



# Intel<sup>®</sup> Ethernet Fabric Suite Software

Release Notes for 12.1.0.0

---

*Rev. 1.0*

*June 2025*



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

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

All 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.

Intel technologies may require enabled hardware, software or service activation.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

Copyright © –2025, Intel Corporation. All rights reserved.

## Contents

---

<b>1.0 Overview of the Release.....</b>	<b>5</b>
1.1 Audience.....	5
1.2 Software License Agreement.....	5
1.3 If You Need Help.....	5
1.4 Updated Features.....	5
1.5 Release Packages.....	7
1.6 Intel Hardware.....	7
1.7 Release Compatibility.....	8
1.8 Operating Systems.....	9
1.9 Intel GPU Support.....	10
1.10 CUDA Support.....	11
1.11 NVIDIA DOCA OFED Support.....	11
1.12 Kernel Modules for Rendezvous Driver.....	12
1.13 Network File Systems.....	13
1.14 Compilers and Libraries.....	14
1.14.1 Supported MPI Libraries.....	14
1.15 Document Versions.....	15
1.16 Installation Requirements.....	15
1.16.1 Best Practices.....	15
1.16.2 Software and Firmware Requirements.....	15
1.16.3 OS Installation Prerequisites.....	15
1.16.4 Installation Instructions.....	19
<b>2.0 Issues.....</b>	<b>20</b>
2.1 Resolved Issues.....	20
2.1.1 Issues Resolved in this Release.....	20
2.1.2 Issues Resolved in Prior Releases.....	21
2.2 Open Issues.....	23
2.2.1 Third-Party Open Issues.....	24
<b>3.0 Related Information.....</b>	<b>26</b>
3.1 Intel® Ethernet Fabric Suite Documentation Library.....	26
3.1.1 How to Search the Intel® Ethernet Fabric Suite Documentation Set.....	27

## Tables

1	Supported Hardware.....	8
2	Supported Intel® Ethernet E810 Network Adapter Software.....	8
3	Supported Operating Systems.....	9
4	Supported OS by CPU.....	9
5	Distro - Intel® GPU Software Support Matrix.....	10
6	Distro - CUDA Software Support Matrix.....	11
7	Distro - DOCA Software Support Matrix.....	12
8	Kernel Modules for Rendezvous Driver.....	12
9	oneAPI Intel GPU Kernel Modules for Rendezvous Driver.....	13
10	CUDA Kernel Modules for Rendezvous Driver.....	13
11	Compiler Versions and Distributions.....	14
12	Intel® EFS Components and Compiler.....	14
13	Supported MPI Libraries.....	14
14	Supported Document Versions.....	15
15	RHEL 8.x, 9.x, 10.x Distribution RPMs.....	16
16	SLES15.x Distribution RPMs.....	17
17	Ubuntu 22.04 Distribution DEBs.....	18
18	Issues Resolved in this Release.....	20
19	Issues Resolved in Prior Releases.....	21
20	Open Issues.....	23
21	Third-Party Open Issues.....	24

## 1.0 Overview of the Release

---

This document provides a brief overview of the changes introduced into the Intel® Ethernet Fabric Suite Software by this release. References to more detailed information are provided where necessary. The information contained in this document is intended as supplemental information only; it should be used in conjunction with the documentation provided for each component.

These Release Notes list the features supported in this software release, open issues, and issues that were resolved during release development.

### 1.1 Audience

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

### 1.2 Software License Agreement

This 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

Creating a technical support ticket for Intel® Ethernet Fabric Suite products is available 24 hours a day, 365 days a year. Contact Intel Customer Support or visit <https://www.intel.com/content/www/us/en/support.html> for additional detail.

### 1.4 Updated Features

The following table lists the updated features included in this release.

Feature	First Introduced
Improved PSM3 MR cache hit rate by allow non-exact hits.	12.1.0.0
Improved PSM3 flow control to avoid rare deadlocks under heavy load.	12.1.0.0
Updated to support Libfabric 2.x.	12.1.0.0
Added initial support for RC QP Failure Recovery.	12.1.0.0
Added NVIDIA DOCA OFED support.	12.1.0.0
Added PSM3 support for Intel® Ethernet E2100 IPU Adapters and 6th Generation Intel® Xeon® Scalable Processors.	12.0.0.0
Enhanced NCCL support.	12.0.0.0
Improved MOFED version detection.	11.7.0.2
<i>continued...</i>	

Feature	First Introduced
Rebuilt Intel® EFS Software for RHEL 8.10 and 9.4 with Intel® Direct GPU Access.	11.7.0.2
Rebuilt Intel® EFS Software for SLES 15.6 with Intel® Direct GPU Access.	11.7.0.1
Improved auto-tuning features for PSM3, including dynamic Credit Flows and detecting the presence of the rv kernel module.	11.7.0.0
Improved Container documentation in the Install Guide.	11.7.0.0
Improved PSM3 intra-node performance for large message sizes.	11.7.0.0
Added support and validation for running CUDA workloads on Intel® MPI Library.	11.7.0.0
Full support for new Intel® oneAPI DPC++/C++ Compiler (formerly only supported on the classic Intel® C++ Compiler).	11.6.0.0
Improved INSTALL process to detect and set GPU and MOFED environment variables.	11.6.0.0
Improved default tuning for Intel GPUs.	11.6.0.0
Technical Preview: Added the ability for PSM3 to consider both GPU and NIC locations within the server when selecting NICs for each process.	11.6.0.0
Rebuilt Intel® EFS Software for SLES 15.5 to add Intel® Direct GPU Access.	11.5.1.2
Improved Intel GPU small message latency and large message pipelining performance.	11.5.1.1
Improved Direct GPU access operations for Intel GPUs.	11.5.1.0
Improved FastFabric tool help text and man pages format and grammar.	11.5.1.0
Fully Supported: PSM3 runtime support for Intel GPUs via Intel® MPI Library (distributed with Intel® oneAPI), with ability to use Direct GPU access.	11.5.0.0
Improved statistics and diagnostic capabilities.	11.5.0.0
Technical Preview: Support for new Intel® oneAPI DPC++/C++ Compiler (formerly only supported on the classic Intel® C++ Compiler).	11.5.0.0
Technical Preview: User-mode memory region (MR) caching.	11.5.0.0
Technical Preview: PSM3 "Wait Mode" feature. Co-release with Intel® MPI Library.	11.5.0.0
Technical Preview: Support for Intel® Data Streaming Accelerator (DSA) memory transfers in 4th Generation Intel® Xeon® Scalable Processors. Shared DSA requires OS support, minimum of RHEL 8.7 or RHEL 9.1. Not available in GPU configurations.	11.5.0.0
Intel® Basic Software for InfiniBand package. Introducing a new software package targeting compute nodes in InfiniBand environments.	11.4.1.0
Improved support and tuning for AI Training fabrics.	11.4.0.0
Improvements for multi-plane deployment and management.	11.4.0.0
Multiple IP Address per individual NIC capability added. See <i>Intel® Ethernet Fabric Suite Host Software User Guide</i> for more info.	11.4.0.0
Technical Preview: Initial support for Ubuntu OS including PSM3 runtime for CPUs, packaging and installation via INSTALL script.	11.4.0.0
Technical Preview: PSM3 runtime support for Intel GPUs via Intel® MPI Library, with ability to use Direct GPU access (limited OS distributions).	11.4.0.0
Update to latest OSU Micro-Benchmarks and Intel® MPI Benchmarks supporting NVIDIA GPUs.	11.3.0.0
Improved <i>Intel® Ethernet Fabric Suite Host Software User Guide</i> documentation.	11.2.0.0
<b>continued...</b>	

