



# **Cornelis™ Omni-Path Express™ Fabric Software**

---

***Release Notes for V10.13.0***

January 2024  
Doc. No. A00206, Rev. 2.0

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Cornelis Networks products described herein. You agree to grant Cornelis Networks 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.

Cornelis Networks technologies may require enabled hardware, software, or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Cornelis, Cornelis Networks, Omni-Path, Omni-Path Express, and the Cornelis Networks logo belong to Cornelis Networks, Inc. Other names and brands may be claimed as the property of others.

Copyright © 2024 Cornelis Networks, Inc. All rights reserved.

# Table of Contents

<b>1. Overview of the Release</b>	<b>4</b>
1.1. Audience	4
1.2. Software License Agreement	4
1.3. If You Need Help	4
1.4. Product Improvements and Changes	4
1.5. Release Packages	5
1.6. Release Compatibility	5
1.7. Operating Systems	5
1.8. CUDA Support	6
1.9. MPI Libraries	7
1.9.1. Supported MPI Libraries	7
1.10. Supported Hardware	7
1.11. Switch Firmware	8
1.12. OPX Provider	8
1.13. Installation Requirements	9
1.13.1. Best Practices	9
1.13.2. Software and Firmware Requirements	9
1.13.3. OS RPMs Installation Prerequisites	9
1.13.4. Installation Instructions	12
1.14. Product Constraints	12
1.15. Product Limitations	13
1.16. Accelerated RDMA Information	13
1.17. hfi1 IRQ Balancing	14
<b>2. Issues</b>	<b>15</b>
2.1. Resolved Issues	15
2.1.1. Issues Resolved in this Release	15
2.1.2. Issues Resolved in Prior Releases	16
2.2. Open Issues	17
2.2.1. Third Party Open Issues	21
<b>3. Related Information</b>	<b>24</b>
3.1. Documentation Library	24
3.1.1. How to Search the Cornelis Omni-Path Express Documentation Set	26

## List of Tables

1. Release Packages	5
2. Supported Operating Systems	6
3. Supported MPI Libraries	7
4. Supported Hardware	8
5. RHEL 7.9 Distribution RPMs	10
6. RHEL 8.6, 8.7, 9.0, and 9.1 Distribution RPMs	10
7. SLES 15.3 and 15.4 Distribution RPMs	11
8. Issues Resolved in this Release	15
9. Issues Resolved in Prior Releases	16
10. Open Issues	17
11. Third Party Open Issues	21

# 1. Overview of the Release

This document provides a brief overview of the changes introduced into the Cornelis Omni-Path Express Fabric Software (OPX Fabric Software) in 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. Software License Agreement

The software provided is 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.

## 1.3. If You Need Help

Technical support for OPX products is available 24 hours a day, 365 days a year.

For additional details:

- Visit the Cornelis Support page at <http://www.cornelisnetworks.com/support>
- Contact [Cornelis Networks Customer Support](mailto:support@cornelisnetworks.com) at [support@cornelisnetworks.com](mailto:support@cornelisnetworks.com)

## 1.4. Product Improvements and Changes

The following improvements are included in this release:


- Support for Genoa Zen 4 CPU-based servers (See [Table 4 "Supported Hardware"](#))
- Support for Sapphire Rapids CPU-based servers (See [Table 4 "Supported Hardware"](#))
- Adds support for Red Hat Enterprise Linux (RHEL) 8.7 and 9.1 (See [Section 1.7 "Operating Systems"](#))
- Adds support for Rocky Linux 8.7 and 9.1 (See [Section 1.7 "Operating Systems"](#))
- Support for libfabric OFI version 1.18.0 (See [Section 1.12 "OPX Provider"](#))
- Support for Intel OneAPI 2022, Update 2
- Support for Open MPI version 4.1.4
- Support for CUDA Toolkit 12.1 (RHEL 8.6, 9.0 and SLES 15.3, 15.4 only) (PSM2 only)
- Support for MVAPICH2-2.3.7 (PSM2 only)
- The Embedded Subnet Manager (ESM) is deprecated

- Starting with RHEL 8.7 and 9.1, snoop capture functionality is deprecated

## 1.5. Release Packages

There are two OPX Fabric Software packages:

**Table 1. Release Packages**

Package	Includes
<b>BASIC</b> <ul style="list-style-type: none"> <li>• For compute, service, and login nodes</li> </ul>	Software that installs the following packages: <div>  <b>NOTE</b>            These packages will replace or update the packages included in the OS distribution.         </div> <ul style="list-style-type: none"> <li>• hfi1-firmware, libpsm2 (for RHEL) and libpsm2-2 (for SLES), hfi1-diagtools-sw</li> <li>• Open MPI and MVAPICH2. See <a href="#">Section 1.9 “MPI Libraries”</a> for details.</li> <li>• mpitests</li> <li>• mpi-selector</li> <li>• Open Fabrics Interface (OFI) libfabric</li> </ul>
<b>OPXS</b> <ul style="list-style-type: none"> <li>• For management nodes</li> </ul>	Includes the BASIC package plus: <ul style="list-style-type: none"> <li>• Fabric Manager               <ul style="list-style-type: none"> <li>– Allows comprehensive control of administrative functions using a mature Subnet Manager</li> <li>– Fabric Manager simplifies subnet, fabric, and individual component management, easing the deployment and optimization of large fabrics</li> </ul> </li> <li>• FastFabric Toolset               <ul style="list-style-type: none"> <li>– Enables rapid, error-free installation and configuration of OPX Host Software and management software tools</li> <li>– Simplified installation, configuration, validation, and optimization of HPC fabrics</li> </ul> </li> </ul>

## 1.6. Release Compatibility

This release is backward compatible with the most recent minor release version. For example, Release 10.N is backward compatible with Release 10.N-1 and Release 10.N-1.x.

## 1.7. Operating Systems

This release of the OPX Fabric Software supports the operating systems listed in the following table.



### NOTE

The 10.13.0 release contains release packages for RHEL 8.6, 8.7, 9.0, and 9.1 (and their Rocky Linux equivalents) as well as SLES 15.3 and 15.4. All other OS release packages are still available with the previous releases shown in [Table 2 “Supported Operating Systems”](#).

**NOTE**

For RHEL 7.9 install the RHEL 7.8 package the supports both RHEL 7.8 and RHEL 7.9 packages.

