Cornelis® Omni-Path Express® Fabric Software

Release Notes for V10.14.5

January 2025
Doc. No. A00615, Rev. 1.0

Table of Contents

1. Overview of the Release

This document provides a brief overview of the changes introduced into the Cornelis Omni-Path Express 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 is provided under license agreements and may contain third-party software under separate third-party licensing. Refer to the license files provided with the software for specific details.

1.3. If You Need Help

Technical support for all Cornelis 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 email: support@cornelisnetworks.com

1.4. Product Improvements and Changes

The following improvements are included in this release:

- Support for Red Hat Enterprise Linux (RHEL) 9.5
- Support for Rocky Linux 9.5
- 5th Generation Intel Xeon Scalable Processor (codename Emerald Rapids)
- 5th Gen AMD EPYC Processor (codename Turin)

1.5. Release Packages

The following table describes the two Omni-Path Express Fabric Software packages.

Table 1. Release Packages

Package	Includes
BASIC • For compute, service, and login nodes	Software that installs the following packages: **NOTE** These packages will replace or update the packages included in the OS distribution. • hfi1-firmware, libpsm2 (for RHEL) and libpsm2-2 (for SLES) • Open MPI and MVAPICH2. See MPI Libraries Support for details. • opa-basic-tools • mpitests • mpi-selector • Open Fabrics Interface (OFI) libfabric
Omni-Path Express Fabric Suite (OPXS) • For management nodes	Includes the BASIC package plus: • Fabric Manager – Allows comprehensive control of administrative functions using a mature Subnet Manager – Fabric Manager simplifies subnet, fabric, and individual component management, easing the deployment and optimization of large fabrics • FastFabric Toolset – Enables rapid, error-free installation and configuration of Omni-Path Express Host Software and management software tools – Simplified installation, configuration, validation, and optimization of HPC fabrics

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 Omni-Path Express Fabric Software supports the operating systems listed in the following table.

Table 2. Supported Operating Systems for Release 10.14.5

Operating System	Update	Base Kernel Version
Red Hat Enterprise Linux (RHEL) 9.5	9.5	5.14.0-503.11.1.el9_5x86_64
Rocky Linux 9.5	9.5	5.14.0-503.14.1.el9_4x86_64

1.8. CUDA Support

GPUDirect RDMA can be used with both PSM2 and the OPX Provider. The OPX Provider and PSM2 also allows for the use of GDRCopy and NCCL. Refer to section "Download and Install NVIDIA Software' in the *Cornelis Omni-Path Express Fabric Software Installation Guide* for installation instructions.

For information on compatible driver versions, refer to NVIDIA's CUDA Compatibility, "Minor Version Compatibility" section.

NOTE

The `-G` option must be used with the INSTALL command to enable Omni-Path Express components with NVIDIA-CUDA support. When loading, CUDA will look for libraries in the `/usr/local/cuda` folder.

1.9. ROCm Support

ROCm is an open software stack that includes programming models, tools, compilers, libraries, and runtimes for AI and HPC solution development on AMD GPUs.

ROCm can be used with the OPX Provider. Refer to section "Download and Install AMD Software" in the *Cornelis Omni-Path Express Fabric Software Installation Guide* for installation instructions.

Note that RHEL 9.5 and Rocky 9.5 requires a minimum ROCm version of 6.3.

IMPORTANT

There are two different versions of the Omni-Path Express software: CPU-only and GPU-enabled.

The GPU-enabled version comes in two types: NVIDIA and AMD.

Each of these three packages is designed to provide optimal performance with its supported hardware.

- If running CPU-only tasks, use the CPU version of the Omni-Path Express software.
- If running GPU tasks, you must use the relevant GPU version.

For information on ROCm, refer to AMD's ROCm Release Notes.

1.10. MPI Libraries Support

The following table lists the MPI libraries supported by Omni-Path Express Fabric Software with the corresponding version, fabric support, and compiler used. Note that the third column indicates if the MPI library is included in the Omni-Path Express Fabric Software package.

Table 3. MPI Implementation

MPI Implementation	Supported versions	In Package Tar file	Runs over	Compiled with
Open MPI	4.1.X	No	OFI MTL (PSM2/OPX Provider), PSM2 MTL	N/A
	4.1.6	Yes		GCC
	4.1.6-cuda	Yes		GCC
	5.0.3	No		N/A
	5.0.3-cuda	No		N/A
MVAPICH	2.3.7	Yes	PSM2	GCC
Intel MPI Library	2021 Update 3 and above	No	OFI (OPX Provider, PSM2)	N/A
Sandia OpenSHMEM	1.5.x	No	PSM2, OPX Provider	N/A

Cornelis recommends that you source the Intel MPI vars.sh or Intel OneAPI setvars.sh script when using Intel MPI. More information may be found at IntelMPI Library documentation.