Feature	First Introduced
Compatibility for PSM3 communication library with NCCL via AWS-OFI-NCCL plugin.	11.2.0.0
Support for TCP/Socket communication layer, runtime selectable.	11.2.0.0
Support for NVIDIA GPUs via CUDA, with the ability to use GPUDirect.	11.1.0.0
Deployment verification for Priority-based Flow Control.	11.1.0.0

See the [Intel® Ethernet Fabric Suite Documentation Library](#) for more information about these features.

## 1.5 Release Packages

There are three Intel® Ethernet Fabric Software packages:

- **Basic Software:** Software bundle for compute, service, and login nodes on Ethernet networks
- **Basic Software for InfiniBand:** Software bundle for compute, service, and login nodes on InfiniBand networks
- **Fabric Suite Software:** Software bundle for management nodes on Ethernet networks

The following table provides an overview of the included software components.

Included in Software Package	Basic Software	Basic Software for InfiniBand	Fabric Suite Software
Kernel rendezvous (RV) module for optimized scalable large message transfers for use with PSM3.	X	X	X
PSM3 OFI (libfabric) provider: libpsm3-fi.	X	X	X
Basic administration tools for use on all cluster nodes.	X	X	X
Pre-built Open MPI component. See <a href="#">Compilers and Libraries</a> for details.	X		X
FastFabric Toolset, which enables rapid, consistent installation and configuration of Intel® EFS host software and management software tools, as well as simplified installation, configuration, validation, and optimization of AI and HPC fabrics on Ethernet networks.			X

## 1.6 Intel Hardware

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

### NOTE

The Intel PSM3 implementation has a limit of 32 NICs per node.

**Table 1. Supported Hardware**

Hardware	Description
Intel® Xeon® Scalable Processors	Codename Sky Lake
2nd Generation Intel® Xeon® Scalable Processors	Codename Cascade Lake
3rd Generation Intel® Xeon® Scalable Processors	Codename Ice Lake
4th Generation Intel® Xeon® Scalable Processors	Codename Sapphire Rapids
5th Generation Intel® Xeon® Scalable Processors	Codename Emerald Rapids
6th Generation Intel® Xeon® Scalable Processors	Codename Granite Rapids
Intel® Ethernet E810 Network Adapters	Codename Columbiaville
Intel® Ethernet E2100 IPU Adapters	Codename Mount Evans
Intel® Data Center GPU Max Series	Codename Ponte Vecchio

## 1.7 Release Compatibility

It is often required that Intel® Ethernet Fabric Suite Software packages have consistent versions throughout a cluster. Occasionally, there exists limited backward compatibility.

For 12.1.0.0 release, the following is **required**:

- PSM3 requires all nodes in a job are running the same version.
- Consistent versions of RV and PSM3 are required.

See the *Intel® Ethernet Fabric Suite Host Software User Guide* for additional information.

Intel® Ethernet Fabric Suite Software packages have been validated against the most current software and firmware releases of Intel® Ethernet E810 Network Adapter at time of release. Those versions are listed below.

**Table 2. Supported Intel® Ethernet E810 Network Adapter Software**

Software Package	Version
Intel® Ethernet Adapter Complete Driver Pack	30.1
Intel® Ethernet Network Adapter Driver for E810 Devices under Linux	1.17.2
Linux RDMA Driver for the Intel® Ethernet Controllers E810 and X722	1.17.31

Compatibility with specific Intel® GPU Software versions is listed in [Intel GPU Support](#).

Compatibility with specific NVIDIA GPU Direct CUDA versions is listed in [CUDA Support](#).



## 1.8 Operating Systems

This release of the Intel® Ethernet Fabric Suite Software supports the operating systems listed in the following table.

**Table 3. Supported Operating Systems**

Operating System	Update/SP	Base Kernel Version
Red Hat Enterprise Linux (RHEL) 8.10 X86_64	Update 10	4.18.0-553.el8_10
Red Hat Enterprise Linux (RHEL) 9.5 X86_64	Update 5	5.14.0-503.11.1.el9_5
Red Hat Enterprise Linux (RHEL) 9.6 X86_64	Update 6	5.14.0-570.12.1.el9_6
Red Hat Enterprise Linux (RHEL) 10.0 X86_64		6.12.0-55.9.1.el10_0
Rocky Linux 8.10 (corresponds to RHEL 8.10)	8.10	4.18.0-553.5.1.el8_10
Rocky Linux 9.5 (corresponds to RHEL 9.5)	9.5	5.14.0-503.14.1.el9_5
Rocky Linux 9.6 (corresponds to RHEL 9.6)	9.6	5.14.0-570.17.1.el9_6
Rocky Linux 10.0 (corresponds to RHEL 10.0)	10.0	6.12.0-55.12.1.el10_0
SUSE Linux Enterprise Server (SLES) 15.6 X86_64	Service Pack 6	6.4.0-150600.16-default
SUSE Linux Enterprise Server (SLES) 15.7 X86_64	Service Pack 7	6.4.0-150700.51-default
Canonical Ubuntu Linux Server 22.04 LTS X86_64	22.04	5.15.0-25-generic

### NOTE

Ubuntu Linux support is a Technical Preview. Not all features of Intel® Ethernet Fabric Suite Software are available.