**NOTE**

OS support for PSM2 GPUDirect RDMA with CUDA is available on limited RHEL versions. Please see the note in [Section 1.8 "CUDA Support"](#).

**Table 2. Supported Operating Systems**

Release	Operating System	Update	Base Kernel Version
10.12.1	Red Hat Enterprise Linux (RHEL) 7.9 X86_64	7.9	3.10.0-1160.el7.x86_64
10.12.1	Red Hat Enterprise Linux (RHEL) 8.5 X86_64	8.5	4.18.0-348.el8.x86_64
10.12.1	Rocky Linux 8.5 (corresponds to RHEL 8.5)	8.5	4.18.0-348.el8.x86_64
10.13.0	Red Hat Enterprise Linux (RHEL) 8.6 X86_64	8.6	4.18.0-372.9.1.el8.x86_64
10.13.0	Red Hat Enterprise Linux (RHEL) 8.7 X86_64	8.7	4.18.0-425.3.1.el8.x86_64
10.13.0	Red Hat Enterprise Linux (RHEL) 9.0 X86_64	9.0	5.14.0-70.13.1.el9_0.x86_64
10.13.0	Red Hat Enterprise Linux (RHEL) 9.1 X86_64	9.1	5.14.0-162.6.1.el9_1.x86_64
10.13.0	Rocky Linux 8.6 (corresponds to RHEL 8.6)	8.6	4.18.0-372.9.1.el8.x86_64
10.13.0	Rocky Linux 8.7 (corresponds to RHEL 8.7)	8.7	4.18.0-372.9.1.el8.x86_64
10.13.0	Rocky Linux 9.0 (corresponds to RHEL 9.0)	9.0	5.14.0-70.13.1.el9_0.x86_64
10.13.0	Rocky Linux 9.1 (corresponds to RHEL 9.1)	9.1	5.14.0-162.6.1.el9_1.x86_64
10.13.0	SUSE Linux Enterprise Server (SLES) 15.3 X86_64	SP 3	5.3.18-57-default_k5.3.18_57-2369.x86_64
10.13.0	SUSE Linux Enterprise Server (SLES) 15.4 X86_64	SP 4	5.14.21-150400.22-default_k5.14.21_150400.22-2369.x86_64

## 1.8. CUDA Support

In general, PSM2 GPUDirect RDMA with CUDA is supported as shown below:

- CUDA Toolkit 12.1 is supported on RHEL 8.6, 9.0, and SLES 15.3 and 15.4

For information on compatible driver versions, refer to NVIDIA's [CUDA Compatibility](#), "CUDA Toolkit and Compatible Driver Versions" table.

**NOTE**

Customers requiring PSM2 GPU Direct RDMA support must use RHEL 8.6 or RHEL 9.0 as support for PSM2 GPUDirect RDMA with CUDA is not yet available on RHEL 8.7 and 9.1.

## 1.9. MPI Libraries

Cornelis recommends that you source the Intel MPI vars.sh or Intel OneAPI setvars.sh script when using Intel MPI . Refer to the Intel MPI Library documentation (<https://software.intel.com/content/www/us/en/develop/tools/oneapi/components/mpi-library.html>).



### NOTE

If you are not using Intel MPI, do not run these scripts. If you need to load the Intel compilers, use the compiler-specific vars.sh included in OneAPI.

### 1.9.1. Supported MPI Libraries

The table below lists the different MPI libraries supported by OPX Fabric Software with the corresponding version, fabric support, and compiler used. Note that the second column indicates if the MPI library is included in the OPX Fabric Software package.

**Table 3. Supported MPI Libraries**

MPI Implementation	Included in Basic Package?	Runs Over	Compiled With
Open MPI 4.1.4 <sup>1</sup>	Yes <sup>2</sup>	PSM2 and OPX provider	GCC
Open MPI 4.1.4-cuda <sup>1</sup>	Yes <sup>2</sup>	PSM2	GCC
MVAPICH2-2.3.7	Yes <sup>2</sup>	PSM2	GCC
Intel OneAPI	No	OFI	N/A
Intel OneAPI 2022 Update 2	No	OFI	N/A

#### NOTE:

1. Open MPI 4.X are known to work, however Open MPI 4.1.0 is not supported at this time.
2. PSM2 is compiled using the Intel Compiler (ICC).
3. MPI libraries that are included in the release and built with PSM2 support were built using the distro default compilers.



### NOTE

Refer to the *Cornelis Omni-Path Express Fabric Host Software User Guide* for setup information when using Open MPI with the SLURM PMI launcher and PSM2.

## 1.10. Supported Hardware

The following table lists the hardware supported in this release. The table does not include OEM-specific hardware, such as custom adapters and switches.



### NOTE

The PSM2 implementation has a limit of four (4) HFIs per server.

**Table 4. Supported Hardware**

Hardware	Description
Cornelis Omni-Path Express Host Fabric Interface Adapter 100HFA016 (x16)	Single Port PCIe Host Fabric Interface (HFI)
Cornelis Omni-Path Express Host Fabric Interface Adapter 100HFA018 (x8)	Single Port PCIe Host Fabric Interface (HFI)
Cornelis Omni-Path Express Accelerated Open Compute Project (OCP) 3.0 Host Fabric Interface Adapter 100HFP016KS (x16)	Single Port OCP Host Fabric Interface Adapter (OCP HFA)
2nd Gen AMD EPYC Processor	Rome Zen 2 CPU-based servers
3rd Gen AMD EPYC Processor	Milan Zen 3 CPU-based servers
4rd Gen AMD EPYC Processor	Genoa Zen 4 CPU-based servers
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 Processor	Skylake CPU-based servers
2nd Generation Intel Xeon Scalable Processor	Cascade Lake CPU-based servers
3rd Generation Intel Xeon Scalable Processor	Ice Lake CPU-based servers
4th Generation Intel Xeon Scalable Processor	Sapphire Rapids CPU-based servers

## 1.11. Switch Firmware

The following firmware is supported for OPX Switches:

- OPX Switch Firmware 10.8.x revision (managed and externally-managed switches)
- OPX Switch Firmware 10.7.x revision (managed and externally-managed switches)

Cornelis recommends using switch firmware 10.8.4 or greater.

Refer to the *Cornelis Omni-Path Express Fabric Switches Release Notes* for more information.

## 1.12. OPX Provider

The OPX provider is written to take full advantage of the libfabric acceleration features while running over existing and future OPX hardware.

