Intel® Ethernet Fabric Suite Fabric Software

Installation Guide

Rev. 1.10

June 2025

Doc. No.: 632484, Rev.: 1.10

Revision History

Date	Revision	Description
June 2025	1.10	Product 12.1.0.0 release - Changes to this document include: • Updated Download and Install DOCA OFED Software (Optional) on page 31 to include new name for NVIDIA's OFED software stack.
February 2025	1.9	Product 12.0.0.0 release - Changes to this document include: • Updated Download and Install Intel® Ethernet Drivers (Optional) to include Intel® Ethernet E2100 IPU Adapters. • Updated Download and Install NVIDIA CUDA Software (Optional) with changes to NVIDIA's Open-Source GPL driver.
July 2024	1.8	Product 11.7.0.0 release - Changes to this document include: • Added section for Intel® EFS Containers Installation
March 2024	1.7	Product 11.6.0.0 release - Changes to this document include: • Added section for Installation Quick Start • Updated references to `-G` option, and ability for installer to auto-detect GPU setup • Added ability for installer to auto-detect MOFED installation and set `MOFED_PATH` • The following sections were updated with information regarding auto-detection: — Download and Install Intel GPU Software (Optional) — Download and Install NVIDIA CUDA Software (Optional) — Download and Install DOCA OFED Software (Optional) — Before You Begin — Using the INSTALL Command Line Options — Install Using the TUI Menus — Install Using CLI Commands — Reinstall Using TUI Menus — Reinstall Using CLI Commands — Reinstall Using Linux Distribution Software Packages Provided by Intel — Before You Begin under Install Host Software on Remaining Servers — Before You Begin under Upgrade the Intel® Ethernet Fabric Suite Software — Intel® EFS-Basic Pre-Installation Checklist — Intel® EFS-FS Pre-Installation Checklist — Intel® Basic-IB Pre-Installation Checklist
September 2023	1.6	Product 11.5.1.0 release - Changes to this document include: • Updated Download and Install Intel GPU Software (Optional) section to include `i915` driver parameters needed for Direct GPU Access. • Added note to Before You Begin about upgrading host software on all nodes when installing with GPU or MOFED support. • Updated title of Download and Install NVIDIA CUDA Software (Optional) section to include 'CUDA' in the name to clarify topic. • Updated Using the INSTALL Command Line Options: — Replaced the word 'ULP' with 'component' to be more accurate. — Reordered arguments to match order in the product. • Assorted grammatical, formatting and style improvements through the document

continued...

Date	Revision	Description
M day 2023	1.5	Product 11.5.0.0 release - Changes to this document include: • Refined Overview section. • Updated Download and Install NVIDIA CUDA Software (Optional) to include info about compatible CUDA Runtime and a link to how to rebuild if needed. • Added note to rebuild Intel® EFS Kernel module after updating dependancies defined in Download and Install Intel GPU Software (Optional), Download and Install NVIDIA CUDA Software (Optional), and Download and Install DOCA OFED Software (Optional) • Updated examples for Using the INSTALL Command Line Options as well as giving reasons why these might be useful. • Added note to Before You Begin about installing host software on remaining nodes when installing with GPU or MOFED support.
February 2023	1.4	Product 11.4.1.0 release - Changes to this document include: • Added information about the Intel® Basic Software for InfiniBand package. Intel Basic IB Package, Upgrade the Intel® Basic-IB Software, Intel® Basic-IB Checklists • Update Install component names to new Basic-IB PackageInstall Using the TUI Menus • Update meta package names. Install Using Linux Distribution Software Packages Provided by Intel and Repository Deployment into the Environment • Updated Intel GPU installation instructions: Download and Install Intel GPU Software (Optional)
October 2022	1.3	Product 11.4.0.0 release - Changes to this document include: • Renamed component *Eth RDMA* to *Eth RoCE* • Updated download instructions Download the Intel® Ethernet Fabric Suite Software • Updated installation and verification instructions with fabric plane support • Updated download and installation instructions: Download and Install NVIDIA CUDA Software (Optional) • Added download and installation instructions: Download and Install Intel GPU Software (Optional) and Download and Install DOCA OFED Software (Optional) • Added technical preview for Ubuntu and Intel® GPU support
March 2022	1.2	Product 11.2.0.0 release - Changes to this document include: • Added instructions for obtaining and installing software required for Intel® Ethernet 800 Series PCIe Adapter NICs. Download and Install Intel® Ethernet Drivers (Optional)
July 2021	1.1	Product 11.1.0.0 release - Changes to this document include: • Added NVIDIA GPU and CUDA support. • Introduced new INSTALL component eth_module.
February 2021	1.0	Product 11.0.0.0 release - Initial pubic release.

