.5.0.18
 - RHEL 6.5

Refer to the *Intel® True Scale Fabric OFED+ Host Software User Guide* for the latest configuration recommendations for optimizing Lustre and GPFS performance with Intel® True Scale Fabric.

1.7 Intel Interface for NVIDIA GPUs

NVIDIA's CUDA parallel computing platform and programming models have been tested for use with this release of the Intel® OFED+ host software using the operating systems listed in [Table 1-3](#):

Table 1-3. NVIDIA's CUDA Tested with OFED+

Distributions	CUDA 6.5
RHEL 6.4	X
RHEL 6.5	X
RHEL 6.6	X
SLES 11 SP3	X

1.8 Compilers Supported

1.8.1 MPI

This release supports the following MPI implementations:

Table 1-4. MPI Compilers

MPI Implementation	Runs Over	Compiled With
Open MPI 1.8.1	PSM Verbs	GCC, Intel, PGI
MVAPICH version 1.2	PSM Verbs	GCC, Intel, PGI
MVAPICH2 version 1.8.1	PSM Verbs	GCC, Intel, PGI
Platform MPI 9.1	PSM Verbs	GCC (default)
Intel MPI version 5.0.1	TMI/PSM, uDAPL	Icc 15.0.0

1.8.2 MVAPICH, MVAPICH2 and Open MPI

MVAPICH, MVAPICH2 and Open MPI have been compiled for PSM to support the following versions of the compilers:



Table 1-5. MVAPICH, MVAPICH2 and Open MPI

Compiler name	Distro	Compiler Version
(GNU) gcc	RHEL6	gcc (GCC) 4.4.4 20100726 (Red Hat 4.4.4-13)
(GNU) gcc	RHEL7	gcc (GCC) 4.8.2 20140120 (Red Hat 4.8.2-16)
(GNU) gcc	SLES11	gcc (SUSE Linux) 4.3.2 [gcc-4_3-branch revision 141291]
(PGI) pgcc	RHEL6	pgcc 10.5-0 64-bit target on x86-64 Linux -tp nehalem-64
(PGI) pgcc	RHEL7	pgcc 14.4-0 64-bit target on x86-64 Linux -tp sandybridge
(PGI) pgcc	SLES11	pgcc 10.5-0 64-bit target on x86-64 Linux -tp nehalem-64
(Intel) icc	RHEL6	icc (ICC) 12.0.4 20110427
(Intel) icc	RHEL7	Version 14.0.2.144 Build 20140120
(Intel) icc	SLES11	icc (ICC) 12.0.4 20110427



1.9 Hardware Supported

Table 1-6 list the hardware supported in this release.

Table 1-6. Hardware Supported

HCA's
QLE7340
QLE7342
QME7342
QME7362
QMH7342
MHQH29-*
MHQH19-*
MHQH19B-XTR
MHQH29B-XTR
MHQH29B-XSR
MCX354A-QCAT
MCX353A-QCAT
NC543i (HP SL390 G7 in-built InfiniBand Host Channel Adapter)
CX-3 LOM down QDR
46M2199
46M2203

1.10 Installation Requirements

1.10.1 Software and Firmware Requirements

All Intel IB software on a given node must be at a compatible release level. Each distribution of the Intel IB Installation Wrapper will have a qualified and compatible version of each Intel IB Software component. Prior to installing Intel IB Software, any versions of the Silverstorm IB stack (and any other vendor's IB stack) must be uninstalled.

1.10.2 Installation Instructions

Note: An upgrade from Intel® True Scale Fabric OFED+ Host Software 7.1 or later may be performed in which case the installation will detect the existing installation and properly upgrade existing installed components and shall be removing components which are no longer supported.

Note: Any versions of an older IB stack must be uninstalled first. If you are downgrading from a newer IntelIB release to an older IntelIB or InfiniServ release, the newer IntelIB release must be stopped and uninstalled, prior to installing the older release. The newer IntelIB must be uninstalled using the `iba_config` tool or the `./INSTALL` tool provided with the newer IntelIB release.