The OPX provider started as a fork of the libfabric Blue Gene Q (BGQ) provider, having all of the hardware-specific parts re-written for OPX. Instruction counts and cache line footprint of most HPC operations show the OPX provider being lighter weight than PSM2 to the host software stack, which provides better overall performance.

The OPX provider provides excellent latency characteristics and message rates at smaller message sizes (under 16K message length). Bulk transfer via SDMA is present in the OPX provider giving some added performance for large messages. However, RDMA support is not implemented in the OPX provider at this time, so performance on large messages is still better with the PSM2 provider.

For further information on the OPX provider, refer to the *Cornelis Omni-Path Express OPX\_Provider Installation and Setup Application Note*.



## 1.13. Installation Requirements

This section provides installation requirements for this release.

Refer to the *Cornelis Omni-Path Express Fabric Software Installation Guide* for related software requirements and complete installation procedures. Refer to the *Cornelis Omni-Path Express Fabric Switches Hardware Installation Guide* for related firmware requirements.

### 1.13.1. Best Practices

Note the following Cornelis recommendations:

- Update to the latest versions of OPX firmware and software to obtain the most recent functional and security updates. These are found in the [Cornelis Customer Center](#).
- To improve security:
  - Administrators should log out users and disable multi-user logins prior to performing provisioning and similar tasks.
  - Update the default HTTPS certificate (refer to the *Cornelis Omni-Path Express Fabric Switches GUI User Guide*, “Updating the Certificate” for details).
  - Configure the `MgmtAllowed` setting and consider limiting access to port configuration changes by limiting access to Userspace Management Datagrams (UMADs). Refer to the *Cornelis Omni-Path Express Fabric Software Installation Guide*, “About User Queries Settings” for more information.

### 1.13.2. Software and Firmware Requirements

[Table 2 “Supported Operating Systems”](#) lists the operating systems supported by this release. For the required packages, refer to [Section 1.13.3 “OS RPMs Installation Prerequisites”](#).

### 1.13.3. OS RPMs Installation Prerequisites

Ensure that the following requirements are met before installing the software.

- Along with normal OS installation options, additional OS RPMs must be installed before you can install the OPX Fabric Software.
- Refer to the applicable section below to verify that all required RPMs are installed for the specific version of your OS distribution.
- Depending on the packages you choose, there may be additional prerequisites. For additional information, refer to the Release Notes for your specific release and installation type.

### 1.13.3.1. RHEL OS RPMs

The tables below list the RPMs for each supported RHEL release.



#### NOTE

- Some RHEL RPMs are available in a Server-Optional repository. Please install the RHEL Server-Optional version from Red Hat, which contains additional, required development packages.
- rdma-ndd is part of rdma-core.

**Table 5. RHEL 7.9 Distribution RPMs**

IB/OPA Centric	System Centric		Other/Generally Installed	Build Requirements
ibacm	atlas	libstdc++-devel	bash	bison
infinipath-psm	bc	ncurses-libs	irqbalance	expat-devel
libibumad	coreutils	numactl-libs	kernel	flex
libibverbs	createrepo	openssl	kmod	libnl3-devel
libnl3	expat	openssl-devel	libgcc	libpfm
librdmacm	expect	openssl-libs	perl	libuuid-devel
opensm-libs	gcc-gfortran	pciutils	perl-Getopt-Long	ncurses-devel
perftest	glibc	redhat-rpm-config	perl-PathTools	numactl-devel
qperf	kernel-devel	rpm-build	perl-Socket	opensm-libs
rdma-core	libatomic	sysfsutils	pkgconfig	openssl-devel (1.0.1 or higher)
rdma-core-devel	libgfortran	tcl	python	tcl-devel
	libgomp	tcsh	systemd	valgrind-devel
	libquadmath	zlib	systemd-libs	
	libstdc++			

**Table 6. RHEL 8.6, 8.7, 9.0, and 9.1 Distribution RPMs**

IB/OPA Centric	System Centric		Other/Generally Installed	Build Requirements
ibacm	atlas	libstdc++-devel	bash	autoconf
libibumad	bc	ncurses-compat-libs	irqbalance	automake
libibverbs	coreutils	ncurses-libs	kernel	bison
librdmacm	createrepo	numactl-libs	kernel-modules-extra	elfutils-libelf-devel
opensm-libs	expat	openssl	kmod	expat-devel
perftest	expect	openssl-devel	libgcc	flex
qperf	gcc-gfortran	openssl-libs	perl	gcc-c++.x86_64
rdma-core	glibc	pciutils	perl-Getopt-Long	kernel-abi-whitelists
rdma-core-devel	kernel-devel	redhat-rpm-config	perl-Socket	kernel-rpm-macros

IB/OPA Centric	System Centric		Other/Generally Installed	Build Requirements
	libatomic	rpm-build	pkgconf	libnl3-devel
	libgfortran	sysfsutils	python2	libpfm
	libgomp	tcl	python3 <sup>1</sup>	libtool
	libquadmath	tcsh	systemd	libuuid-devel
	libstdc++	zlib		ncurses-devel
				numactl-devel
				opensm-libs
				openssl-devel (1.1.1 or higher)
				tcl-devel

<sup>1</sup>Required for RHEL 9.0 and above

### 1.13.3.2. SLES OS RPMs

The tables below list the RPMs for each supported SLES release.



#### NOTE

Some SLES RPMs are available in SLES Software Development Kit (SDK). Please install the SLES SDK iso from SUSE, which contains additional, required development packages.