Contents

Revision History....3

Preface.... 10

Intended Audience.... 10
Intel® Ethernet Fabric Suite Documentation Library.... 10
How to Search the Intel® Ethernet Fabric Suite Documentation Set....11
Documentation Conventions.... 11
Best Practices.... 12
License Agreements....12
Technical Support....12

1.0 Introduction....13

1.1 Document Organization.... 13

2.0 Overview.... 14

2.1 Intel® Ethernet Fabric Suite Overview....14
2.1.1 Network Interface Card....16
2.2 Intel® Ethernet Fabric Suite Software Overview.... 16
2.3 Installation Overview....18
2.4 Installation Packages....18
2.4.1 Intel Eth Basic Package....19
2.4.2 Intel Eth FS Package.... 19
2.4.3 Intel Basic IB Package.... 20

Part 1: Installing the Software.... 21

3.0 Installation Quick Start....22

4.0 Installation Getting Started.... 24

4.1 Pre-Installation Requirements....24
4.1.1 Fabric Design Prerequisites....24
4.1.2 Fabric Setup Prerequisites.... 25
4.1.3 OS Packages Installation Prerequisites.... 27
4.1.4 MPI Compiler Prerequisites....27
4.1.5 Performance Tuning Prerequisites....27
4.2 Download and Install Intel® Ethernet Drivers (Optional)....27
4.3 Download and Install Intel GPU Software (Optional).... 28
4.4 Download and Install NVIDIA CUDA Software (Optional).... 29
4.5 Download and Install DOCA OFED Software (Optional)....31
4.6 Download the Intel® Ethernet Fabric Suite Software.... 32
4.7 Unpack the Tar File.... 32

5.0 Install the Intel® Ethernet Fabric Suite Software....34

5.1 Before You Begin.... 34
5.1.1 Using the INSTALL Command Line Options.... 35
5.2 Install Using the TUI Menus....39
5.3 Install Using CLI Commands....41
5.4 Install Using Linux Distribution Software Packages Provided by Intel....42

5.4.1 Repositories Included in the IEFS Package....42
5.4.2 Repository Deployment into the Environment....43
5.4.3 Using IEFS Repository on Linux OS....44
5.5 Reinstall Software After a Kernel Update....46
5.5.1 Reinstall Using TUI Menus....46
5.5.2 Reinstall Using CLI Commands....47
5.5.3 Reinstall Using Linux Distribution Software Packages Provided by Intel....48
5.6 Install Kernel Module with DKMS....51
5.7 Verifying Intel® EFS GPU Software is Installed....51

Part 2: Configuring the Software....53

6.0 Configuration Getting Started....54
6.1 Pre-Configuration Requirements....54
6.1.1 Chassis Configuration Prerequisites....54
6.2 How to Use the FastFabric TUI....55

7.0 Install Host Software on Remaining Servers....58
7.1 Before You Begin....58
7.2 Install the Host Software on the Remaining Hosts Using the FastFabric TUI Menu....60
7.3 Verify the Host Software on the Remaining Servers Using the FastFabric TUI Menu....63
7.4 Install the Host Software on the Remaining Hosts Using the FastFabric CLI Commands...67
7.5 Verify the Host Software on the Remaining Servers Using CLI Commands....69

8.0 Perform Initial Health Check....73
8.1 Before You Begin....73
8.2 Configure and Initialize Health Check Tools Using FastFabric TUI Menu....73
8.3 Configure and Initialize Health Check Tools Using FastFabric CLI Commands....75

9.0 Installation Verification and Additional Settings....76
9.1 ARP Neighbor Table Setup for Large Clusters....76

Part 3: Upgrading the Software....77

10.0 Upgrade Getting Started....78
10.1 Upgrade Prerequisites....78
10.2 Download the Intel® Ethernet Fabric Suite Software....78
10.3 Unpack the Tar File....79

11.0 Upgrade the Intel® Ethernet Fabric Suite Software....80
11.1 Before You Begin....80
11.2 Upgrade Scenarios....80
11.3 Upgrade the Intel® EFS-Basic Software....80
11.4 Upgrade the Intel® EFS-FS Software....81
11.4.1 Upgrade the Software on the Management Nodes....82
11.4.2 Upgrade the Software on the Remaining Servers....82
11.5 Upgrade from Intel® EFS-Basic to Intel® EFS-FS Software Package....84
11.6 Upgrade the Intel® Basic-IB Software....85

Appendix A Software Installation Checklists....87
A.1 Intel® EFS-Basic Checklists....87

A.1.1 Intel® EFS-Basic Pre-Installation Checklist..... 87
A.1.2 Install the Intel® EFS-Basic Software Checklist..... 88
A.1.3 Upgrade the Intel® EFS-Basic Software Checklist..... 88
A.2 Intel® EFS-FS Checklists.....89
A.2.1 Intel® EFS-FS Pre-Installation Checklist..... 89
A.2.2 Install and Configure the Intel® EFS-FS Software Checklist.....89
A.2.3 Upgrade the Intel® EFS-FS Software Checklist..... 90
A.3 Intel® Basic-IB Checklists..... 91
A.3.1 Intel® Basic-IB Pre-Installation Checklist.....91
A.3.2 Install the Intel® Basic-IB Software Checklist.....91
A.3.3 Upgrade the Intel® Basic-IB Software Checklist.....92

Appendix B Intel® EFS Software Components to Packages Mapping..... 93

Appendix C Intel® EFS Containers Installation.....94

Figures

1 Intel® EFS Fabric.....15
2 Intel® EFS Host Fabric Software Stack.....16
3 Intel® EFS Fabric and Software Components..... 17

Tables

1 Intel® EFS Software Components Mapping..... 93

Preface

This manual is part of the documentation set for the Intel® Ethernet Fabric Suite Fabric (Intel® EFS Fabric), which is an end-to-end solution consisting of Network Interface Cards (NICs), fabric management, and diagnostic tools.

The Intel® EFS Fabric delivers the next generation, High-Performance Computing (HPC) network solution that is designed to cost-effectively meet the growth, density, and reliability requirements of HPC and AI training clusters.

Intended Audience

The intended audience for the Intel® Ethernet Fabric Suite (Intel® EFS) document set is network administrators and other qualified personnel.