Note: The installation process attempts to uninstall any existing 3rd party versions of OFED, however some packagings of OFED may not be completely uninstalled. If using a 3rd party OFED installation, it is recommended to uninstall it prior to installing IntelIB.



Note: The installation process attempts to uninstall any existing distribution versions of OFED, however some rpms included in the distribution packaging of OFED may not be completely uninstalled. It is recommended to uninstall any OFED rpms which come with the distribution prior to installing IntelIB.

Note: FastFabric may be used to install the IntelIB-Basic package on all nodes in the cluster.

Refer to the Intel® True Scale Fabric Software Installation Guide for installation information and procedures.

1.11 Changes for this Release

The following sections describe the changes that have been made to the Intel® True Scale Fabric Suite Software package between versions 7.2.2.0.8 and 7.3.0.0.26.

For detailed information about any of the previous releases listed, refer to the Release Notes for the specific version.

1.11.1 Changes to Software Components

Table 1-7 shows the new software components supported for the releases listed.

Table 1-7. Changes to Software Component Support

Release	Supported Software Added or Changed
7.3.1.0.12	Intel® True Scale Fabric OFED+ Host Software
7.3.0.0.26	Intel® True Scale Fabric OFED+ Host Software
7.2.2.0.8	Intel® True Scale Fabric OFED+ Host Software
7.2.1.1.22	Intel® True Scale Fabric OFED+ Host Software

1.11.2 Changes to Industry Standards Compliance

Table 1-8 shows each Basic OFED version that is supported and the Intel® OFED+ Releases that include each

Table 1-8. Changes to Industry Standards Compliance

Basic OFED Software Package Supported	Intel® OFED+ Host Software Package
Version 1.5.4.1	Versions 7.2.1.1.22, 7.2.2.0.8
Version 3.5-2	Version 7.3.0.0.26, 7.3.1.0.12

1.12 Product Constraints

The following is a list of product constraints for this release:

- The libgoto BLAS library included in the MPI sample applications for use by HPL is no longer supported by the developer and does not support arbitrary combinations of OS and CPU types.
 - Mixing AMD and Intel CPUs is not supported
 - AMD “Bulldozer” CPUs are not supported
 - Mixing different Linux distributions may not work reliably.

An alternative to libgoto is the ATLAS library, which must be manually compiled for each desired CPU type and Linux distribution. A sample version of ATLAS has been



provided in the same directory as the other sample MPI applications, and the latest version may be found online at <http://math-atlas.sourceforge.net/>

Examples of using the ATLAS library can be found in the hpl and hpl-2.0 directories with the other sample applications.”

- The version of Open MPI shipped with the Intel® True Scale Fabric Suite Software is incompatible with the Performance Application Programming Interface (“papi”) libraries optionally available in Red Hat Enterprise Linux version 6.x. If you have installed the optional papi RPMs and try to use FastFabric to recompile Open MPI on RHEL 6.x, you will first have to uninstall any installed version of papi if it is greater than version 3.x. Older versions of papi (for example, papi-3.x) are still compatible with the shipped version of OpenMPI
- All installation and uninstallation of Intel® OFED+ Host software package components must be performed using the `./INSTALL` or `iba_config` commands. If software is manually installed or uninstalled using other methods (RPM, other scripts, and so on), the installation on the system could become inconsistent and cause unreliable operation, in which case subsequent runs of `./INSTALL` or `iba_config` may make incorrect conclusions about the configuration of the system and consequently make incorrect recommendations. If the system becomes inconsistently configured, Intel recommends running the `./INSTALL TUI` and selecting `ReInstall` on all components. Once the re-installation has started, carefully review all prompts and choices.
- OFED SDP has not been qualified for this release. IPoIB is recommended for data transfers.

1.13 Product Limitations

The following is a list of product limitations for this release:

- Intel products will auto-negotiate with devices that utilize IBTA-compliant auto-negotiation. When attaching Intel products to a third-party device, the bit error rate is optimized if the third-party device utilizes attenuation-based tuning.