**Table 7. SLES 15.3 and 15.4 Distribution RPMs**

IB/OPX Centric	System Centric		Other/Generally Installed	Build Requirements
ibacm	bc	libnuma1	bash	bison
libibmad5	coreutils	libopenssl1_1	glibc	flex
libibumad3	createrepo_c	libopenssl-devel	grep	kernel-devel
libibverbs1	expect	libosmcomp3	irqbalance	libexpat-devel
libnl3-200	gcc-fortran	libquadmath0	kmod	libnuma-devel
libpsm_infinipath1	kernel-devel	libz1	libedit0	libopenssl-devel (1.0.1 or higher)
libquadmath0	kernel-syms	openssl	libgcc_s1	libuuid-devel
librdmacm1	libatomic1	rpm-build	libncurses5	ncurses-devel
libverbs1	libexpat1	tcl	libstdc++6	opensm-libs3
mpi-selector	libgfortran4	tcsh	perl	tcl-devel
opensm-devel	libgomp1	libudev-devel <sup>1</sup>	perl-base	valgrind-devel
opensm-libs3	libncurses6	systemd-devel <sup>2</sup>	pkg-config	
perftest			systemd	
qperf			udev	
rdma-core			python-base <sup>1</sup>	

IB/OPX Centric	System Centric		Other/Generally Installed	Build Requirements
rdma-core-devel			python3-base <sup>2</sup>	
rdma-ndd				
<sup>1</sup> Required for SLES15.3 only				
<sup>2</sup> Required for SLES15.4 only				

### 1.13.3.3. Example

The example below shows the installation of a group of OS RPMs on a SLES OS server. Not all required OS RPMs are included in this example and some OS RPMs in this example might not be needed.

```
# yum install libibmad libibverbs librdmacm qperf perftest
rdma infinipath-psm expat libstdc++-devel gcc-gfortran atlas tcl expect
tcsh sysfsutils pciutils bc libibumad libibumad-devel libibumad
libibumad-devel libibverbs-devel libibmad-devel librdmacm-devel ibacm-devel
openssl-devel libuuid-devel expat-devel infinipath-psm-devel valgrind-devel
libgnome libibverbs openssl-libs ncurses-devel hwloc hwloc-gui
```

Cornelis recommends that you build your own list of OS RPMs for installation.

### 1.13.4. Installation Instructions

There are two OPX Fabric Software packages:

- `CornelisOPX-OPXS.<distro>-x86_64.<version>.tgz` for the management node.
- `CornelisOPX-Basic.<distro>-x86_64.<version>.tgz` for compute, service, and login nodes.

The packages in the tgz file are RPMs. Installing individual RPMs is not supported in this release.



#### IMPORTANT

If you want to install the CUDA versions of the OPXS software using the `INSTALL` command, you must use the `-G` option.

Refer to the *Cornelis Omni-Path Express Fabric Software Installation Guide* for related software requirements and complete installation procedures. Refer to the *Cornelis Omni-Path Express Fabric Switches Hardware Installation Guide* for related firmware requirements.

## 1.14. Product Constraints

- Power class 2 AOCs are supported. You must use 10.5 (or newer) host software and 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.
- Some newer distros include an updated version of IPoIB which has an incompatibility with Fabric Managers prior to this release. This could result in the inability to establish IPoIB connections. Ensure that OS versions are consistent across the fabric.

## 1.15. Product Limitations

This release has the following product limitations:

- The embedded version of the Fabric Manager supports a maximum of 100 HFI ports involving less than 20 Switch ASICs. Calculate the number of Switch ASICs in your fabric as follows:
  - One ASIC per Cornelis Omni-Path Express Edge Switch 100 Series
  - Two ASICs per OPX Chassis Leaf module
  - Two ASICs per OPX Chassis Spine module
- The OPX Libfabric provider has a limit on the number of processes (or MPI ranks) per Omni-Path host fabric interface (HFI) card due to direct binding of processes to hardware contexts on the HFI. This limit may be encountered when running with a single HFI and one MPI rank per core on some high core count, multi-socket servers. For more detailed information about this limit, please contact [Cornelis Networks Customer Support](#).

If this issue impacts your workload, please consider:

- Using the legacy PSM2 fabric provider to allow multi-way sharing of HFI contexts across the entire set of CPU cores (which may cause resource contention under heavy load).
  - Executing your workload using fewer MPI ranks per server node, either using a reduced MPI rank count or in combination with an increased number of threads (such as pthreads or OpenMP threads) where applicable.
  - Deploying additional HFIs in the server to increase the total available hardware contexts.
- Performance Administration (PA) Failover should **not** be enabled with Fabric Managers 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.

## 1.16. Accelerated RDMA Information

Accelerated RDMA is a Verbs protocol extension to improve the performance of RDMA write and RDMA read operations on OPX hardware.

This extension improves the efficiency of large message transfers to provide performance benefits for storage protocols and other Verbs-based protocols. The benefits include increased achievable bandwidth with reduced CPU utilization. The Accelerated (or Token ID (TID)) RDMA protocol accelerates the OpenFabrics Alliance (OFA) Verbs API with no changes required to API consumers. The acceleration technique is performed by the host driver and the application running over the OFA Verbs API does not need to make any code changes.

Accelerated RDMA is off by default.

To enable it, add `cap_mask=0x4c09a09cbba` to the `/etc/modprobe.d/hfi1.conf` file. Instructions on how to do this are in the *Cornelis Omni-Path Express Fabric Performance Tuning User Guide*, “Setting HFI1 Driver Parameters” section.

## 1.17. hfi1 IRQ Balancing

`irqbalance` daemon and `udev` can be used to ensure optimal HFI IRQ => CPU core mapping. Optimal HFI IRQ => CPU core mapping may yield better performance. There are two parts needed for `hfi1` IRQ balancing to work:

- `irqbalance` policy script that is called for each IRQ belonging to an HFI. Tells `irqbalance` which CPU core(s) a given HFI IRQ should be processed on.
- An `udev.rules` file that tells `irqbalance` to rescan IRQs when an HFI appears or disappears from the system bus.

OPXS provides the `irqbalance` script but not the `udev.rules` file.

Here is an example `udev.rules` file to trigger `irqbalance` when an HFI appears or disappears:

```
ACTION=="add|change|remove", SUBSYSTEM=="hfi1_user", RUN+="/usr/bin/killall -SIGHUP irqbalance"
```

Important notes about the above rule:

- Keys off of `hfi1_user` virtual device instead of 'hfi1' module load. This rule will be triggered for each HFI added rather than just the initial `hfi1` module load/unload.
- Necessary to key off of `hfi1_user` since this is created at about the same time the IRQ affinity hints needed by our policy script are created. Keying off of 'hfi1' module load runs the risk of a race between when the affinity hints are populated and when `irqbalance` rescan is triggered.
- Sends `SIGHUP` to the `irqbalance` daemon rather than doing `systemctl restart irqbalance`. `SIGHUP` was the recommended method to rescan from our distro's documentation. `systemctl restart` did not work correctly in our testing.

See the *Cornelis Omni-Path Express Fabric Performance Tuning User Guide* for information about `irqbalance` and custom HFI IRQ affinity.

## 2. Issues

This section lists the resolved and open issues in the OPX Fabric Software.