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	*Intel® Ethernet Fabric Suite Software Installation Guide*	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	*Intel® Ethernet 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 Text-based User Interface (TUI) menus and command line interface (CLI) commands.
Running MPI applications on Intel® EFS Running middleware that uses Intel® EFS	*Intel® Ethernet Fabric Suite Host Software User Guide*	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

continued...

Task	Document Title	Description
		this document includes cluster administrators and those running or implementing Message-Passing Interface (MPI) programs.
Optimizing system performance	*Intel® Ethernet Fabric Performance Tuning Guide*	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	*Intel® Ethernet Fabric Suite Software Release Notes*	

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.

Documentation Conventions

The following conventions are standard for Intel® Ethernet Fabric Suite documentation:

- *Note:* provides additional information.
- ***Caution:*** indicates the presence of a hazard that has the potential of causing damage to data or equipment.
- ***Warning:*** indicates the presence of a hazard that has the potential of causing personal injury.
- Text in blue font indicates a hyperlink to a figure, table, or section in this guide. Links to websites are also shown in blue. For example:

  See License Agreements for more information.

  For more information, visit www.intel.com.
- Text in **bold** font indicates user interface elements such as menu items, buttons, check boxes, key names, key strokes, or column headings. For example:

Click the **Start** button, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.

Press **CTRL+P** and then press the **UP ARROW** key.

- Text in `Courier` font indicates a file name, directory path, or command line text. For example:

  Enter the following command: `sh ./install.bin`

- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:

  Refer to *Intel® Ethernet Fabric Suite Software Installation Guide* for details.

  In this document, the term *chassis* refers to a managed switch.

Procedures and information may be marked with one of the following qualifications:

- **(Linux)** – Tasks are only applicable when Linux is being used.
- **(Host)** – Tasks are only applicable when Intel® Ethernet Host Software or Intel® Ethernet Fabric Suite is being used on the hosts.
- Tasks that are generally applicable to all environments are not marked.

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.

License Agreements

This software is provided under one or more license agreements. Refer to the license agreement(s) provided with the software for specific detail. Do not install or use the software until you have carefully read and agree to the terms and conditions of the license agreement(s). By loading or using the software, you agree to the terms of the license agreement(s). If you do not wish to so agree, do not install or use the software.

Technical Support

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 details.

1.0 Introduction

This guide provides instructions for installing the Intel® Ethernet Fabric Suite (Intel® EFS) software and configuring the system for the first time. It also provides instructions for upgrading the software.

For details about the other documents for the Intel® Ethernet Fabric Suite product line, refer to Intel® Ethernet Fabric Suite Documentation Library in this document.

You install the software using one of the following methods:

- Guided installation using Text User Interface (TUI) menus (recommended)
- Command Line Interface (CLI) commands
- Linux Distribution Software

1.1 Document Organization

This manual is organized as follows:

- Introduction
- Overview

Part 1: Installing the Software

- Installation Getting Started
- Install Intel® Ethernet Fabric Suite Software

Part 2: Configuring the Software

- Configuration Getting Started
- Install Host Software on Remaining Servers
- Perform Initial Health Check
- Installation Verification and Additional Settings

Part 3: Upgrading the Software

- Upgrade Getting Started
- Upgrading the Intel® Ethernet Fabric Suite Software

Appendices

- Software Installation Checklists
- Intel® EFS Software Components to Packages Mapping

2.0 Overview

This section provides an overview of the Intel® Ethernet Fabric Suite and software installation process.

2.1 Intel® Ethernet Fabric Suite Overview

The Intel® Ethernet Fabric Suite (Intel® EFS) interconnect fabric design enables a broad class of multiple node computational applications requiring scalable, tightly-coupled processing, memory, and storage resources. With open standard APIs developed by the OpenFabrics Alliance (OFA) Open Fabrics Interface (OFI) workgroup, NICs in the Intel® EFS family are optimized to provide the low latency, high bandwidth, and high message rate needed by High Performance Computing (HPC) and AI training applications.

The following figure shows a sample Intel® EFS-based fabric, consisting of different types of nodes and servers.