**Table 4. Supported OS by CPU**

CPU Type	RHEL				SLES	
	8.10	9.5	9.6	10.0	15.6	15.7
Intel® Xeon® Scalable Processors (codename Sky Lake)	•	•	•	•	•	•
2nd Generation Intel® Xeon® Scalable Processors (codename Cascade Lake)	•	•	•	•	•	•
3rd Generation Intel® Xeon® Scalable Processors (codename Ice Lake)	•	•	•	•	•	•
4th Generation Intel® Xeon® Scalable Processors (codename Sapphire Rapids)	•	•	•	•	•	•
5th Generation Intel® Xeon® Scalable Processors (codename Emerald Rapids)	•	•	•	•	•	•
6th Generation Intel® Xeon® Scalable Processors (codename Granite Rapids)	•	•	•	•	•	•

CPU Type	Ubuntu	Rocky			
	22.04	8.10	9.5	9.6	10.0
Intel® Xeon® Scalable Processors (codename Sky Lake)	•	•	•		
2nd Generation Intel® Xeon® Scalable Processors (codename Cascade Lake)	•	•	•		
3rd Generation Intel® Xeon® Scalable Processors (codename Ice Lake)	•	•	•		
4th Generation Intel® Xeon® Scalable Processors (codename Sapphire Rapids)	•	•	•		
5th Generation Intel® Xeon® Scalable Processors (codename Emerald Rapids)	•	•	•		
6th Generation Intel® Xeon® Scalable Processors (codename Granite Rapids)	•	•	•		

**NOTE**

Rocky Linux 9.6 and 10.0 were not released at time of 12.1.0.0 validation, but the equivalent RHEL 9.6 and 10.0 were fully validated.

## 1.9 Intel GPU Support

Intel® Direct GPU Access is supported by Intel® EFS via separate RPM modules provided.

Runtime, Intel® GPU driver, and distro support compatibility matrix are interleaved, as the Intel® GPU toolkit supports many forward and backward compatibilities.

The table below shows the Intel® GPU software versions that are compatible with this release of the Intel® Ethernet Fabric Suite Software.

**Table 5. Distro - Intel® GPU Software Support Matrix**

OS Distro (See Note)	Intel® GPU Driver	Level Zero Runtime	Intel® GPU Level Zero Interface library
RHEL 8.10	1.23.10.83.231129.91-132	1.14.0-803.123.el8_10	1.3.27642.67-803.134.el8_10
RHEL 9.5	1.23.10.83.231129.91-132	1.14.0-803.123.el9_5	1.3.27642.67-803.134.el9_5
RHEL 9.6	Intel® GPU not supported in current release.		
RHEL 10.0	Intel® GPU not supported in current release.		
SLES 15 SP6 (Note)	1.23.10.83.231129.91-124	1.14.0-i803.123	1.3.27642.67-803.126
SLES 15 SP7	Intel® GPU not supported in current release.		
Ubuntu 22.04	Distro support in Technical Preview -Intel® GPU not supported		

**NOTE**

For 12.1.0.0, only SLES 15 SP6 was validated on Intel® GPUs, other distros are supported, but not validated.

## 1.10 CUDA Support

CUDA GPUDirect is supported by Intel® EFS via separate RPM modules provided. CUDA Runtime, NVIDIA driver, and distro support compatibility matrix are interleaved, as the CUDA toolkit supports many forward and backward compatibilities. Refer to NVIDIA's [CUDA Compatibility](#), CUDA Toolkit, and Compatible Driver Versions table for more details.

The table below shows the CUDA software versions that are compatible with this release of the Intel® Ethernet Fabric Suite Software. Given the variety of CUDA compatible software and distros, some CUDA versions within the ranges specified below are supported but not validated.

NCCL is supported via OFI plugin (aws-ofi-nccl-version). This plugin enables libfabric as a network provider while running NVIDIA's NCCL based applications. The table below has been extended to also show compatibility among NCCL, distros, CUDA driver, and CUDA runtime versions.

**Table 6. Distro - CUDA Software Support Matrix**

OS Distro	NVIDIA GPU Driver	CUDA Runtime	NCCL / plugin
RHEL 8.10	555.42.06 - 575.51.03	CUDA 12.5.1 - 12.9	NCCL 2.27.3 aws-ofi-nccl-1.15.0
RHEL 9.5	570.86.10 - 575.51.03 (See Note)	CUDA 12.8.0 - 12.9	NCCL 2.27.3 aws-ofi-nccl-1.15.0
RHEL 9.6	CUDA not supported in current release.		
RHEL 10.0	CUDA not supported in current release.		
SLES 15 SP6	535.86.10 - 575.51.03 (See Note)	CUDA 12.6.1 - 12.9	NCCL 2.27.3 aws-ofi-nccl-1.15.0
SLES 15 SP7	CUDA not supported in current release.		
Ubuntu 22.04	Distro support in Technical Preview - CUDA not supported		

- Red Hat Enterprise Linux 9.0 and SLES 15 SP4 (and newer) are required to use NVIDIA's open driver which is only supported on Ampere GPUs (A100) and later.
- With the latest CUDA releases, NVIDIA's open driver is now the default driver to use when installing.

## 1.11 NVIDIA DOCA OFED Support

NVIDIA DOCA OFED is supported by Intel® EFS via rebuilding of the RPM modules provided. Mellanox OFED has been replaced with NVIDIA DOCA OFED. The DOCA OFED package supports various compatibilities. Refer to the DOCA OFED <https://docs.nvidia.com/doca/sdk/index.html> page for more details.

The table below shows the DOCA and MOFED software versions that are validated with this release of the Intel® Ethernet Fabric Suite Software. Testing was done on InfiniBand NICs.

**Table 7. Distro - DOCA Software Support Matrix**

OS Distro	DOCA OFED	MOFED
RHEL 8.10	2.9.2-0.1.5	24.10-2.1.8.1
RHEL 9.5	2.9.2-0.1.5	24.10-2.1.8.1
RHEL 9.6	DOCA OFED not supported in current release.	
RHEL 10.0	DOCA OFED not supported in current release.	
SLES 15 SP6	2.9.2-0.1.5	24.10-2.1.8.1
SLES 15 SP7	DOCA OFED not supported in current release.	
Ubuntu 22.04	Distro support in Technical Preview - MOFED not supported	

#### NOTE

DOCA OFED refers to the top-level `doca-ofed` meta package and MOFED refers to the underlying kernel modules used during install of Intel® EFS.

#### NOTE

Currently Intel® Ethernet Fabric Suite FastFabric does not support InfiniBand, so Intel® recommends installing the Intel® Basic Software for InfiniBand (`Intel-Basic-IB.<distro>-x86_64.<version>.tgz`) package, when using InfiniBand.

## 1.12 Kernel Modules for Rendezvous Driver

This release of the Intel® Ethernet Fabric Suite Software contains the rendezvous driver kernel modules listed in the following table.

#### NOTE

The version number is shown in bold.