### 2.1. Resolved Issues

#### 2.1.1. Issues Resolved in this Release

The following table lists issues that are resolved in this release.

**Table 8. Issues Resolved in this Release**

ID	Description	Resolved in Release
STL-64617 STL-64651 STL-64652	Performance issues have been observed on AMD Milan on RHEL 9.	10.13.0
STL-59919	Creating a ram disk using <code>dmsetup</code> on RHEL 7.7 will cause a kernel panic. The creation sequence is: <ol style="list-style-type: none"> <li><code>modprobe brd rd_nr=1 rd_size=33554432 max_part=1</code></li> <li><code>size=\$(blockdev --getsize /dev/ram0)</code></li> <li><code>dmsetup create mdevice0 --table "0 \${size} linear /dev/ram0 0"</code></li> </ol>	10.13.0
STL-65531	The PSM2 build date displays an incorrect timestamp on RHEL9 systems.	10.13.0

## 2.1.2. Issues Resolved in Prior Releases

The following table lists issues that were resolved in prior releases.

**Table 9. Issues Resolved in Prior Releases**

ID	Description	Resolved in Release
STL-63687	<p>PSM2 prior to version 11.2.230 has an issue where silent data errors will occur when all four of the following conditions are met:</p> <ul style="list-style-type: none"> <li>• OPA Congestion Control Architecture (CCA) is enabled in the FM configuration.</li> <li>• The application receives into an unaligned receive buffer, a buffer whose start address is not a 4B multiple.</li> <li>• The unaligned receive is for a large receive, one where PSM2 will use the TID receive mechanism.</li> <li>• One or more packets for the unaligned TID receive encounters congestion and is marked with the forward explicit congestion notification (FECN) bit.</li> </ul> <p>When these conditions are met, PSM2 is responsible for copying the payload to the destination memory rather than the TID hardware. But the PSM2 code did not take into account the unaligned buffer start. As a result, the data was copied to the wrong location in the receive buffer.</p> <p>This PSM2 issue has been resolved starting with PSM2 11.2.230 to correctly account for the unaligned start.</p>	10.12.0
STL-63732	<p>When IPoIB multicast (MC) join, requests were issued to an existing IPoIB MC group, via host port(s) configured to be limited-member(s) of a Virtual Fabric (VF), the FM/SA would not permit VF limited-members to join the MC group and would report an error message. This action taken by the FM/SA was incorrect. The IPoIB Architecture spec (RFC4392) states, that both MC full-members and MC limited-members may join an 'existing' MC group. The following is an example of the error message reported by the FM/SA:</p> <pre>sa_McMemberRecord_Set: Failing group join request for PortGUID 0x001175010170c3fc from opahsx74 hfi1_0 Port 1, PortGUID 0x001175010170c3fc, LID 0x00000001 because it is not a full member of virtual fabric with MGid 0xff12401b81010000:00000000ffffffff</pre>	10.12.0
STL-63901	<p>When the Fabric Manager SM has discovered and is configuring a switch, the SM must initialize the SCVL_nt Table of each neighboring port. This is done via a MAD request to the neighboring port. Currently, when the SM encounters a problem with that MAD request, the SM quarantines the switch port-0 (instead of quarantining the neighboring port that failed the request). This eventually causes the SM to report the disappearance of all end-nodes connected to the quarantined switch.</p>	10.12.0
STL-59413	<p>Using PSM2 with CUDA enabled on top of a non-CUDA enabled hfi1 driver may result in invalid behavior.</p>	10.12.1
STL-64558	<p>Removes redundant option from FastFabric TUI to build separate openmpi package for OPX Provider.</p>	10.12.1
STL-64578	<p>Using FastFabric in 10.12.0 on Rocky 8.6 - OpenMPI BUILD Fails</p>	10.12.1
STL-64591	<p>Using FastFabric in 10.12.0 on Rocky 8.6 - Mvapich2 BUILD Fails</p>	10.12.1



## 2.2. Open Issues

The following table lists the open issues for this release.

**Table 10. Open Issues**

ID	Description	Workaround
134819 (STL-49733)	In KNL-F EFI shell, the command <code>ifconfig -l</code> does not correctly display the IP address after being assigned via DHCP.	Launch a newer version of the EFI shell from the embedded shell.
136822 (STL-49728)	The 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.
139613 (STL-49732)	The Subsystem Vendor and Subsystem Device ID in the PCI configuration space of OPX 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.
139995 (STL-49724)	When installing OPX Fabric Software on a clean system (with no prior versions installed), the following error message may be displayed:  <code>cat: //etc/opa/version_delta: No such file or directory</code>	This message can be safely ignored. The installer is looking for an OPXS version file before it has been created on the system. The installation is not impacted.
143174 (STL-47003)	Due to a SLES 15 kernel setting, <code>hfdiags</code> cannot work while the HFI driver is loaded. The tool and driver are mutually exclusive.	Boot the kernel with <code>iomem=relaxed</code> . This turns off the resource exclusivity check.
STL-48921	On some Distributions, if you attempt to unload or restart the <code>hfi1</code> driver while the <code>ib_ipoib</code> driver is loaded, it may fail with a message similar to this:  <code>modprobe: FATAL: Module hfi1 is in use</code>	To avoid this issue, unload the <code>ib_ipoib</code> driver before unloading or restarting the <code>hfi1</code> driver.
STL-49313	On OS versions (RHEL 7.7 and newer and SLES 15, SLES 12.4 and newer) where it is available, configuring IPoIB module parameter <code>ipoib_enhanced=0</code> disables AIP on same node.	To run with AIP, ensure enhanced IPoIB is enabled ( <code>ipoib_enhanced=1</code> ).