Figure 1. Intel® EFS Fabric

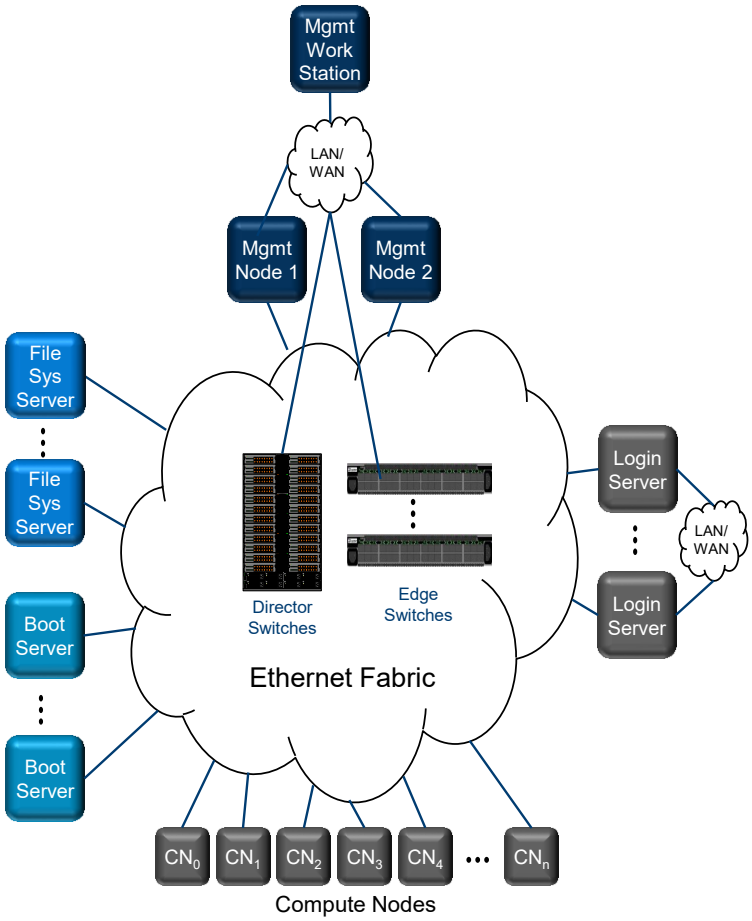

The software ecosystem is built around OFA software and includes three key APIs.

1. The OFA OFI represents a long-term direction for high-performance user-level and kernel-level network APIs.
2. OFA Verbs provides support for existing remote direct memory access (RDMA) applications.
3. Sockets are supported and permits many existing applications to immediately run on Intel® Ethernet Fabric Suite as well as provide TCP/IP features such as IP routing and network bonding.

Higher-level communication libraries, such as the Message Passing Interface (MPI), are layered on top of these low level OFA APIs. This permits existing HPC applications to immediately take advantage of advanced Intel® Ethernet Fabric Suite features.

Intel® Ethernet Fabric Suite combines the Network Interface Card (NIC), standard third-party Ethernet switches, and fabric management tools into an end-to-end solution. The host fabric software stack is shown in the following figure.

Figure 2. Intel® EFS Host Fabric Software Stack

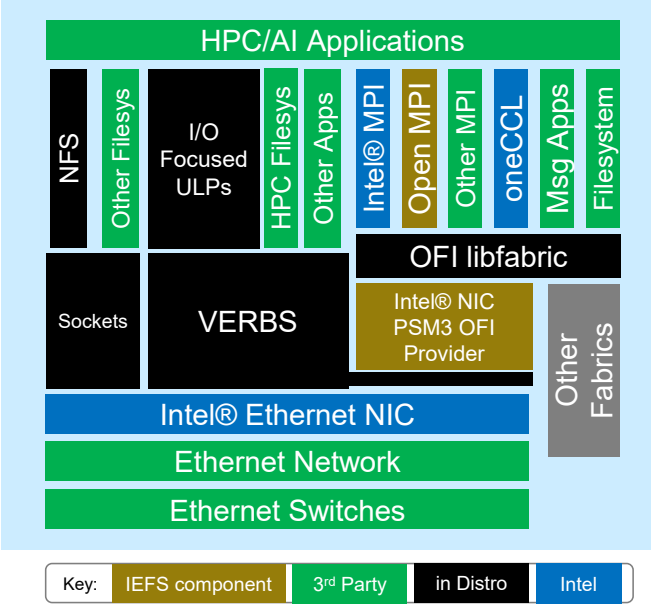

2.1.1 Network Interface Card

Each host is connected to the fabric through a Network Interface Card (NIC). The NIC translates instructions between the host processor and the fabric. It includes the logic necessary to implement the physical and link layers of the fabric architecture, so that a node can attach to a fabric and send and receive packets to other servers or devices. NICs also include specialized logic for executing and accelerating upper layer protocols, such as RDMA transport layers.

2.2 Intel® Ethernet Fabric Suite Software Overview

For software applications, Intel® EFS maintains consistency and compatibility with existing standard APIs through the open source OpenFabrics Alliance (OFA) software stack on Linux distribution releases.

Software Components

The key software components and their usage models are shown in the following figure and described in the following table.