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.

Refer to the following table for the list of Cornelis components.

Table 4. Components

Component	Git Repository	Tag
Kernel Driver	opa-distro-drivers	10.14.5.0_rhel9.5
libpsm2	opa-psm2	PSM2_12.0.1
OPX Provider	ofiwg/libfabric	v1.22.0*
Fabric Manager	opa-fm	v10.14.5.0.20
FastFabric	opa-ff	v10.14.5.0.20
*Plus additional patches; contact Cornelis Networks Customer Support for more information.		

1.11. 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.

Table 5. 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
4th Gen AMD EPYC Processor	Genoa Zen 4 CPU-based servers
5th Gen AMD EPYC Processor	Turin Zen 5 CPU-based servers
AMD Instinct MI300A Accelerated Processing Unit (APU)	AMD Instinct APUs
AMD Instinct MI series Graphic Processing Units (GPUs)	AMD Instinct MI100, 200, and 300 series GPUs
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
5th Generation Intel Xeon Scalable Processor	Emerald Rapids CPU-based servers

Hardware	Description
NVIDIA Graphics Processor Units (GPUs)	NVIDIA Tesla V100 and later

NOTE

The PSM2 provider has been tested with up to eight (8) HFIs per server when running in non-multirail mode (`PSM2_MULTIRAIL=0`) and four (4) HFIs per server when running with `PSM2_MULTIRAIL!=0`.

1.12. OPX Libfabric Provider

The OPX Libfabric Provider (hereafter called the OPX Provider) is written to take full advantage of the libfabric acceleration features while running over existing and future Omni-Path Express hardware.

The OPX Provider delivers excellent latency characteristics and message rates at smaller message sizes (under 16K message length). Bulk transfer using SDMA is present in the OPX Provider giving some added performance for large messages. RDMA support is implemented in the OPX Provider using the environment variable `FI_OPX_EXPECTED_RECEIVE_ENABLE=1`.

Use `FI_LOG_LEVEL=trace FI_LOG_SUBSYS=core` to print the file location info and OPX Provider library file in use. The output will look similar to:

```
Using opx Provider: Library file location is *.so file location
```

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.

There are two Omni-Path Express 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

The `-G` option must be used with the INSTALL command to install Omni-Path Express components with NVIDIA-CUDA support.

IMPORTANT

There are two different versions of the Omni-Path Express software: CPU-only and GPU-enabled.

The GPU-enabled version comes in two types: NVIDIA and AMD.

Each of theses three packages is designed to provide optimal performance with its supported hardware.

- If running CPU-only tasks, use the CPU version of the Omni-Path Express software.
- If running GPU tasks, you must use the relevant GPU version.

1.13.1. Best Practices

Note the following Cornelis recommendations:

- Update to the latest versions of Omni-Path Express 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.

1.14. Product Constraints

Power class 2 AOCs are supported. This support requires a minimum Host software version of 10.5 and a minimum UEFI version of 1.5. Integrated HFI (-F) requires a specific BIOS level to support power class 2 AOC. Contact your BIOS vendor for more information.

1.15. Product Limitations

This release has the following product limitations:

- On systems where the total number of CPU cores exceeds the limit of the HFI Driver contexts (160 for the HFI), Accelerated IPoFabric (AIP) does not enable correctly. For example, on dual-socket 96 core CPUs totaling 192 cores, a message similar to the following may be seen in dmesg: `hfi1_0: No receive contexts available for netdevs`. If AIP Is required for IPoFabric performance, the hfi1 driver must be configured to allow AIP to function properly by restricting the hfi1 module parameter `num_user_contexts` to a lower value such as 149 (this works with the default `hfi1 krcvqs` value of 2). See the *Cornelis Omni-Path Express Fabric Performance Tuning User Guide*, "Setting HFI1 Module Parameters" for more detail. Contact Cornelis Networks Customer Support for more technical detail or guidance for your exact system.
- The OPX Libfabric provider has a limit on the number of processes (or MPI ranks) per 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 details about this limit, contact Cornelis Networks Customer Support.