ID	Description	Workaround
STL-57127	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>Irqlbalance 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> <p><b>NOTE:</b> Additional steps (like stopping Lustre) may be required on more complicated fabrics.</p>
STL-59663	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>Refer to the <i>Cornelis Omni-Path Express Fabric Performance Tuning User Guide</i>, 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 <a href="#">Cornelis Networks Customer Support</a> if you need further guidance.</p>
STL-60932	On 3rd Generation Intel Xeon Scalable Processors, IPoFabric small message latency may measure higher than expected depending on the tunings being used.	<p>Cornelis recommends that you use the Tuned latency-performance profile to significantly reduce IPoFabric small message latency. Tuned is a utility (for example, tuned-2.10.0-15.el8.noarch) that allows for dynamic and adaptive tuning of CPU behavior.</p> <p>Note that using Tuned latency-performance profile may have a negative performance impact to other bandwidth/message rate measurements.</p>
STL-61176	<p>Running perftest, the following text may be displayed:</p> <pre>Couldn't post send: qp 0 scnt=189 Failed to complete run_iter_bw function successfully</pre> <p><b>NOTE:</b> This failure is only seen on perftest and does not seem to impact any other upper layer protocols (ULPs). However, it is worse on RHEL 8.3.</p>	Disable TID RDMA.
STL-61215	In rare instances, the host may boot and the HFI device is not visible on the PCI bus, and is therefore not usable.	Reboot/power cycle the host.
STL-61244	Servers with 2nd Gen AMD EPYC Processor may experience an unexpected server reboot when running under load with a high number of processes per node.	Contact <a href="#">Cornelis Networks Customer Support</a> for more information.

ID	Description	Workaround
STL-61256	On 2nd Gen AMD EPYC Processors, AIP interrupts do not map optimally in all situations, especially when certain BIOS settings (such as NPS) are modified. This can impact IPoFabric performance.	Manually remap the AIP interrupts to give the best performance. Refer to the <i>Cornelis Omni-Path Express Fabric Performance Tuning User Guide</i> , "Driver IRQ Affinity Assignments" section for more details.  Contact <a href="#">Cornelis Networks Customer Support</a> if you need further guidance.
STL-61329	On RHEL 7.x servers with two HFIs installed (dual-rail), unbinding both HFIs from the hfi1 driver may result in a kernel warning trace.	None.
STL-61331	When installing on SLES 12.x from a yum/zypper repository, the installation may fail due to missing dependencies.	Do not use <code>-y</code> when performing the yum/zypper installation and select the option to ignore the dependencies.
STL-61527	Bidirectional host to host workloads using a GPU OPXS installation do not achieve maximum performance.	Use a non-GPU OPXS install for host-host workloads, or set <code>PSM2_CUDA=0</code> , <code>PSM2_GPUDIRECT=0</code> in the job environment for host-only workloads.
STL-61590	Open MPI (OMPI) has changed the support for one-sided operations since OMPI 4.0.5. Because of this change, OMPI is now launching both PSM2 MTL and OFI Byte Transfer Layer (BTL), which consumes more contexts than it actually needs. This can result in runtime warnings/errors on some high PPN jobs.	Run <code>-mca btl self,vader</code> . This string explicitly tells OMPI to use <code>-mca btl self,vader</code> as the BTL instead of OFI and prevents OMPI from consuming more contexts than it actually needs.  <b>Note:</b> Applications that require BTL to communicate may be adversely affected by this change. These are limited to a subset of MPI applications. Intel MPI benchmarks (IMB) will not be impacted.
STL-62119	Issues running the allreduce IMB benchmark using older versions of Intel MPI (ie: 2020.4) on AMD Milan systems have been observed.	Upgrade to use a more recent version of IMPI and IMB if this issue is seen.
STL-62604 STL-62605	Running certain IMB benchmarks using the the IMB source code packaged with oneAPI 2021.3 (or earlier) can result in benchmark crashes.	Upgrade to use a more recent version of IMB or upgrade to oneAPI 2021.4
STL-62613	The capmask for Accelerated IP (AIP) has changed starting in OPXS 10.11.1.0. This can cause AIP to be disabled if you are using the old cap mask referenced in the performance tuning guide, which is used to enable Accelerated RDMA.	The correct capmask is to use is now <code>0x4c09a09cbba</code> .
STL-63374	False credit loops may be reported by <code>opareport -o validatecreditloops</code> on a fabric configured in non-fat tree topology.	None, <code>opareport -o validatecreditloops</code> is only supported on fat tree topologies.
STL-64192	The OPXS repo based installation fails on RHEL 9 and SLES 15.4 due to missing dependencies for Python2 which has been deprecated in both OSes. OPXS packages a utility, <code>hfiutils</code> , which requires Python2. The utility is deprecated and will be removed in an upcoming OPXS release.	Perform installation using packaged <code>INSTALL</code> script

ID	Description	Workaround
STL-64193	In SLES15 distros, libpsm_infinipath1 is designated to obsolete libpsm2-compat even though the rpms can coexist. For some use cases (such as rebuilding MPI from source), both rpms must be installed. When a user installs OPXS using zypper, the process fails when installing libpsm2-compat due to this conflict.	The host software can be installed using the provided INSTALL script instead.
STL-64620	If using the OPXS repo-based installation on SLES15.4 and the suggested zypper workarounds are followed to remove the conflicting libpsm_infinipath1 package in favor of the libpsm2-compat package, a segfault may result when running libfabric due to a missing link to libpsm_infinipath.so.1.	<p>The libpsm2-compat package provides libpsm_infinipath.so.1 that is required by libfabric. To instruct libfabric on how to find the path to the library:</p> <ul style="list-style-type: none"> <li>• Add a config file for psm2-compat that points to the location of the libpsm_infinipath.so.1 provided by libpsm2-compat           <pre>cat /etc/ld.so.conf.d/psm2compat.conf /usr/lib64/psm2-compat</pre> </li> <li>• Load the new config file using ldconfig</li> </ul>
STL-64974	Failures have been observed when running the IMB-EXT "Window" benchmark with >20 ppn over the OPX Provider.	<p>To run the IMB-EXT benchmark with the OPX provider and higher PPN, reload the hfi1 module with an increased num_user_contexts value. For example, if the server has 40 physical cores and you want to run one IMB-EXT process per core, increase the number of hfi1 contexts to 80:</p> <pre># rmmod hfi1 # modprobe hfi1 num_user_contexts=80 # cat /sys/module/hfi1/parameters/num_user_contexts 80 80</pre> <p>The benchmark can now be run with #contexts/2 = 40 PPN, one for each CPU core.</p>
STL-65516	On AMD systems running RHEL 9.1, IMB benchmarks running over libfabric may fail with an "OFI get address vector map" error.	IOMMU should be disable in the bios. See <i>Cornelis Omni-Path Express Fabric Performance Tuning User Guide</i> for recommended IOMMU settings for AMD systems.
STL-65811	During installation, the following warning messages can be observed:  Not Found: ./CornelisOPX-Tools-FF.RHEL87-x86_64.10.13.0.0.6/RPMS/*/ opa-address-resolution-devel user  Not Found: ./CornelisOPX-Tools-FF.RHEL87-x86_64.10.13.0.0.6/RPMS/*/ opa-snapconfig user	None. These warnings can be safely ignored and installation will proceed successfully without impact to functionality.
STL-65954	OPX does not currently support separating traffic to specified Virtual Lanes.	None. This feature will be available in a future release.