Figure 3. Intel® EFS Fabric and Software Components

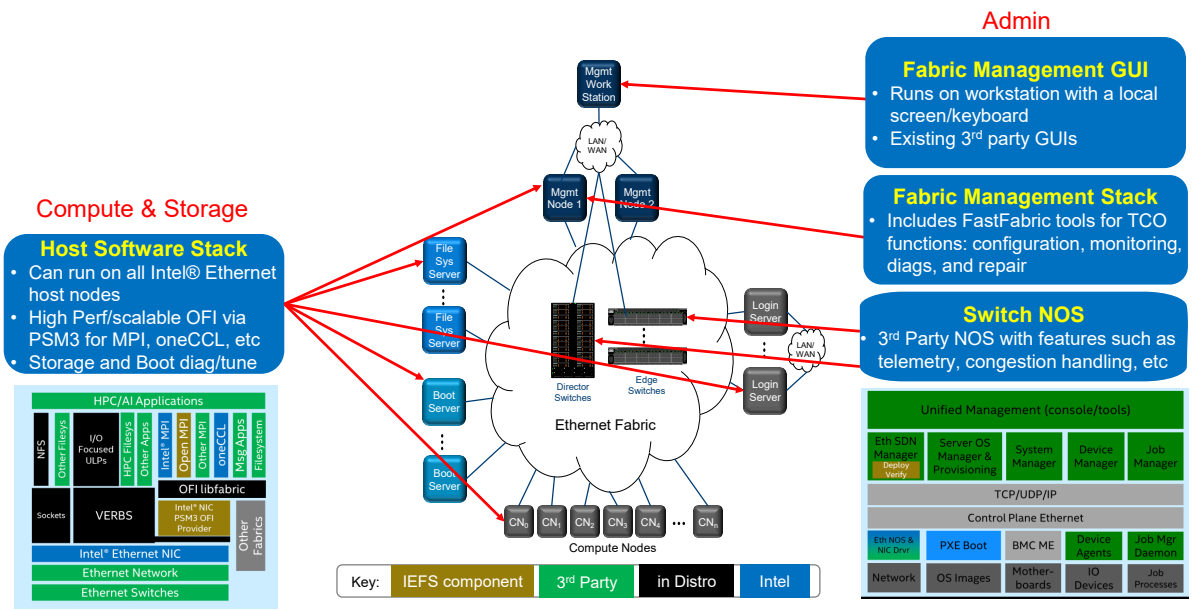

Software Component Descriptions
Switch Network Operating System (NOS) Intel® EFS supports a variety of third-party NOS solutions on standard Ethernet switches. Each of these switches may include features such as: • An embedded processor that runs switch management and control functions. • System management capabilities, including signal integrity, thermal monitoring, and voltage monitoring, among others. • Ethernet port access using command line interface (CLI) or graphical user interface (GUI).
Host Software Stack • Runs on all Intel® EFS-connected host nodes and supports compute, management, and I/O nodes. • Provides a rich set of APIs including OFI, sockets, and OFA verbs. • Provides high performance, highly scalable MPI implementation through the Intel PSM3 OFI (also known as libfabric) provider, and multiple MPI middlewares. • Includes Boot over Fabric mechanism for configuring a server to boot over the Intel® Ethernet Fabric using the NIC Unified Extensible Firmware Interface (UEFI) firmware. User documents: • *Intel® Ethernet Fabric Suite Host Software User Guide* • *Intel® Ethernet Fabric Performance Tuning Guide*
Fabric Management Stack Intel® EFS supports a variety of third-party Ethernet management solutions including popular Software Defined Networking (SDN) stacks. As part of the management solution, the Intel® EFS FastFabric tools are provided to aid deployment verification, fabric tuning, and diagnosis. • Runs on Intel® EFS-connected management nodes. • Includes a toolkit for configuration, monitoring, diagnostics, and repair. User documents: • *Intel® Ethernet Fabric Suite FastFabric User Guide*

2.3 Installation Overview

The standard software installation process takes you through installing the Intel® Ethernet Fabric (Intel® EFS-Basic) or Intel® Ethernet Fabric Suite (Intel® EFS-FS) software package, configuring the system, and verifying the system settings.

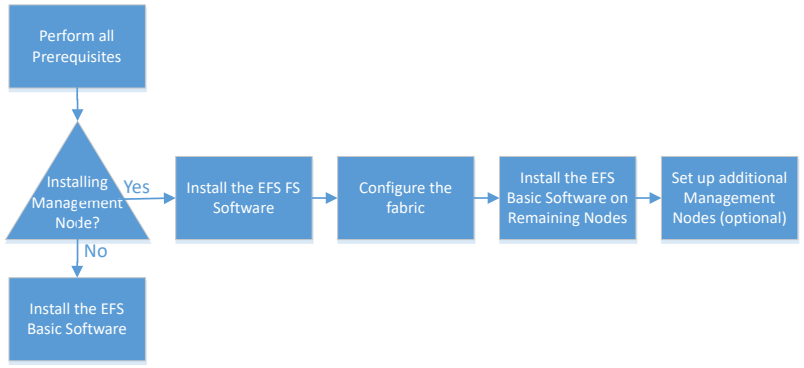

Intel recommends that you install the Intel® EFS software on the Management Node using the Install TUI, and then use FastFabric to configure the Management Node.

NOTE

Without proper configuration on the Management Node, some tools or applications may not work. For example, MPI applications may require password-less SSH, and some FastFabric functions depend on proper SNMP setup. It is crucial to configure the Management Node with FastFabric TUI or CLI commands after Intel® EFS software installation.

Once the Management Node has been configured, the Basic software can be installed on all the remaining hosts using either the FastFabric TUI or a provisioning or diskless boot mechanism.

NOTE

If you are using a provisioning system, consult the documentation that comes with the provisioning system.

2.4 Installation Packages

The following software installation packages are available for an Intel® Ethernet Fabric.

2.4.1 Intel Eth Basic Package

The `IntelEth-Basic.`*`DISTRO.VERSION`*`.tgz` installation package (where *`DISTRO.VERSION`* is the OS distribution and software version) installs the Intel® Ethernet Host Software components needed to set up compute, I/O, and Service nodes with drivers, stacks, and basic tools for local configuration and monitoring.

The installation package includes the following components:

- Basic Tools
- PSM3
- Kernel Module
- Eth RoCE
- OpenMPI (ofi, gcc)
- MPI Source
- OFA Debug Info

NOTE

A separate Intel® Ethernet Host Software installation package is available for each of the supported Linux distributions. Refer to the release notes of the package version being installed for a list of supported Linux distributions.

Related Links

Intel EFS Software Components to Packages Mapping on page 93

2.4.2 Intel Eth FS Package

The `IntelEth-FS.`*`DISTRO.VERSION`*`.tgz` installation package (where *`DISTRO.VERSION`* is the OS distribution and software version) provides the Intel® Ethernet Host Software package along with the Intel® Ethernet Fabric Suite FastFabric toolset.

The installation package includes the following components:

- Intel Eth Basic package

  For the list of components, refer to Intel Eth Basic Package.
- Intel® Ethernet Fabric Suite FastFabric

  Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more details.