  If this issue impacts your workload, consider:

  - Using the PSM2 fabric provider to allow sharing contexts between ranks (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 with different 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.
- Cannot `ssh` to OPA100 switches with newer operating systems. Users should explicitly re-enable SHA-1 algorithms. Contact Cornelis Networks Customer Support for more information.
- The Subsystem Vendor and Subsystem Device ID in the PCI configuration space of discrete HFI cards may not indicate the correct OEM vendor and device (STL-49732). 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. Refer to TA0053 available in the Cornelis Customer Center.

2. Issues

This section lists the resolved and open issues in the Omni-Path Express Fabric Software.

2.1. Issues Resolved in this Release

There are no resolved issues in this release.

2.2. Open Issues

The following table lists the open issues for this release.

Table 6. Open Issues

ID	Description	Workaround
STL-61527	Host-based workloads using a GPU Omni-Path Express Fabric Suite installation do not achieve maximum performance.	Use a non-GPU Omni-Path Express Fabric Suite installation for host-to-host workloads, or use a non-GPU-enabled libfabric at application runtime.

ID	Description	Workaround
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 `-mca btl self,vader`. This string explicitly tells OMPI to use `-mca btl self,vader` as the BTL instead of OFI and prevents OMPI from consuming more contexts than it actually needs. **Note:** 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-64620	If using the Omni-Path Express Fabric Suite 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`.	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: • Add a configuration file for `psm2-compat` that points to the location of the `libpsm_infinipath.so.1` provided by `libpsm2-compat`: `cat /etc/ld.so.conf.d/ psm2compat.conf /usr/lib64/psm2-compat` • Load the new configuration file using ldconfig.
STL-64974	Failures have been observed when running the IMB-EXT "Window" benchmark using the OPX Provider and ppn greater than (num_user_contexts/2).	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. See the *Cornelis Omni-Path Express Fabric Performance Tuning User Guide*, section "HFI1 Driver Module Parameters" for instructions. The benchmark can now be run with #contexts/2 = 40 PPN, one for each CPU core.

ID	Description	Workaround
STL-65757	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: `hfi_userinit: assign_context command failed: Device or resource busy PSM2 can't open hfi unit: -1 (err=23)In release 10.13, the oversubscription case may result in a segfault with the following errormessage: *** Process received signal *** Signal: Segmentation fault (11) Signal code: (-6)`	Set PSM2_MULTI_EP=0 (user environment variable) before or during job launch.

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-66199	On RHEL 9.x, the `opaswitchadmin` command used to update unmanaged switch firmware fails with the following error:
	`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`
		Perform the following:
		1. Uninstall the OPXS version.
		2. Install the inbox `opa` rpms that is part of the OS.
		`yum install opa-*`
		3. Use `opaswitchadmin` to update the unmanageed switch.
		4. Reinstall the OPXS standalone release.

STL-66910	Executables compiled with the `-pg` flag to enable extra profile information for debuggers may cause failures when attempting to run with Omni-Path and the PSM2 library.	Download latest release of PSM2 from GitHub (currently here: PSM2). In psm_config.h, change `#define HAL_CONTEXT_OPEN_RETRY_MAX 3` to `#define HAL_CONTEXT_OPEN_RETRY_MAX 100`, and compile PSM2. Use this version of PSM2 by setting the path to the library with LD_LIBRARY_PATH. Contact Cornelis Networks Customer Support for more information.
STL-67178	NVIDIA GPU support is unavailable when the NVIDIA driver is not loaded prior to loading the hfi1 driver.	Review boot configurations to ensure NVIDIA GPU driver is available prior to running the `CornelisOPX-OPXS` INSTALL script or when loading the hfi1 driver.