**Table 8. Kernel Modules for Rendezvous Driver**

OS	RPM Version
RHEL 8.10	kmod-iefs-kernel-updates-4.18.0_553.el8_10.x86_64- <b>6364</b> .x86_64.rpm
RHEL 9.5	kmod-iefs-kernel-updates-5.14.0_503.11.1.el9_5.x86_64- <b>6364</b> .x86_64.rpm
RHEL 9.6	kmod-iefs-kernel-updates-5.14.0_570.12.1.el9_6- <b>6364</b> .x86_64.rpm
RHEL 10.0	kmod-iefs-kernel-updates-6.12.0_55.9.1.el10_0.x86_64- <b>6364</b> .x86_64.rpm
SLES 15.6	iefs-kernel-updates-kmp-default-6.4.0_150600.16_default_k6.4.0_150600.21- <b>6364</b> .x86_64.rpm
SLES 15.7	iefs-kernel-updates-kmp-default-6.4.0_150700.51_default_k6.4.0_150700.51- <b>6364</b> .x86_64.rpm
Ubuntu 22.04	kmod-iefs-kernel-updates_5.15.0.25- <b>6364</b> _amd64.deb

**Table 9. oneAPI Intel GPU Kernel Modules for Rendezvous Driver**

OS	oneAPI Intel GPU RPM Version
RHEL 8.10	kmod-iefs-kernel-updates-4.18.0_553.el8_10.x86_64- <b>6364oneapize</b> .x86_64.rpm
RHEL 9.5	kmod-iefs-kernel-updates-5.14.0_503.11.1.el9_5.x86_64- <b>6364oneapize</b> .x86_64.rpm
RHEL 9.6	Intel® GPU not supported in current release.
RHEL 10.0	Intel® GPU not supported in current release.
SLES 15.6	iefs-kernel-updates-kmp-default-6.4.0_150600.16_default_k6.4.0_150600.21- <b>6364oneapize</b> .x86_64.rpm
SLES 15.7	Intel® GPU not supported in current release.
Ubuntu 22.04	Ubuntu support is a Technical Preview - oneAPI not supported

**Table 10. CUDA Kernel Modules for Rendezvous Driver**

OS	CUDA RPM Version
RHEL 8.10	kmod-iefs-kernel-updates-4.18.0_553.el8_10.x86_64- <b>6364cuda</b> .x86_64.rpm
RHEL 9.5	kmod-iefs-kernel-updates-5.14.0_503.11.1.el9_5.x86_64- <b>6364cuda</b> .x86_64.rpm
RHEL 9.6	CUDA not supported in current release.
RHEL 10.0	CUDA not supported in current release.
SLES 15.6	iefs-kernel-updates-kmp-default-6.4.0_150600.16_default_k6.4.0_150600.21- <b>6364cuda</b> .x86_64.rpm
SLES 15.7	CUDA not supported in current release.
Ubuntu 22.04	Ubuntu support is a Technical Preview - CUDA not supported

**NOTE**

Ubuntu Linux support is a Technical Preview. Not all features of Intel® Ethernet Fabric Suite Software are available.

## 1.13 Network File Systems

Network shared file systems are common in environments where Intel® Ethernet Fabric Suite Software is deployed. This section reports the network file systems that have been tested with Intel® Ethernet E810 Network Adapters as shown in [Release Compatibility](#).

The following parallel file systems have been validated with RDMA:

- IBM Spectrum Scale<sup>1</sup>:
  - Version 5.2.2.1 on supported distros
- BeeGFS
  - Version 8.0 on RHEL all supported distros, not validated on SLES

Additionally, in-distro NFS capability has also been validated for the supported distros.

<sup>1</sup> Formerly known as General Parallel File System (GPFS)

## 1.14 Compilers and Libraries

Intel® Ethernet Fabric Suite Software package provides both source and binary files. The binary files are built using either the GCC (OS distribution) or Intel compilers. The following tables capture the specific versions of the compilers that were used in this release and the binary components that they correspond.

**Table 11. Compiler Versions and Distributions**

Compiler	Distribution	Compiler Version
(Intel) icx	Intel® oneAPI	Intel® oneAPI DPC++/C++ Compiler 2025.1.1 (2025.1.1.20250418)
(GNU) gcc	RHEL 8.10	gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-21)
(GNU) gcc	RHEL 9.5	gcc (GCC) 11.5.0 20240719 (Red Hat 11.5.0-2)
(GNU) gcc	RHEL 9.6	gcc (GCC) 11.5.0 20240719 (Red Hat 11.5.0-5)
(GNU) gcc	RHEL 10.0	gcc (GCC) 14.2.1 20250110 (Red Hat 14.2.1-7)
(GNU) gcc	SLES 15 SP6	gcc (SUSE Linux) 7.5.0
(GNU) gcc	SLES 15 SP7	gcc (SUSE Linux) 7.5.0
(GNU) gcc	Ubuntu 22.04	gcc (Ubuntu 11.3.0-1ubuntu1~22.04) 11.3.0

**Table 12. Intel® EFS Components and Compiler**

Component	Compiled With
PSM3	ICX
RV Module (Rendezvous Module)	GCC
Fast Fabric CLI tools	GCC
Open MPI 4.1.8	GCC (See Note)

### NOTE

FastFabric provides a simple tool to rebuild this MPI library using either the GCC compiler or the Intel® oneAPI DPC++/C++ Compiler (not included in the software package). Refer to *Intel® Ethernet Fabric Suite FastFabric User Guide*, Rebuilding MPI Library and Tools for additional information to properly setup the environment for compiling MPI libraries and applications.

### 1.14.1 Supported MPI Libraries

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

**Table 13. Supported MPI Libraries**

MPI Implementation	Included in Basic Package?	Runs Over
Open MPI 4.1.8	Yes (See note)	PSM3 (via libfabric)
Intel® MPI Library 2021, Update 16	No	PSM3 (via libfabric)

Intel® MPI Library 2021, Update 16, is recommended to be used with Intel® EFS, and can be obtained via the following URL (with proper credentials): [oneAPI HPC Toolkit](#)

---

**NOTE**

Open MPI libraries are not included for Ubuntu Linux as support is in Technical Preview.

---



---

**NOTE**

This release of Intel® EFS has been validated against the in-distro versions of libfabric and the version of libfabric included with Intel® MPI Library 2021, Update 16. Other versions of libfabric (1.7.2 and later) are supported.

---

## 1.15 Document Versions

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

**Table 14. Supported Document Versions**

Title	Doc. Number	Revision
<i>Intel® Ethernet Fabric Suite Software Installation Guide</i>	632484	1.10
<i>Intel® Ethernet Fabric Suite Host Software User Guide</i>	632489	1.11
<i>Intel® Ethernet Fabric Suite FastFabric User Guide</i>	632487	1.8
<i>Intel® Ethernet Fabric Performance Tuning Guide</i>	632488	1.10
<i>Intel® Ethernet Fabric Suite Software Release Notes</i>	N78966	1.0

See [Intel® Ethernet Fabric Suite Documentation Library](#) for document descriptions.