1.14 Other Information


The following is a list of “need to know” information for this release:

- The Dell PowerEdge M1000e Blade System has been updated with a new backplane (version 1.1), which requires different QME734x transmitter tuning settings. In order to facilitate proper transmitter settings, a new module parameter `qme_bp` has been added to the QIB driver. The `qme_bp` parameter should be set by the user to one of two values for the version of the installed backplane. A value of 0 means the backplane is version 1.0 and the default value of 1 means the backplane is version 1.1.

The module parameter can be set by editing the `/etc/modprobe.d/ib_qib.conf` file for nodes installed with RHEL operational environment or `/etc/modprobe.conf.local` file for nodes installed with SLES operational environment. The string “`qme_bp=value`” needs to be added to the “`options ib_qib ...`” option line. The value contained in this option, in combination with the chassis slot number, is used by the driver and support script to select the correct settings and program the QME734x transmitter.

To determine the version of the Dell backplane use the following procedure:

- 1.Login to the CMC of the Dell PowerEdge M1000e Blade System
- 2.Type “`getsysinfo`”
- 3.Under the Chassis Information, look for Chassis Midplane Version, refer to the following example:.



Output from getsysinfo:

Chassis Information:

```
System Model           = PowerEdgeM1000
System AssetTag       = 123abc
Service Tag           =
Chassis Name          = dell-cmc
Chassis Location      = [UNDEFINED]
Chassis Midplane Version = 1.1 ←
Power Status          = ON
```

- For information on Oracle* Remote Data Service (RDS) support, refer to [RAC Technologies Matrix for Linux Platforms](#)
- The OpenSHMEM effort (see <http://www.openshmem.org>) is defining a standardized API specification for SHMEM. Although it is premature to claim compliance, Intel® SHMEM aims to be compliant with the OpenSHMEM 1.0 specification. Intel provides a SHMEM API that is compatible with the OpenSHMEM 1.0 specification, other than any omissions or bugs documented in these release notes. If compliance with OpenSHMEM's passive progress statements are required, Intel® SHMEM's passive progress mechanism must be enabled.
- When using Mellanox HCAs, any changes to Virtual Fabrics (vFabrics) in the Fabric Manager, may require a reboot of the hosts with Mellanox HCAs. This limitation relates to the Mellanox HCAs not properly responding to changes to the Fabric Manager service level (SL). For some vFabric configuration changes, if the Fabric Manager SL changes or is mapped to a different Virtual Lane (VL) than previously, the Mellanox HCA can continue to use the previous VL. If that VL is presently disabled by the Fabric Manager, future uses of applications which use the Fabric Manager SL may hang or timeout because there are no VL Arbitration cycles for that VL. As a result, anytime vFabric configuration is changed, it is recommended to reboot all hosts with Mellanox HCAs so that the desired Quality of Service (QoS) configuration changes fully take effect. Any hosts with Intel® HCAs will not need to be rebooted.
Due to Mellanox HCAs not correctly handling changes to the Fabric Manager SL, Intel recommends that all the hosts using Mellanox ConnectX or ConnectX-2 HCAs be rebooted when used in a virtual fabric configuration.
- When Dispersive Routing is enabled, it allows packets sent using an MPI program run over PSM to take any one of several routes through a fabric, thus often increasing performance. The number of routes is determined by the value of 2 to the power of the Lid Mask Control setting (LMC). Because LMC defaults to 0, the default number of routes through the fabric is 2^0 or 1. LMC can be set as high as 3, allowing a total number of 2^3 or 8 routes through the fabric. Providing these additional routes can reduce fabric congestion, and thus improve performance. Dispersive Routing is supported when the Fabric Manager is used in the fabric. Dispersive Routing is not supported when using OpenSM.



- When running MVAPICH2, Intel recommends turning off RDMA fast path. To turn off RDMA fast path, specify `MV2_USE_RDMA_FAST_PATH=0` in the `mpirun_rsh` command line or set this option in the `parameter` file for `mvapich2`.
- The `ib_send_bw` benchmark, when run in UC mode, is written such that it will hang if even one packet is dropped.