ID	Description	Workaround
STL-67821	On AMD CPUs with SMT enabled and IOMMU disabled, the hfi1 driver does not allocate one context per physical CPU core as intended. The PSM2 provider will run in context-sharing mode, but the OPX Provider will not run due to lack of context sharing support.	Disable SMT or run with the PSM2 provider in context sharing mode.
STL-68208	When using IPoIB in connected mode with high traffic and bouncing links at the same time, errors may occur and, in rare circumstances, result in a kernel panic.	Use IPoIB in datagram mode or use connected mode but do not bounce links unnecessarily.
STL-69062	The PSM2 error `Unhandled error in TID Update: Bad address` may occur when running applications that use CUDA Unified Memory.	None. GPUDirect does not support CUDA Unified Memory. Contact your application vendor to inquire about not using CUDA managed memory for psm2 operations.
STL-69509	Unexpected behavior including segmentation faults have been observed while running one-sided MPI applications with the OPX Libfabric Provider.	Use the PSM2 provider.
STL-70448	Mixing use of CUDA-managed and unmanaged memory may result in a crash.	The OPX Provider does not support mixing CUDA-managed and unmanaged memory at this time. • If your GPU application uses both CUDA-managed and unmanaged memory, use PSM2. • If your GPU application uses CUDA-managed memory, use the non-GPU build of libfabric and the OPX Provider. • If your GPU application only uses unmanaged memory, use the GPU build of libfabric and the OPX Provider.
STL-70552	Rebuilding Open MPI or mvapich2 with FastFabric utilities (`do_openmpi_build` and `do_mvapich2_build`) will return an error if using the Intel icx compiler.	These utilities are compatible only with the Intel ICC compiler. Rebuild Open MPI or mvapich2 using the ICC compiler, or rebuild manually.
STL-70990	Errors detecting libfabric occur when configuring Open MPI or other software on AMD GPU-enabled nodes.	Make sure libhsa-runtime64.so.1 and all other libraries are discovered with `ldd /usr/lib64/libfabric.so`. If not, consult your system administrator or add the location to your LD_LIBRARY_PATH, such as `export LD_LIBRARY_PATH=/opt/rocm-6.2.2/lib/:$LD_LIBRARY_PATH` before configuring.

ID	Description	Workaround
STL-71335	On systems using OpenSSH 9.4 or newer, `opacmdall` and related Fast Fabric utilities return an error when attempting to execute remote commands via SSH.	For simultaneous SSH connectivity to multiple servers, use 'pdsh' or another cluster management utility. For assistance with Fast Fabric utilities, contact Cornelis Networks Customer Support.
STL-71697	For non-GPU installations, the repo-based install method erroneously installs AMD GPU kernel RPMs.	Use the traditional ./INSTALL command method to install. Contact Cornelis Networks Customer Support for additional information.

2.3. Interoperability Issues

The following table lists the open interoperability issues for this release.

Table 7. Interoperability Issues

ID	Description	Workaround
129563 STL-47095	Memory allocation errors with MVAPICH2-2.1/ Verbs.	**NOTE:** To avoid this issue, use MPI over PSM2. If you are using MPIs over Verbs, the following workaround is required: • When running MVAPICH2 jobs with a large number of ranks (for example, > 36 ranks but ≤ 72 ranks), you must set the following parameters in `/etc/security/limits.conf`: – hard memlock unlimited – soft memlock unlimited • Increase the lkey_table_size:LKEY table size in bits (2^n, where $1 \leq n \leq 23$) from its default of 16 to 17. For instructions on setting module parameters, refer to the *Cornelis Omni-Path Express Fabric Performance Tuning User Guide*, "HFI1 Driver Module Parameters" chapter.
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 `ipoib_enhanced=0` disables AIP on same node.	To run with AIP, ensure enhanced IPoIB is enabled (`ipoib_enhanced=1`).
STL-61762	osu_mbw_mr can report bandwidth greater than line-rate in OSU micro benchmarks version 4.1 and higher.	None.
STL-67113	Timeouts, hangs, and memory allocation issues can occur with MPI jobs if the `memlock` limit is not sufficient.	Set the following parameters in `/etc/security/limits.conf` • `hard memlock unlimited` • `soft memlock unlimited`

ID	Description	Workaround
STL-68521	Running `modprobe -r hfi1` results in error message, `modprobe: FATAL: Module nvidia is in use.`	This issue concerns NVIDIA's module unloading and is outside of Cornelis' control. It does not prevent unloading/reloading hfi1. This message can be ignored.
STL-70076	OPX libfabric provider initialization failure with AWS OFI NCCL v1.10.0.	Customers who want to use the AWS OFI NCCL plug-in with the OPX libfabric provider for NVIDIA's NCCL-based applications should use AWS OFI NCCL v1.9.2.
STL-71651	GDR copy compile errors in GDRCopy driver may be encountered when using RHEL 9.5 with GDRCopy versions older than 2.4.3.	Update GDRCopy to version 2.4.3 or newer. This is available for download at https://github.com/NVIDIA/gdrcopy/releases.

3. Related Information

3.1. Documentation Library

Go to the Cornelis Customer Center to download the publications from the Download Library.