## 1.16 Installation Requirements

This section provides installation requirements for this release.

### 1.16.1 Best Practices

- Intel recommends that users update to the latest versions of Intel® Ethernet Fabric Suite software to obtain the most recent functional and security updates.
- To improve security, the administrator should log out users and disable multi-user logins prior to performing provisioning and similar tasks.

### 1.16.2 Software and Firmware Requirements

Operating systems supported by this release are listed in [Table 3](#). For the required packages, refer to [OS Installation Prerequisites](#).

### 1.16.3 OS Installation Prerequisites

Ensure that the following requirements are met prior to installing the software.

- In addition to normal OS installation options, OS packages must be installed before you can install the Intel® Ethernet Fabric Suite software.

- Refer to the applicable section below to verify that all required packages 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 OS Distribution Release Notes for your specific release and installation type.

### Red Hat Enterprise Linux OS RPMs

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

#### NOTE

Some RHEL rpms may require packages included in the Server-Optional Install disks or additional online repositories such as the CRB repository.

#### NOTES

- RHEL 8.4+ adds `kernel-abi-stablelists`.
- RHEL 10.0+ removes `atlas` and replaces `zlib` with `zlib-ng-compat`.

**Table 15. RHEL 8.x, 9.x, 10.x Distribution RPMs**

Ethernet Centric	System Centric		Other/Generally Installed	Build Requirements
libibverbs	atlas (NOTE)	ncurses-compat-libs	bash	autoconf
librdmacm	bc	ncurses-libs	kernel	automake
rdma-core	coreutils	numactl-libs	kernel-modules-extra	bison
rdma-core-devel	createrepo	openblas-devel	kernel-tools	elfutils-libelf-devel
libfabric (1.7.2 or higher)	expat	openssl	kmod	expat-devel
	expect	openssl-devel	libgcc	flex
	gcc-gfortran	openssl-libs	mpi-selector	gcc-c++-x86_64
	glibc	pciutils	perl	kernel-abi-stablelists (NOTE)
	kernel-devel	redhat-rpm-config	perl-Getopt-Long	kernel-abi-whitelists
	libatomic	rpm-build	perl-Socket	kernel-rpm-macros
	libgfortran	sysfsutils	pkgconf	libfabric-devel (1.7.2 or higher)
	libgomp	tcl	python2	libnl3-devel
	libquadmath	tcsh	systemd	libpfm
	libstdc++	zlib (NOTE)		libtool
	libstdc++-devel	zlib-ng-compat (NOTE)		libuuid-devel
		hwloc-libs		ncurses-devel
				numactl-devel
continued...				



Ethernet Centric	System Centric		Other/Generally Installed	Build Requirements
				openssl-devel (1.1.1 or higher)
				perl-generators
				tcl-devel

### SUSE Linux Enterprise Server 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). Install the SLES SDK iso which contains additional, required development packages.

**Table 16. SLES15.x Distribution RPMs**

Ethernet Centric	System Centric		Other/Generally Installed	Build Requirements
libfabric (1.7.2 or higher)	bc	libopenssl1_1	bash	bison
libibverbs1	coreutils	libopenssl-devel	cpupower	flex
libibverbs-utils	createrepo_c	libosmcomp3	glibc	kernel-devel
librdmacm1	expect	libquadmath0	grep	libexpat-devel
librdmacm-utils	gcc-fortran	libudev-devel (SLES 15.3)	irqbalance	libnuma-devel
net-snmp	kernel-devel	libz1	kmod	libopenssl-devel (1.0.1 or higher)
rdma-core	kernel-syms	openssl	libedit0	libuuid-devel
rdma-core-devel	libatomic1	openblas-devel (SLES 15.3)	libgcc_s1	ncurses-devel
	libexpat1	openblas-common-devel (SLES 15.4)	libstdc++6	tcl-devel
	libgfortran4	rpm-build	perl	valgrind-devel
	libgomp1	systemd-devel (SLES 15.4)	perl-base	libfabric-devel (1.7.2 or higher)
	libncurses6	tcl	pkg-config	
	libnuma1	tcsh	python3-base	
		libhwloc15	systemd	
			udev	
			mpi-selector	

## Ubuntu Linux Server OS DEBs

The tables below list the DEBs for each supported Ubuntu release.

### NOTE

Some Ubuntu package are available in additional repositories. Enable Universe and Multiverse repositories that contains additional, required development packages.

**Table 17. Ubuntu 22.04 Distribution DEBs**

Ethernet Centric	System Centric		Other/Generally Installed	Build Requirements
expect	bc	systemd	bash	autoconf
gfortran	coreutils	tcl	debhelper	automake
libatomic1	debconf	tcsh	flex	bison
libfabric1	expat	uuid-dev	gcc	libelf-dev
libgfortran5	libatlas3-base	zlib1g	libfabric-dev	libexpat1-dev
libgomp1	libgcc-s1	libhwloc15	libnl-3-dev	
libibverbs1	libgetopt-long-descriptive-perl		libpfm4	
libibverbs-dev	libncurses6		linux-image-generic	
libncurses-dev	libnuma1		linux-modules-extra-<kver>-generic	
libnuma1	libopenblas-dev		linux-tools-common	
libquadmath0	libsocket-perl		kmod	
librdmacm-dev	libssl-dev			
librdmacm1	libtool			
libssl-dev	openssl			
libstdc++-11-dev	pciutils			
libstdc++6	perl			
linux-headers	pkgconf			
rdma-core	python3			
tcl-dev	sysfsutils			

### Example

The example below shows how to install a group of OS RPMs on a RHEL OS server. Not all required OS RPMs are included in this example and some OS RPMs in this example might not be needed.

```
# dnf install qperf perftest rdma expat libstdc++-devel gcc-gfortran atlas tcl
expect tcsh sysfsutils pciutils bc libibverbs-devel librdmacm-devel ibacm-devel
openssl-devel libuuid-devel expat-devel valgrind-devel libgnum libibverbs*
ncurses-devel hwloc
```

Intel recommends that you build your own list of OS packages for installation.

#### 1.16.4 Installation Instructions

There are three Intel® Ethernet Fabric Software packages:

- `IntelEth-FS.<distro>-x86_64.<version>.tgz` for the management node on Ethernet Networks.
- `IntelEth-Basic.<distro>-x86_64.<version>.tgz` for compute, service, and login nodes on Ethernet networks
- `Intel-Basic-IB.<distro>-x86_64.<version>.tgz` for compute, service, and login nodes on InfiniBand networks

The following methods are available to install the software:

- INSTALL TUI
- INSTALL CLI
- INSTALL createrepo feature
- FastFabric multi-server installation

Refer to the *Intel® Ethernet Fabric Suite Software Installation Guide* for related software requirements and complete installation procedures.