1.14.1 Verbs Performance

Whenever possible, Intel recommends installing at least one True Scale HCA on a PCIe bus connected to the CPU 1 Socket for optimal verbs performance. For details on finding the correct PCIe slot, consult the documentation for your server chassis or motherboard.

1.15 Documentation

Table 1-9 lists the end-user documentation for the current release. All related documentation is available on the Intel® download site.

Documentation for Intel® Partners is available at the vendors web site.

Table 1-9. Related Documentation for this Release

Document Title
Intel® Hardware Documents
<i>Intel® True Scale Fabric Switches 12000 Series Hardware Installation Guide</i>
<i>Intel® True Scale Fabric Switches 12000 Series Users Guide</i>
<i>Intel® True Scale Fabric Switches 12000 Series CLI Reference Guide</i>
<i>Intel® True Scale Fabric Adapter Hardware Installation Guide</i>
Intel® OFED+ Documents
<i>Intel® True Scale Fabric Software Installation Guide</i>
<i>Intel® True Scale Fabric OFED+ Host Software User Guide</i>
<i>Intel® True Scale Fabric OFED+ Host Software Release Notes</i>
Intel® IFS Documents
<i>Intel® True Scale Fabric Suite FastFabric User Guide</i>
<i>Intel® True Scale Fabric Suite Fabric Manager User Guide</i>
<i>Intel® True Scale Fabric Suite FastFabric Command Line Interface Reference Guide</i>
<i>Intel® True Scale Fabric Suite Software Release Notes</i>
Intel® Fabric Viewer Documents
<i>Intel® True Scale Fabric Suite Fabric Viewer Online Help</i>
<i>Intel® True Scale Fabric Suite Fabric Viewer Release Notes</i>



2.0 System Issues for Release 7.3.1.0.12

2.1 Introduction

This section provides a list of the resolved and open issues in the OFED+ Host Software that were verified by this release. The list of open issues include a description and workaround.

2.2 Resolved Issues in this Release

Table 2-1 is a list of issues that are resolved in this and the previous two releases.

Table 2-1. Resolved Issues

Product	Release	Description
True Scale11/19/14	7.3.1.0.12	Open MPI 1.8.1 packaged with the True Scale software release version 7.3.0 no longer experiences a performance reduction as compared to Open MPI 1.6.5.
IFS/ FastFabric	7.2.2.0.8	Result of iba_verifyhosts for C-states are no longer misleading on SLES 11.
IFS/ HCA	7.2.2.0.8	OFED+ now works properly with SLES11SP3 kernel 3.0.93-0.8.



2.3 Known Issues

The subsections below catalog the known open issues for this release, as well as a description and a workaround by component.

2.3.1 Open Issues Table

Table 2-2 is the list of open issues for Release 7.3.1.