NOTE

A separate EFS installation package is available for each of the supported Linux distributions. Refer to the release notes of the version being installed for a list of supported Linux distributions.

Related Links

Intel EFS Software Components to Packages Mapping on page 93

2.4.3 Intel Basic IB Package

The `Intel-Basic-IB.`*`DISTRO.VERSION`*`.tgz` installation package (where *`DISTRO.VERSION`* is the OS distribution and software version) provides the Intel® Basic Software for InfiniBand package.

The installation package includes the following components:

- Basic Tools
- PSM3
- Kernel Module
- OFA Debug Info

NOTE

A separate installation package is available for each of the supported Linux distributions. Refer to the release notes of the version being installed for a list of supported Linux distributions.

Part 1: Installing the Software

3.0 Installation Quick Start

This section provides quick start instructions to install the Intel® Ethernet Fabric Suite Software

Installing Prerequisites

The Intel® Ethernet Fabric Suite Software Release Notes document provides the list of prerequisite packages that are required to be installed on the system before IEFS can be installed.

The INSTALL program will report any missing dependencies.

Default Installation - CPU

To install the default setup for CPU-Only servers, enter the following command:

```
./INSTALL -a
```

Default Installation - GPU

To install the default setup for servers with Intel® or NVIDIA GPUs, enter the following command:

```
./INSTALL -a -G
```

The INSTALL program will automatically detect GPU drivers and software, and install the appropriate IEFS components for that type of GPU.

Default Installation - User Space

To install the default setup for only user-space components, as is performed when installing into a container, enter the following command:

```
./INSTALL -a [-G] --user-space
```

For more information, see Intel® EFS Containers Installation.

As above, if GPU support is desired, use '-G' on the command line as well.

Default Installation - specifying a specific kernel

To install the default setup to be aligned with a specific kernel, not necessarily the running kernel, enter the following command:

```
./INSTALL -a -R <kernel-version> [-G]
```

Further Information

For more detailed instructions for advanced installation capabilities, consult the section for Install the Intel® Ethernet Fabric Suite Software. Further adjustments (post-install) may be made using the "iefsconfig" tool.

4.0 Installation Getting Started

This section provides instructions and information for getting started with the Intel® Ethernet Fabric Suite Software installation.

4.1 Pre-Installation Requirements

This section provides the information and procedures needed prior to installing the fabric software. Typically, the Site Implementation Engineer performs the setup tasks described in this section to ensure that the fabric is ready for the software installation.

4.1.1 Fabric Design Prerequisites

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

It is important that the design and installation of the hardware be planned carefully prior to the installation and setup of the fabric. The design plan must include the following information:

- Identification of servers that will function as the administration or Management Nodes, where the Intel® Ethernet Fabric Suite (IEFS) will be installed.
  - Server memory requirements based on the software being used.
  - Swap disk space allowance should follow recommendations for the given version of Linux.
  - Intel recommends, but does not require the following:
    - Using Intel® Xeon® Processor dual-socket server.
    - 32 GB or more of ECC memory (ideally fully populating DIMM channels on processor to insure maximum memory bandwidth)
    - 256 GB or more of storage (ideally RAIDed for resiliency)
- Plan the cabling of the fabric and create a cable planning spreadsheet using the sample `.xlsx` files installed into `/usr/share/eth-tools/samples/` on the management node. After a plan is established, `ethxlattopology` may be used to convert the spreadsheet into a topology XML file that can be used by FastFabric during fabric verification.

  NOTE

  The use of cable planning spreadsheets and the resulting topology XML files is highly recommended by Intel as an effective and efficient way to ensure the cluster is assembled and installed as intended.

  For more information, refer to the **ethxlattopology** section of the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

- Plan the naming conventions for hosts and switches in the fabric. Intel recommends all switches and hosts be given unique names. Having unique names simplifies operations that are performed using host and switch names.

4.1.2 Fabric Setup Prerequisites

Ensure that the following requirements are met prior to installing and setting up the fabric.

NOTE

For information about the configuration files used by FastFabric, refer to *Intel® Ethernet Fabric Suite FastFabric User Guide*.

1. Ensure that all hardware is installed:
   - Servers
   - Core and edge switches
   - Fabric cables
2. Ensure that a Network Interface Card (NIC) is installed in each server.
3. The hardware configuration should be reviewed to ensure everything has been installed according to the plan.

   Following the software installation, Intel® Ethernet Fabric Suite FastFabric tools may be used to help verify the installation.
4. **(Linux)** Ensure that the required Operating System (OS) version (with the same kernel version) is installed on all hosts with the following options:
   - Root user command prompt ends in "`#`" or "`$`" with a space after it.
   - Fancy and colored prompts are disabled.
   - TCL and Expect packages are installed on all Fabric Management Nodes.

   The Management Node(s) that run FastFabric should have a full installation and must include the TCL and Expect OS packages.

   For MPI clusters, install the C and Fortran compilers, along with their associated tools, on each Management Node.

   NOTE

   Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for a list of supported OS versions.

5. **(Linux)** Enable remote login as `root` on each host.

   In order for FastFabric to manage the hosts, the Management Node must be able to securely log in as `root` to each host. This can be accomplished using SSH.

   NOTE

   FastFabric includes the `ethsetupssh` tool, which can help perform the key exchange to establish secure password-less login from the FastFabric node to all other nodes in the fabric. To simplify the use of this tool, Intel recommends initially configuring the same root password on all hosts. After root access through SSH has been set up using FastFabric, the root passwords can be changed.