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

Task	Document Title	Description
Using the Omni-Path Express documentation set	*Cornelis Omni-Path Express Fabric Quick Start Guide*	Provides a roadmap to Cornelis' comprehensive library of publications describing all aspects of the product family. This document outlines the basic steps for installing your Omni-Path Express cluster and ensuring it is operational.
Setting up an Omni-Path Express cluster	*Cornelis Omni-Path Express Fabric Setup Guide*	Provides a high-level overview of the steps required to stage a customer-based installation of the Omni-Path Express Fabric. Procedures and key reference documents, such as Omni-Path Express 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	*Cornelis Omni-Path Express Fabric Switches Hardware Installation Guide*	Describes the hardware installation and initial configuration tasks for Omni-Path Express Director Class Switches and Omni-Path Express Edge Switches.
	Cornelis Omni-Path Express Host Fabric Interface Installation Guide	Contains instructions for installing the Omni-Path Express HFI in an Omni-Path Express cluster.

Task	Document Title	Description
	Cornelis Omni-Path Express Gateway Installation and Setup Guide	Describes the hardware installation and set up tasks for the Omni-Path Express Gateways.
Installing Host Software Installing HFI firmware Installing switch firmware (externally-managed switches)	*Cornelis Omni-Path Express Fabric Software Installation Guide*	Describes using command line interface (CLI) commands to perform the installation or install using the Linux distribution software. You have the option of using a Text-based User Interface (TUI) to guide you through the installation process.
Managing a switch using Chassis Viewer GUI Installing switch firmware (managed switches)	*Cornelis Omni-Path Express Fabric Switches GUI User Guide*	Describes the graphical user interface (GUI) of the Omni-Path Express Chassis Viewer GUI. This document provides task-oriented procedures for configuring and managing the Omni-Path Express Switch family. Help: GUI embedded help files
Managing a switch using the CLI Installing switch firmware (managed switches)	*Cornelis Omni-Path Express Fabric Switches Command Line Interface Reference Guide*	Describes the CLI task information for the Omni-Path Express Switch family. Help: -help for each CLI
Managing a fabric using FastFabric	*Cornelis Omni-Path Express Fabric Suite FastFabric User Guide*	Provides instructions for using the set of fabric management tools designed to simplify and optimize common fabric management tasks. The management tools consist of CLI commands and TUI menus. 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	*Cornelis Omni-Path Express Fabric Suite Fabric Manager User Guide*	Describes the Fabric Manager, a well-defined management protocol that communicates with management agents in every Omni-Path Express HFI and switch. Through these interfaces the Fabric Manager can discover, configure, and monitor the fabric.
	Cornelis Omni-Path Express Fabric Suite Fabric Manager GUI User Guide	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 Omni-Path Express	*Cornelis Omni-Path Express Fabric Host Software User Guide*	Describes how to set up and administer the Omni-Path Express HFI after the software has been installed. This document is for cluster administrators and Message-Passing Interface (MPI) application programmers.

Task	Document Title	Description
Writing and running middleware that uses Omni-Path Express software	*Cornelis Performance Scaled Messaging 2 (PSM2) Programmer's Guide*	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.
	Cornelis Omni-Path Express OPX Provider Installation and Setup Application Note	Provides a reference for programmers working with the OPX Provider.
Optimizing system performance	*Cornelis Omni-Path Express Fabric Performance Tuning User Guide*	Describes BIOS settings and parameters that have been shown to ensure best performance, or make performance more consistent, on the Omni-Path Express 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 Omni-Path Express	*Cornelis Omni-Path Express IP and LNet Router Design Guide*	Describes how to install, configure, and administer an IPoIB router solution (Linux IP or LNet) for inter-operating between an Omni-Path Express Fabric and a legacy InfiniBand fabric.
Building Containers for Omni-Path Express Fabrics	*Building Containers for Cornelis Omni-Path Express Fabrics using Docker and Singularity Application Note*	Provides information for building and running Docker and Singularity containers on Linux-based computer platforms that include Omni-Path Express networking technology.
Writing management applications that interface with Omni-Path Express	*Cornelis Omni-Path Express Management API Programmer's Guide*	Contains a reference for programmers working with the Omni-Path Express 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 Omni-Path Express	*Configuring Non-Volatile Memory Express (NVMe) over Fabrics on Cornelis Omni-Path Express Application Note*	Describes how to implement a simple Omni-Path Express 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	*Cornelis Omni-Path Express Fabric Software Release Notes*	
	Cornelis Omni-Path Express Fabric Manager GUI Software Release Notes	
	Cornelis Omni-Path Express Fabric Switches Release Notes (includes managed and externally-managed switches)	
	Cornelis Omni-Path Express Fabric Unified Extensible Firmware Interface (UEFI) Release Notes	
	Cornelis Omni-Path Express Fabric Thermal Management Microchip (TMM) Release Notes	
	Cornelis Omni-Path Express Fabric Firmware Tools Release Notes	

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.