Table 2-2. Open Issues

Product/Component	Description	Workaround
IFS/ Rolls/Kits	The IFS 7.3 release supports PCM 4.1.1 (and higher). For PCM 4.1.1 (and higher), the GUI shows the version as 3.5.2 for the OFED+ and IFS kit. The release version is no longer displayed.	The Kit name (.bz2) file shows appropriate release version (7.3.0.0.26) after OFED version (3.5.2). kit-intel_ofed-3.5.2-7.3.0.0.26-rhels-6-x86_64.tar.bz2 kit-intel_ifs-3.5.2-7.3.0.0.26-rhels-6-x86_64.tar.bz2
IFS/ Install/Uninstall	The SRP Target contained in IFS for release 7.2 nodes cannot be installed on nodes running RHEL6.x.	Download and Install the required SRP drivers from http://sourceforge.net/projects/scst/files/srpt on designated SRP target node
IFS/ MPI	If LD_LIBRARY_PATH is exported inconsistently with the version of Open MPI being used, applications may build or run incorrectly. This issue can impact FastFabric tools that use MPI, rebuilding of mpi apps, or rebuilding Open MPI itself using the do_build or do_openmpi_build tools.	When using Open MPI, make sure PATH and LD_LIBRARY_PATH are not exported specifying a different path than the Open MPI path that is being used. The mpi-selector can configure a LD_LIBRARY_PATH for subsequent logins. Open MPI does not require the LD_LIBRARY_PATH to be set.
IFS/ MPI	MVAPICH2 jobs run between PCIe HCAs and Third Party or Intel HCAs, may not complete successfully. The test may abort with an ibv_post_recv error.	Any MVAPICH2 job attempted on fabrics with PCIe HCAs and Third Party or Intel HCAs, must zero the MV2_USE_SRQ environment variable as show in the example of the NAS CG benchmark: cd directory_containing_benchmark /usr/mpi/gcc/mvapich2-1.7/bin/ mpirun_rsh -np 4 -hostfile mpi_hosts \ MV2_USE_RDMA_FAST_PATH=0 MV2_USE_SRQ=0 ./cg.B.4
IB Third Party/ Other	Some applications using Platform MPI 8.1 or 8.2 over PSM exhibit poor scalability or worse performance at larger core counts.	Use the Platform MPI mpirun command with the -intra=nic option. This means that PSM's shared memory communications will be used instead of Platform MPI's native shared memory communications.



Table 2-2. Open Issues (Continued)

Product/ Component	Description	Workaround
Intel IB HCA/ SHMEM	SHMEM collective calls using PE subsets can hang. This problem does not affect collective calls using the entire PE set, which is the more common case. A collective call specifies the entire PE set when the parameters are set as follows: <pre>PE_start = 0</pre> <pre>logPE_stride = 0</pre> <pre>PE_size = num_pes()</pre>	None
Intel IB HCA/ SHMEM	The SHMEM API reductions on complex types are not implemented: <pre>shmem_complexd_sum_to_all</pre> <pre>shmem_complexf_sum_to_all</pre> <pre>shmem_complexd_prod_to_all</pre> <pre>shmem_complexf_prod_to_all</pre>	None
IFS/ IBAccess	When a port is down and does not have a LID assigned, <code>clear_p1stats</code> or <code>clear_p2stats</code> will fail against the given port	None
IFS/ IBAccess	When using vFabric, the OFED <code>saquery</code> command may use the wrong P-Key and timeout waiting for responses.	Intel recommends using the <code>iba_saquery</code> tool, which is included with IntelIB-Basic or IntelIB-IFS. <code>iba_saquery</code> will work properly when vFabric is configured.
IFS/ Open SM	When using <code>opensm</code> , after bouncing ports on a node, the port may not return to an active state for a period of time. As a result, commands that issue an SA query such as OFED's <code>saquery</code> command, or various FastFabric tools such as <code>iba_report</code> and <code>iba_saquery</code> , may hang waiting for the port to become active and the SA to respond.	Restart <code>opensm</code> . Intel recommends using the Intel® Fabric Manager, which has much greater resiliency and quicker handling of port state changes.
IFS/ IPoIB	When using vFabric to change an IPoIB application from Networking to Non-Networking, the IPoIB interface may remain in a running state.	After changing the application, restart the network services or bring the interface down/up to force IPoIB to re-query the SM and correct the situation.




Table 2-2. Open Issues (Continued)

Product/ Component	Description	Workaround
IFS/ Rolls/Kits	When using the Platform HPC and installing the Intel OFED+ or IFS kits, some messages appear in the system log file regarding unknown symbol and symbol version disagreement. These are due to the way and order in which the Intel OFED+ or IFS kits are installed and are resolved by the time OFED is started.	None
IFS/ Rolls/Kits	When the IFS is uninstalled and then reinstalled, some messages appear in Platform HPC's CFM log about issues with dependency resolution. These issues do not affect the operation of any IFS utilities including FastFabric and Fabric Manager.	None
True Scale	A continuous reboot occurs after loading a gPXE image.	Use Ctrl-B . This allows a user to enter into a command shell of gPXE. Once in the command shell, using the command autoboot starts a netboot connection using IPoIB.
IFS	While installing IntelIB-IFS.<osid>.7.3.0.0.26 on "osid=RHEL7" system, the following rpm packages may need to be removed (if installed) before running the INSTALL script: <pre> gemu-kvm gnome-boxes libvirt-daemon-kvm libmlx5 glusterfs-rdma </pre>	None.