## 2.0 Issues

This section lists the resolved and open issues in the Intel® Ethernet Fabric Suite 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 18. Issues Resolved in this Release**

ID	Description	Resolved in Release
HED-828	RDMA renaming techniques do not reliably work, causing a mismatch between the net device name and the RDMA device name on boot. Setting proper udev rules will now solve this issue:  <pre>ACTION=="add", KERNELS=="0000:31:00.0", \ SUBSYSTEM=="net", NAME=="cvl0"</pre> <pre>ACTION=="add", KERNELS=="0000:31:00.0", \ SUBSYSTEM=="infiniband", PROGRAM=="rdma_rename %k NAME_FIXED irdma-cvl0"</pre>	12.1.0.0 (Not reproducible with v31.0 Intel® Ethernet Adapter Complete Driver Pack)
HED-5078	When running a Multi-Threaded job with Intel® MPI Library 2021.13, workloads can hang.  <pre>mpirun -genv OMP_NUM_THREADS 16 \   -ppn 16 -np 32 \   IMB-MT -thread_level multiple</pre>	12.1.0.0 (Resolved with 2021.16 Intel® MPI)
HED-5332 HED-5333	Occasional hangs can be observed when running IMB-RMA, which focuses on one-sided operations, with 32 or more ranks across several nodes.	12.1.0.0
HED-5361	When using PSM3_RDMA=0 and PSM3_GPUDIRECT=0 it is possible to see a hang when running on Intel GPUs.	12.1.0.0
HED-5369	Occasionally, when running PSM3 with CUDA Managed Memory (e.g. osu_multi_lat MD MH), PSM3_GPRUDIRECT=1 may not work and will abort:  <pre>reg_mr failed: 0x14e056002000:0x14e056003fff (len 0x2000) access 0x80000003: Bad address</pre>	12.1.0.0
HED-5316	When running PSM3 with IPv6 addressing, i.e. PSM3_ADDR_FMT=6, on Intel E810 cards, very poor MPI performance may happen if these cards are under network pressure, such as incast traffic with large number of QPs.	12.1.0.0 Refer to <i>Intel® Ethernet Fabric Performance Tuning Guide</i>

## 2.1.2 Issues Resolved in Prior Releases

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

**Table 19. Issues Resolved in Prior Releases**

ID	Description	Resolved in Release
HED-5095	When running large MPI jobs with many ranks per node (e.g. 2048 ranks with 128 ranks per node), occasionally the following may occur and the job will be aborted:  reg_mr failed: 0x2f000965ca00:0x2f000966c9ff (len 0x10000) access 0x80000003: Bad address	12.0.0.0
HED-1714	After upgrading the kernel with Dynamic Kernel Module Support (DKMS), the irdma module is removed. The irdma module now supports weak-updates which should cover most cases.	12.0.0.0
HED-4895	The usage text for the INSTALL program does not match the text in the Installation Guide that describes the command line options to INSTALL, in particular, the discussion of the "-G" option and auto-detection of GPU drivers and software.	11.7.0.0
HED-4873	When running CPU AI workloads and using PSM3's intra-node (shm) mechanisms, with multiple threads (workers or openMP), the default PSM3 shm modes' use of Linux Cross-Memory-Attach (CMA) may exhibit significant AllReduce latency jitter.	11.7.0.0
HED-4195	When compiling a C program or library against CUDA 12.0 and 12.1 it may fail to compile if compiled with -Werror -Wundef.	11.7.0.0
HED-4509	When running Intel MPI with Intel GPU, I_MPI_OFFLOAD_RDMA=1, I_MPI_FABRICS=ofi and the PSM3 provider, the pingping test can hang.	11.6.0.0
HED-4486	On Intel GPU, when running PSM3 with Intel MPI I_MPI_OFFLOAD_RDMA=1, PSM3_RDMA=1 and PSM3_GPUDIRECT=0 PSM3 can repeatedly report the following error:  reg_mr failed: ... access 0x90000003: Invalid argument"  Similar errors may occur with other middleware.	11.6.0.0
HED-4392	For applications using Intel GPUs, when using MPI_AlltoAll with large messages (such as with GROMACS and 512K-2M messages), low performance can be observed; also when using large messages, significant latency jitter may be observed.	11.6.0.0
HED-1155 HED-2980	Prior to version 1.12.1, libfabric may layer the ofi_rxd provider above PSM3 and result in lower than expected OpenMPI performance when running with PSM3 (psm3;ofi_rxd).	11.6.0.0 Fixed in Libfabric included in supported OSes
HED-977	After upgrading the rdma-core version as part of Intel® Ethernet E810 Controller software installation, attempts to rebuild PSM3 may fail due to rdma header file conflicts.	11.6.0.0 Fixed in Supported OSes
HED-3875	Intel® oneAPI Collective Communications Library (oneCCL) examples may fail to build.	11.5.0.0 (Resolved with Intel® oneCCL 2021.9)
HED-2742	After building and installing IRDMA driver, then attempting to install Intel® EFS package, the system may experience a kernel panic if not rebooted immediately after installing irdma.	11.4.1.0 (Resolved with irdma 1.8)
HED-3611	Intel® EFS ethcreatrepo does not generate proper repositories when executed on non-GPU environments.	11.4.1.0
HED-3639	On Ubuntu 22.04, INSTALL reports an error indicating dracut is not found.	11.4.1.0
continued...		