ID	Description	Workaround
STL-66199	<p>On RHEL 9.x, the <code>opaswitchadmin</code> command used to update unmanaged switch firmware fails with the following error:</p> <pre>Invalid MAD: Invalid Session ID Error: Failed to get ini descriptors - status 3 Error: Failed to acquire node description - status -1 Invalid MAD: Invalid Session ID Error: Failed to get ini descriptors - status 3</pre>	Contact Cornelis Customer Support
STL-66659	<p>Accelerated IPoFabric (AIP) does not enable correctly on systems with high core count CPUs if the total number of CPU cores exceeds the limit of HFI driver contexts (160 for 100-Series Omni-Path). For example, on dual-socket AMD 96 core CPUs totaling 192 cores, a message similar to the following may be seen in <code>dmesg</code>:</p> <pre>hfi1_0: No receive contexts available for netdevs.</pre> <p>If AIP is required for IPoFabric performance, the <code>hfi1</code> driver must be configured to allow AIP to function properly.</p>	<p>Restrict the <code>hfi1</code> module parameter <code>num_user_contexts</code> to a lower value such as 149 (this works with the default <code>hfi1 krcvqs</code> value of 2). See the <i>Cornelis Omni-Path Express Fabric Performance Tuning User Guide</i>, "Setting HFI1 Module Parameters" for more detail. Contact <a href="#">Cornelis Networks Customer Support</a> for more technical detail or guidance for your exact system.</p>

### 2.2.1. Third Party Open Issues

The following table lists the third party open issues for this release.

**Table 11. Third Party Open Issues**

ID	Description	Workaround
129563 (STL-47095)	Memory allocation errors with MVAPICH2-2.1/Verbs.	<p><b>NOTE:</b> To avoid this issue, use MPIs over PSM2. 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>Increase the <code>lkey_table_size</code>: LKEY 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>Cornelis Omni-Path Express Fabric Performance Tuning User Guide</i>, "HFI1 Driver Module Parameters" chapter.</li> </ul>

ID	Description	Workaround
141273 (STL-46935)	The in-distro version of perftests has bugs.	Use the upstream version of perftest from <a href="https://github.com/linux-rdma/perftest">https://github.com/linux-rdma/perftest</a> .
STL-47571 STL-657	<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>hfi_userinit: assign_context command failed: Device or resource busy PSM2 can't open hfi unit: -1 (err=23)</pre> <p>In release 10.13, the oversubscription case may result in a segfault with the following error message:</p> <pre>*** Process received signal *** Signal: Segmentation fault (11) Signal code: (-6)</pre>	<p>Set <code>PSM2_MULTI_EP=0</code> (user environment variable) before or during job launch.</p> <p>For details, see the <i>Cornelis Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>.</p>
STL-57040	If Intel MPI 2019, Update 2 and OpenMPI are installed on the same cluster, OpenMPI jobs could link to the Intel MPI internal libfabric library after running Intel MPI <code>mpivars.sh</code> .	<p>Before running an OpenMPI job, perform the following to reset the libfabric environment variables that had been previously set by Intel MPI.</p> <pre>source &amp;lt;installdir&gt;/intel64/bin/ mpivars.sh -ofi_internal=0</pre>
STL-59449	<p>Occasionally, an unload of the hfi1 driver will produce a message similar to the following:</p> <pre>WARNING: CPU: 15 PID: 119002 at ../kernel/workqueue.c:4091 destroy_workqueue+0x61/0x240()</pre> <p>This message will be followed by a call trace showing the hfi1 unload.</p> <p>This issue is related to memory reclaim activity just before the unload.</p>	<p>None.</p> <p><b>NOTE:</b> This is a bug in the Linux kernel. Cornelis Networks is working with the Linux community to resolve this issue.</p>
STL-60633	When you install <code>gpfs.callhome-ecc-client-5.0.5-0</code> under SLES using zypper, you may receive an error code if there are no existing credentials.	<p>Ignore the error code.</p> <p><b>NOTE:</b> To avoid this error, you can touch an empty file <code>/usr/lpp/mmfs/bin/ecc/ECCBase/com.ibm.ws.webservices.thinclient_7.0.0.jar</code> to force zypper to exit with a success code. For example, run the following command before you install <code>gpfs.callhome-ecc-client-5.0.5-0</code>:</p> <pre>touch /usr/lpp/mmfs/bin/ecc/ECCBase/ com.ibm.ws.webservices.thinclient_7.0.0.jar</pre>
STL-61762	<code>osu_mbw_mr</code> can report bandwidth greater than line-rate in OSU micro benchmarks version 4.1 and higher.	None.

ID	Description	Workaround
LU-14733 (STL-61911)	<p>The lustre Lnet layer can emit the following message:</p> <pre>dk.opal63.llnl.gov.7:00000001:00020000:43.0 :1622598261.714620:0:129525:0: (brw_test.c:415:brw_bulk_ready()) BRW bulk READ failed for RPC from 12345-192.168.128.126@o2ib18: -103</pre> <p>The issue happens when an FRMR invalidation operation is failed in driver. RHEL 8.4 is exposed to the issue due to the release removing the older FMR kernel registration scheme. Lustre is sending an old previously used invalidate.</p>	<p>Use Lustre client release 2.12.8 or Build Lustre from source since the patches are in the Lustre source</p>
STL-63861	Large size IPv6 pings (greater than 15,975) may experience packet loss on SLES 15.4 and RHEL 9.0.	Contact <a href="#">Cornelis Networks Customer Support</a> or Linux distro provider for the latest information on the issue.
STL-64175	When Installing OPXS software, dracut warnings can be observed on the console. These warnings are not a result of OPXS software and do not impact functionality.	None. See SLES documentation for more information.

## 3. Related Information

### 3.1. Documentation Library

Go to the [Cornelis Customer Center](#) to download the publications from the Release Library.

Use the tasks listed in this table to find the corresponding document.