6. Resolve the TCP/IP Host Names.

   FastFabric and TCP/IP must resolve host names to the management network.

Name resolution is accomplished by configuring a DNS server on the management network, with both management network and Ethernet addresses for each host.

Alternatively, an `/etc/hosts` file needs to be created on the Management Node; FastFabric can then propagate this `/etc/hosts` file to all the other hosts.

If you are using the `/etc/hosts` file approach and *not* using Domain Name System (DNS):

- On the master node, add all the Ethernet addresses into the `/etc/hosts` file.
- The `localhost` line should not be edited.
- The `/etc/hosts` file should not have any node-specific data.
- Copy the file to every node, as follows:

```
ethscpall -p -f hostfile /etc/hosts /etc/hosts
```

If you are using Domain Name System (DNS):

- Refer to the documentation for the domain name system (DNS) server being used. Make sure to edit the `/etc/resolv.conf` configuration file on the Management Node to use the proper DNS server.
- Refer to the Linux OS documentation for more information about configuring the `/etc/resolv.conf` file. This file is typically configured during OS installation.
- If `/etc/resolv.conf` must be manually configured for each host, FastFabric can aid in copying the file to all the hosts.
- The `/etc/resolv.conf` file created on the Management Node must not have any node-specific data and must be appropriate for use on all hosts.
- Copying the `/etc/resolv.conf` file to all the nodes is accomplished during the OS installation.
- If the `/etc/resolv.conf` file was not set up on all the hosts during the OS installation, the **FastFabric Copy a file to all hosts** operation can be used during the Install Host Software on Remaining Servers procedures to copy the `/etc/resolv.conf` file from the Management Node to all the other nodes.

7. Set up a Network Time Protocol (NTP) server.

   Configure an NTP server for the cluster, and set all Linux hosts and internally managed chassis to sync to the NTP server.

8. Set up sshd config.

   In configurations with more than four NICs per server, it may be necessary to increase the sshd *MaxSessions* and *MaxStartups* settings to avoid failures of some FastFabric tests due to premature closing of ssh sessions by the destination server.

   Adjust `/etc/ssh/sshd_config` with the following lines:

```
MaxSessions 50
MaxStartups 50:30:100
```

Related Links

Install Host Software on Remaining Servers on page 58

4.1.3 OS Packages Installation Prerequisites

If you are using the INSTALL script for installation, OS packages must be installed before you can install the Intel® Ethernet Fabric Suite software.

Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for the list of OS packages required for each supported OS in this release.

4.1.4 MPI Compiler Prerequisites

The Open MPI that is downloaded with the Intel® EFS-Basic Software package is selectable in the Intel Intel® EFS Install Menu. This MPI was built with the GCC compiler included with the relative OS distribution. To use this MPI for compiling and running your applications and benchmarks, you need to install the appropriate development libraries.

4.1.5 Performance Tuning Prerequisites

Intel recommends that you pre-configure servers and settings to tune fabric performance to meet the needs of your system. These tasks can be performed before or after the installation. Refer to the *Intel® Ethernet Fabric Performance Tuning Guide* that describes settings and parameters that have been shown to improve MPI/HPC performance on Intel® Ethernet Fabric Suite. If you are interested in benchmarking the performance of your system, these tips may help you obtain better performance.

4.2 Download and Install Intel® Ethernet Drivers (Optional)

If your cluster has Intel® Ethernet Network Adapters such as the Intel® Ethernet E810 Network Adapter, you should install the latest available drivers provided by Intel. You may also wish to upgrade to the latest adapter firmware for best stability and compatibility with the latest drivers, though it is not always required.

For Intel® Ethernet E810 Network Adapter

The following is a high-level overview of the install process that must be performed on each node in the fabric:

1. Visit Download Drivers and Software and search for "Ethernet Adapter Complete Driver Pack" (quoted) to find the latest drivers and NVM firmware. Either download the full release package (e.g. `Release_<version>.zip`) or find the individual drivers.
   - `ice` (Network Adapter Driver)
   - `irdma` (RDMA Driver)
   - *Non-Volatile Memory (NVM) Update Utility for Intel® Ethernet Network Adapters E810 Series*.
2. Once each item is downloaded follow the included instructions to install them.

NOTE

After updating `irdma`, it is also recommenced to update and patch `rdma-core` as defined in the `README` provided with `irdma-<version>.tgz`.

For Intel® Ethernet E2100 IPU Adapter

The following is a high-level overview of the install process that must be performed on each node in the fabric:

1. Visit Resource and Documentation Center (RDC) Collection for the IPU and download the latest release packages.
   - `intel-ipu-docs-<version>.tar.gz` (Latest Documentation).
   - `intel-ipu-<pkg>-<version>.tar.gz` (Other Packages, refer to documentation on which packages are needed).
2. Carefully follow the instructions to install the various components that are required for your system.

4.3 Download and Install Intel GPU Software (Optional)

Intel® Ethernet Fabric Suite Software comes with pre-built packages that support oneAPI. To take advantage of this, you must install the correct oneAPI libraries and modules beforehand.

To install Intel® oneAPI, refer to https://dgpu-docs.intel.com/installation-guides/index.html

NOTE

In newer kernels, the DMA Buf kernel module included with the kernel is sufficient. However, older kernels require an out-of-tree module (`intel-dmabuf-dkms`) included in the oneAPI install process.