ID	Description	Resolved in Release
HED-3468	On RHEL 9.0, Intel® EFS Fast Fabric commands that rely on the common <code>scp</code> tool may encounter errors in the case where source and destination are same filename and same server. See <a href="https://bugzilla.redhat.com/show_bug.cgi?id=2108409">https://bugzilla.redhat.com/show_bug.cgi?id=2108409</a>	11.4.1.0 (Resolved in openssh 8.7p1-10)
HED-87 Whamcloud LU-13976	Lustre is not supported on Intel® Ethernet E810 Network Adapters.	Resolved with Lustre LTS version 2.12.8
HED-2742	After building and installing IRDMA driver, then attempting to install Intel® EFS package, the system may experience a kernel panic if not rebooted immediately after installing IRDMA.	11.4.0.0 (Resolved with v27.7 Intel® Ethernet Adapter Complete Driver Pack)
HED-2958	On RHEL 8.5, when running <code>hostverify</code> , <code>intel_pstate</code> will show as <code>intel_cpufreq</code> .	11.4.0.0
HED-2966	Intel® MPI v2021.6.0 will crash when running with <code>I_MPI_OFI_PROVIDER=psm3</code> .	11.4.0.0 (Resolved with Intel® MPI v2021.7)
HED-1018	If the network interface experiences a link down event (for example, a disconnected cable or <code>ip link set &lt;interface&gt; down</code> ), the associated RDMA device is removed and no longer available to RDMA applications. As a result, running RDMA jobs may fail. When the network interface link is restored, the RDMA device is automatically re-added. At this time, RDMA jobs may be restarted.	11.2.0.0 (Resolved with v26.8 Intel® Ethernet Adapter Complete Driver Pack)
HED-1286	In upstream and other distributions of OpenMPI, <code>MPI_Ssend</code> and <code>MPI_Issend</code> do not work properly. As a result, additional 0-length messages may be presented to the MPI application.	11.2.0.0 (Resolved in OpenMPI 4.1.2)
HED-1993	A bug was introduced in the OSU Micro-Benchmarks in version 5.7.1 for benchmarks with CUDA enabled. When the managed memory ("MD" or "MH") mode is mixed with non-managed memory ("H" or "D") modes, the error <code>[./../util/osu_util_mpi.c:1691] CUDA call 'cudaMemPrefetchAsync(buf, length, devid, um_stream)' failed with 1: invalid argument (or similar)</code> may occur and cause an application segfault. The bug has been reported upstream to OSU maintainers.	11.2.0.0 (Resolved in OSU 5.8)
HED-1144	The <code>mpicheck</code> tool, included in <code>mpi_apps</code> , uses <code>/dev/random</code> to seed the random number generator. <code>/dev/random</code> blocks when it runs out of entropy, so reading from <code>/dev/random</code> can halt process execution.	11.2.0.0
HED-1441	Error encountered when attempting to install <code>openmpi rpm</code> built from <code>do_build</code> script in <code>/usr/src/eth/MPI</code> . Reports conflicts with <code>mpitests_openmpi_gcc_hfi-4.2-932.x86_64</code> .	11.2.0.0
HED-2002	Early user termination (Ctrl-C) of certain HPC applications using PSM3 with CUDA and GPUDirect cause nodes to crash.	11.2.0.0 (Resolved in CUDA 11.4.1)
HED-1265 (previously HED-1256)	The <code>osu_latency_mp</code> test (fork and do comms only from the parent) hangs even if <code>IBV_FORK_SAFE=1</code> .	11.1.0.0 (Resolved in irdma v1.4.47)
HED-1410	Attempting to build <code>libpsm3-fi</code> without installing <code>iefs-kernel-updates-devel</code> fails with the message: <code>configure: error: RV Driver Headers requested but &lt;rv/rv_user_ioctl.h&gt; not found</code>	11.1.0.0
HED-1440	If executing <code>do_build</code> from <code>/usr/src/eth/MPI</code> without having installed <code>libfabric-devel rpm</code> , the build will attempt to build PSM2 API and is likely to fail or use an unintended library, yielding an <code>openmpi</code> binary rpm that will not run on Ethernet.	11.1.0.0
HED-1446	Help and Usage output for <code>ethcaptureall</code> still mentions "chassis" and "switch" in help text.	11.1.0.0
<b>continued...</b>		

ID	Description	Resolved in Release
HED-1456 (previously HED-1465)	Occasionally, when running PSM3 using Open MPI in multi-rail configurations (PSM3_MULTIRAIL=1 or 2 and multiple NICs per process, or PSM3_QP_PER_NIC>1), jobs may hang in MPI_Finalize. This behavior has not been observed when running Intel® MPI Library.	11.1.0.0
HED-1481	When following the installation guide and using the ethfastfabric tool in the order shown in the menus, the "Configure SNMP" step may fail due to a lack of allhosts file.	11.1.0.0
HED-1482	When running ethfabricinfo directly or via ethfastfabric the "Summary of Fabric Components" fails, ethfabricinfo reports it cannot find the configuration file /etc/eth-tools/mgt_config.conf.	11.1.0.0
HED-1484	On some systems, when running Perform Single Host Verification via ethfastfabric, the hostverify tool reports a timeout.	11.1.0.0
HED-1485	Various fastfabric operations may fail. This is most obvious is when running Perform Single Host Verification via ethfastfabric. The hostverify tool reports that snmpget is not found.	11.1.0.0

## 2.2 Open Issues

The following table lists the open issues for this release.

**Table 20. Open Issues**

ID	Description	Workaround
HED-898	Occasional hangs are observed during Intel® MPI Benchmarks, IMB-EXT Accumulate benchmark.	Intel® EFS is not validated for one-sided MPI operations, as such scenarios are not anticipated to be widely deployed.
HED-3645	On Ubuntu 22.04, reinstalling Intel® EFS will overwrite configuration files regardless of options passed to <code>INSTALL (-O/-N)</code> .	Intel recommends that you manually back up configuration files before reinstalling software.
HED-3889	<p>INSTALL of IEFS may fail if installing IEFS on a SLES 15.x machine with MOFED installed.</p> <p>Lacking OS Prereqs for PSM3 --&gt; psm3 requires libibverbs1</p> <p>This is caused by the MOFED not using the same package naming style as SLES</p>	Once all other SLES 15.x prerequisites listed in <a href="#">Table 16</a> are met, rerun INSTALL with <code>--without-depcheck</code> to skip the dependency checking.
HED-4969	<p>When running large MPI jobs with many ranks per node (e.g. 2048 ranks with 128 ranks per node), occasionally the NMI watchdog will kick in during memory deregistration and depending on system settings, a kernel panic can occur.</p> <p>NMI watchdog: Watchdog detected hard LOCKUP on cpu 146</p>	<p>Several possible workarounds can be used:</p> <ul style="list-style-type: none"> <li>Disable iommu by removing <code>intel_iommu=on</code> from the kernel command line.</li> <li>Either increase the NMI watchdog timeout or disable it. However, this may cause unrelated issues to not be detected in a timely manner.</li> </ul>
HED-5035	When using <code>PSM3_RDMA=2</code> or <code>3</code> for jobs with a high process per node count and/or a high node count, the number of RC QPs required can be much larger and may exceed NIC HW or software limits. In addition, <code>PSM3_RDMA=3</code> may also require more memory.	On Intel E810 NICs, the number of QPs available may be increased by setting the <code>irdma</code> driver's module parameters to include <code>limits_sel=6</code> and rebooting. Setting <code>FI_PSM3_LAZY_CONN=1</code> , may also reduce the number of QPs and memory needed for some applications.
HED-5490	When using <code>PSM3_CUDA=1</code> and <code>PSM3_GPUDIRECT=1</code> for jobs on MLX NICs, the resulting performance for medium sized messages may suffer.	Setting <code>PSM3_GPUDIRECT_SDMA=0</code> , will stop PSM3 from attempting to register an MR the NIC cannot support.