Task	Document Title	Description
Using the OPX documentation set	<i>Cornelis Omni-Path Express Fabric Quick Start Guide</i>	A roadmap to Cornelis' comprehensive library of publications describing all aspects of the product family. This document outlines the basic steps for installing your OPX cluster and ensuring it is operational.
Setting up an OPX cluster	<i>Cornelis Omni-Path Express Fabric Setup Guide</i>	Provides a high-level overview of the steps required to stage a customer-based installation of the OPX Fabric. Procedures and key reference documents, such as OPX user and installation guides, are provided to clarify the process. Additional commands and best known methods are defined to facilitate the installation process and troubleshooting.
Installing hardware	<i>Cornelis Omni-Path Express Fabric Switches Hardware Installation Guide</i>	Describes the hardware installation and initial configuration tasks for OPX Director Class Switches and OPX Edge Switches.
	<i>Cornelis Omni-Path Express Host Fabric Interface Installation Guide</i>	Contains instructions for installing the OPX HFI in an OPX cluster.
	<i>Cornelis Omni-Path Express Gateway Installation and Setup Guide</i>	Describes the hardware installation and set up tasks for the OPX Gateways.
Installing Host Software Installing HFI firmware Installing switch firmware (externally-managed switches)	<i>Cornelis Omni-Path Express Fabric Software Installation Guide</i>	Describes using a Text-based User Interface (TUI) to guide you through the installation process. You have the option of using command line interface (CLI) commands to perform the installation or install using the Linux distribution software.
Managing a switch using Chassis Viewer GUI Installing switch firmware (managed switches)	<i>Cornelis Omni-Path Express Fabric Switches GUI User Guide</i>	Describes the graphical user interface (GUI) of the OPX Chassis Viewer GUI. This document provides task-oriented procedures for configuring and managing the OPX Switch family. Help: GUI embedded help files
Managing a switch using the CLI Installing switch firmware (managed switches)	<i>Cornelis Omni-Path Express Fabric Switches Command Line Interface Reference Guide</i>	Describes the CLI task information for the OPX Switch family. Help: -help for each CLI



Task	Document Title	Description
Managing a fabric using FastFabric	<i>Cornelis Omni-Path Express Fabric Suite FastFabric User Guide</i>	Provides instructions for using the set of fabric management tools designed to simplify and optimize common fabric management tasks. The management tools consist of TUI menus and CLI commands.  Help: -help and man pages for each CLI. Also, all host CLI commands can be accessed as console help in the Fabric Manager GUI.
Managing a fabric using Fabric Manager	<i>Cornelis Omni-Path Express Fabric Suite Fabric Manager User Guide</i>	The Fabric Manager uses a well-defined management protocol to communicate with management agents in every OPX HFI and switch. Through these interfaces the Fabric Manager can discover, configure, and monitor the fabric.
	<i>Cornelis Omni-Path Express Fabric Suite Fabric Manager GUI User Guide</i>	Provides an intuitive, scalable dashboard and set of analysis tools for graphically monitoring fabric status and configuration. This document is a user-friendly alternative to traditional command-line tools for day-to-day monitoring of fabric health.  Help: Fabric Manager GUI embedded help files
Configuring and administering HFI and IPoIB driver Running MPI applications on OPX	<i>Cornelis Omni-Path Express Fabric Host Software User Guide</i>	Describes how to set up and administer the OPX HFI after the software has been installed. This document is for cluster administrators and Message-Passing Interface (MPI) application programmers.
Writing and running middleware that uses OPX software	<i>Cornelis Performance Scaled Messaging 2 (PSM2) Programmer's Guide</i>	Provides a reference for programmers working with the Performance Scaled Messaging 2 (PSM2) Application Programming Interface (API). The PSM2 API is a low-level user-level communications interface.
	<i>Cornelis Omni-Path Express OPX_Provider Installation and Setup Application Note</i>	Provides a reference for programmers working with the OPX provider.
Optimizing system performance	<i>Cornelis Omni-Path Express Fabric Performance Tuning User Guide</i>	Describes BIOS settings and parameters that have been shown to ensure best performance, or make performance more consistent, on the OPX Architecture. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.
Designing an IP or LNet router on OPX	<i>Cornelis Omni-Path Express IP and LNet Router Design Guide</i>	Describes how to install, configure, and administer an IPoIB router solution (Linux IP or LNet) for inter-operating between an OPX Fabric and a legacy InfiniBand fabric.
Building Containers for OPX Fabrics	<i>Building Containers for Cornelis Omni-Path Express Fabrics using Docker and Singularity Application Note</i>	Provides basic information for building and running Docker and Singularity containers on Linux-based computer platforms that incorporate OPX networking technology.

Task	Document Title	Description
Writing management applications that interface with OPX	<i>Cornelis Omni-Path Express Management API Programmer's Guide</i>	Contains a reference for programmers working with the OPX Architecture Management (OPAMGT) Application Programming Interface (API). The OPAMGT API is a C-API permitting in-band and out-of-band queries of the FM's Subnet Administrator and Performance Administrator.
Using NVM over Fabrics on OPX	<i>Configuring Non-Volatile Memory Express (NVMe) over Fabrics on Cornelis Omni-Path Express Application Note</i>	Describes how to implement a simple OPX Architecture-based point-to-point configuration with one target and one host server.
Learning about new release features, open issues, and resolved issues for a particular release	<i>Cornelis Omni-Path Express Fabric Software Release Notes</i>	
	<i>Cornelis Omni-Path Express Fabric Manager GUI Software Release Notes</i>	
	<i>Cornelis Omni-Path Express Fabric Switches Release Notes</i> (includes managed and externally-managed switches)	
	<i>Cornelis Omni-Path Express Fabric Unified Extensible Firmware Interface (UEFI) Release Notes</i>	
	<i>Cornelis Omni-Path Express Fabric Thermal Management Microchip (TMM) Release Notes</i>	
	<i>Cornelis Omni-Path Express Fabric Firmware Tools Release Notes</i>	

### 3.1.1. How to Search the Cornelis Omni-Path Express Documentation Set

Many PDF readers, such as Adobe Reader and Foxit Reader, allow you to search across multiple PDFs in a folder.

Follow these steps:

1. Download and unzip all the publications into a single folder.
2. Open Acrobat Reader and use **CTRL-SHIFT-F** to open the Advanced Search window.
3. Select **All PDF documents in...**
4. Select **Browse for Location** in the dropdown menu and navigate to the folder containing the PDFs.
5. Enter the string you are looking for and click **Search**.

Use advanced features to further refine your search criteria. Refer to your PDF reader Help for details.