Table 2-2. Open Issues (Continued)

Product/ Component	Description	Workaround
IFS	<p>The install/uninstall of IFS throws out warnings related to xxx remove failed. For example:</p> <pre>Uninstalling compat-rdma-3.5-3.10.0_123.e 17.x86_64.x86_64... warning: file /lib/modules/3.10.0-123.e17. x86_64/updates/drivers/infin iband/ulp/srp/ib_srp.ko:remo ve failed: No such file or directory</pre>	None. These are harmless warnings and can be ignored.
IB/Third Party	<p>OFED3.5-2 adds a newer version of <code>ibacm</code>, that causes <code>ibacm</code> service to load on boot time, causing additional usage count on the <code>qib</code> driver, preventing <code>openibd</code> to be restarted/stopped.</p>	Stop the <code>ibacm</code> service using the command <code>service ibacm stop</code> .
IFS	<p>If the <code>libedit</code> rpm is not present on a host, it should be installed prior to installing IFS. If the <code>libedit</code> libraries are not installed, the following error is displayed. Additionally, commands such as <code>iba_manage_switch</code> and <code>iba_switch_admin</code> would also fail with an error message similar to:</p> <pre>#/opt/iba/ib_tools/L8simlnx /opt/iba/ib_tools/L8simlnx: error while loading shared libraries: libedit.so.0: cannot open shared object file: No such file or directory</pre> <p>If the <code>libedit</code> rpm is not installed prior to the IFS installation, the capture operation would fail. However, installing <code>libedit</code> rpm would fix the issue. A reinstall of IFS is not needed.</p>	Install <code>libedit</code> rpm if not already present on a host.

§ §





Appendix A Performance Gain Conditions Test

The example shows how to determine if the following conditions are true:

1. The OS and CPUs support NUMA (Non-Uniform Memory Architecture) and NUMA node to I/O device binding.
2. Two HCAs connect to different PCIe root complexes which, in turn, connect to different NUMA nodes.

```
$ numactl --hardware

available: 2 nodes (0-1)

node 0 cpus: 0 1 2 3 4 5 6 7

...

node 1 cpus: 8 9 10 11 12 13 14 15

...
```

If `numactl --hardware` shows more than 1 NUMA node, then your OS supports NUMA.

To see whether your system supports NUMA node to IO device binding and whether your HCAs connect to different NUMA nodes, look at the files in the `/sys` directories to see if the `numa_node` field is populated correctly. The following steps indicate how to do this.

3. Change directory to `/sys/class/infiniband`

```
$ cd /sys/class/infiniband
```

4. List all files in the `/infiniband` directory in long format:

```
$ ls -la
```

This list the symbolic links to the HCA devices with the pci bus, slot and function number (in the following example 6:00.0 and 82:00.0):

```
lrwxrwxrwx 1 root root 0 Jul  9 11:24 qib0 -> ../../devices/pci0000:00/0000:00:03.0/0000:06:00.0/infiniband/qib0/
```

```
lrwxrwxrwx 1 root root 0 Jul  9 11:24 qib1 -> ../../devices/pci0000:80/0000:80:02.0/0000:82:00.0/infiniband/qib1/
```

5. Print the numa node id for the respective devices:

```
[infiniband]$ cat ../../devices/pci0000:00/0000:00:03.0/0000:06:00.0/numa_node
```



0

```
[infiniband]$ cat ../../devices/pci0000:80/0000:80:02.0/  
0000:82:00.0/numa_node
```

1

The HCAs are bound to the two NUMA nodes first shown in this appendix with the 'numactl --hardware" command: 0 and 1.

§ §