## 2.2.1 Third-Party Open Issues

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

**Table 21. Third-Party Open Issues**

ID	Description	Workaround
HED-727	The <code>osu_mbw_mr</code> benchmark as provided by OSU OMB v3.9 and newer (including 5.9 in this release) may return bandwidth greater than 100 Gbps (line rate). This is due to the average time being used instead of the maximum time between communication pairs. See <a href="http://mvapich.cse.ohio-state.edu/static/media/mvapich/CHANGES-OMB.txt">http://mvapich.cse.ohio-state.edu/static/media/mvapich/CHANGES-OMB.txt</a> for the change log in version 3.9.	Use <code>mpi_multibw</code> included in <code>mpi_apps</code> . (This is a modified version of the OSU v3 benchmark.)
HED-1170	If incorrect NIC filtering is applied via the environment variables: <code>PSM3_HAL</code> , <code>PSM3_NIC</code> , <code>PSM3_ADDR_FMT</code> , <code>PSM3_SUBNETS</code> , <code>PSM3_NIC_SPEED</code> , Open MPI may silently use the <code>tcp btl</code> instead of <code>ofi mtl</code> .	Specify the Open MPI parameter <code>-mca btl ^tcp</code> which should cause a failure if an incorrect filter is specified.
HED-1962	As of Open MPI 4.1.0, OFI has been added to the BTL layer. By default, this will result in an additional OFI context being opened/allocated per rank on every node. In high rank count jobs, this can result in excessive memory use and potentially other issues. Intel recommends that you disable OFI BTL if not required.	Add <code>--mca btl ^ofi</code> to <code>mpirun</code> command args to disable OFI BTL.
HED-2057	When using NVIDIA Quadro GPUs and multiple processes per node, PCI BAR space exhaustion may result in incorrect mappings of GPUDirect pinned memory and resulting application data corruption. The issue is believed to be in the NVIDIA GPU drivers and has not been observed for Tesla GPUs.	<p>PSM3 11.1 includes automated detection of this issue when the IO triggering the issue proceeds. Upon detection, PSM3 causes immediate job exit with this message:</p> <pre>Incorrect Physical Address (0xa0000000) returned by nVidia driver. PSM3 exiting to avoid data corruption. Job may be rerun with PSM3_GPUDIRECT=0 to avoid this issue.</pre> <p>If this occurs, data produced immediately prior to the error may also be suspect and should be discarded. The job may be re-run successfully without <code>PSM3_GPUDIRECT</code> or with one process per node.</p>
HED-2191 HED-5261	With clusters using RDMA via RoCE running Spectrum Scale (GPFS) file system in an environment with <code>verbsRdmaCm</code> enabled, individual nodes will fail to establish a connection to other nodes when the IP address is configured to include a non-link local IPv6 address.	<ul style="list-style-type: none"> <li>IPv6 must be enabled to use RoCE, if interfaces are selected by using the port name.</li> <li>IPv6 local link address needs to be conformed with <code>IPV6_ADDR_GEN_MODE=eui64</code></li> </ul>
HED-2491	NVIDIA's Pre-Built Driver RPMs do not accurately list the contents they "provide". This results in a missing dependency when attempting to install Intel® EFS kernel module.	<p>Install Intel® EFS kernel module rpm without dependencies.</p> <pre>rpm -U --force --nodeps ./iefs-kernel-updates-kmp- default-5.3.18_57_default_k5.3.18_57-6252cuda.x86_ 64.rpm</pre>
HED-3571	An unhandled <code>cuda error</code> may occur when running Horovod (0.25.0) with the CUDA 11.7 release resulting in stopped jobs. More details are available here: <a href="https://github.com/horovod/horovod/issues/3625">https://github.com/horovod/horovod/issues/3625</a> .	No workaround at this time.
continued...		



ID	Description	Workaround
HED-3601	Restarting <code>lldpad</code> causes segfault.	Restarting <code>lldpad</code> works without error. The segfault can be ignored.
HED-4011	In some environments, the PSM InfiniPath library, which is no longer supported, exports internal symbols that can conflict with running/compiling the PSM3 OFI Provider.	Ensure that PSM1/Intel® True Scale provider library is removed before installing IEFS / PSM3
HED-4469	Starting with SLES 15.5, Nvidia's Open-source G06 driver is added to SLES's in-distro repo. However, this driver does not include the necessary files to rebuild Intel® EFS RV Kernel module and is not claimed to be supported by Nvidia CUDA Release Notes.	Uninstall SLES's Nvidia driver and install Nvidia's Open Driver: <pre>zypper remove nvidia-open-driver-G06-signed-kmp-default zypper install nvidia-open-gfxG05-kmp-default</pre>
HED-5086	When attempting to run GPU Direct test on an Intel GPU a kernel panic in i915 driver can sometimes be observed.	Remove the following from the grub cmdline: <pre>initcall_blacklist=sync_debugfs_init,dma_buf_init</pre>

## 3.0 Related Information

### 3.1 Intel® Ethernet Fabric Suite Documentation Library

Intel® Ethernet Fabric Suite publications are available at the following URL:

<https://www.intel.com/content/www/us/en/support/articles/000088090/ethernet-products/intel-ethernet-software.html>

Use the tasks listed in this table to find the corresponding Intel® Ethernet Fabric Suite document.

Task	Document Title	Description
Installing host software Installing NIC firmware	<i>Intel® Ethernet Fabric Suite 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 fabric using FastFabric	<i>Intel® Ethernet 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 Text-based User Interface (TUI) menus and command line interface (CLI) commands.
Running MPI applications on Intel® EFS Running middleware that uses Intel® EFS	<i>Intel® Ethernet Fabric Suite Host Software User Guide</i>	Describes how to set up and administer the Network Interface Card (NIC) after the software has been installed and provides a reference for users working with Intel PSM3. Performance Scaled Messaging 3 (PSM3) is an Open Fabrics Interface (OFI, also called libfabric) provider which implements an optimized user-level communications protocol. The audience for this document includes cluster administrators and those running or implementing Message-Passing Interface (MPI) programs.
Optimizing system performance	<i>Intel® Ethernet Fabric Performance Tuning Guide</i>	Describes BIOS settings and parameters that have been shown to ensure best performance, or make performance more consistent, on Intel® Ethernet Fabric Suite Software. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.
Learning about new release features, open issues, and resolved issues for a particular release	<i>Intel® Ethernet Fabric Suite Software Release Notes</i>	

### 3.1.1 How to Search the Intel® Ethernet Fabric Suite 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 Intel® Ethernet Fabric Suite PDFs into a single folder.
2. Open your PDF 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.