Configure `i915` Driver

To use Direct GPU Access, you must load the `i915` Graphics Driver. Direct GPU Access is handled through the Intel® EFS RV Kernel Module. You need to load the `i915` Driver with the option `prelim_override_p2p_dist` set to 1.

```
options i915 prelim_override_p2p_dist=1
```

OR

```
modprobe i915 prelim_override_p2p_dist=1
```

In-Tree DMA Buf Kernel Module

The in-tree module should be installed and loaded at this point. This module is required to rebuild and/or install Intel® Ethernet Fabric Suite Software with oneAPI support.

```
export INTEL_GPU_DIRECT="CURRENT_KERNEL"
```

Out-of-Tree DMA Buf Kernel Module

The `intel-dmabuf-dkms` module should be installed and loadable at this point. This module is required to rebuild and/or install Intel® Ethernet Fabric Suite Software with oneAPI support.

Optionally, run the command `export INTEL_GPU_DIRECT=<path to Intel dmabuf module src build directory>` so the Intel® Ethernet Fabric Suite Software installer knows where to find the oneAPI module source when re-building the kernel module. If INTEL_GPU_DIRECT is not set in the environment, the installer will set it by automatically detecting GPU drivers and software installed on the system.

NOTE

The `<path to Intel dmabuf module src build directory>` is the directory where the `Module.symvers` file is present after the module is built.

If `Module.symvers` is not present, it must be built manually. To build it, run the following commands:

```
cd <path to Intel dmabuf module src build directory>
export LEX=flex
export YACC=bison
cp defconfigs/dmabuf .config
make -j KLIB=/lib/modules/$(uname -r) olddefconfig
make -j KLIB=/lib/modules/$(uname -r)
```

Common path for `INTEL_GPU_DIRECT` is:

- On most Distros:

```
export INTEL_GPU_DIRECT=/usr/src/intel-dmabuf-dkms-<version>
```

NOTE

The Intel® EFS kernel module must be rebuilt after updating the out-of-tree DMA Buf kernel module. Refer to Reinstall Software After a Kernel Update about how to rebuild.

4.4 Download and Install NVIDIA CUDA Software (Optional)

Intel® Ethernet Fabric Suite Software comes with pre-built packages which support CUDA. To take advantage of this, you must install the correct CUDA libraries and drivers beforehand.

To install CUDA, refer to http://docs.nvidia.com/cuda/cuda-installation-guide-linux/.

NOTE

In newer kernels, Open-Source GPL drivers (the Intel® EFS kernel module) cannot load symbols from non-GPL drivers (the proprietary NVIDIA driver). Therefore, the NVIDIA Open-Source GPL driver must be used. Refer to *Intel® Ethernet Fabric Suite Software Release Notes* for which OS Distributions require the Open-Source GPL driver.

NOTE

The NVIDIA Open-Source GPL driver requires running on NVIDIA Ampere Architecture or newer.

The `nvidia` or `nvidia-open` driver should be installed and loadable at this point. This driver is required to rebuild and/or install Intel® EFS with CUDA support.

Optionally, run the command `export NVIDIA_GPU_DIRECT=<path to NVIDIA driver src build directory>` so the Intel® EFS installer knows where to find the NVIDIA driver source when re-building the kernel module. If `NVIDIA_GPU_DIRECT` is not set in the environment, the installer will set it by automatically detecting GPU drivers and software installed on the system.

NOTE

The `<path to NVIDIA driver src build directory>` is the directory where the `Module.symvers` file is present after the NVIDIA driver is built.

If `Module.symvers` is not present, it must be built manually. To build it, run the following commands:

```
cd <path to NVIDIA driver src build directory>
make
```

NOTE

As of NVIDIA driver R515, the default driver installed is now the Open-Source GPL driver. The directory where the `Module.symvers` file is now moved to a new subdirectory where the NVIDIA driver is built (e.g. `/usr/src/nvidia-open-<version>/kernel-open/Module.symvers`).

Common paths for `NVIDIA_GPU_DIRECT` are:

- On RHEL with latest Open Source GPL Driver:

  ```
  export NVIDIA_GPU_DIRECT=/usr/src/nvidia-open-<version>/kernel-open
  ```

- On SLES with latest Open Source GPL Driver:

  ```
  export NVIDIA_GPU_DIRECT=/usr/src/kernel-modules/nvidia-<version>-default/
  kernel-open
  ```

NOTE

The Intel® EFS kernel module must be rebuilt after updating the NVIDIA GPU driver. Refer to Reinstall Software After a Kernel Update about how to rebuild.

CUDA Runtime Version

For the version of CUDA that Intel® Ethernet Fabric Suite Software was validated against, refer to *Intel® Ethernet Fabric Suite Software Release Notes*. Generally, CUDA-enabled applications and libraries will only run against the same major version of the library it was built against. This is also true for PSM3. Ensure the CUDA Runtime Library's version is compatible with the version PSM3 is built with. To support running on a different CUDA version, PSM3 will have to be rebuilt. Refer to chapter "Building the PSM3 rpm" in *Intel® Ethernet Fabric Suite Host Software User Guide* to learn how to rebuild PSM3.

Related Links

Verifying Intel EFS GPU Software is Installed on page 51

4.5 Download and Install DOCA OFED Software (Optional)

Intel® Ethernet Fabric Suite Software comes with packages which support NVIDIA DOCA OFED. To take advantage of this, you must install the correct MOFED libraries and modules beforehand.

NOTE

NVIDIA DOCA OFED should only be used with NVIDIA NICs.

To install DOCA OFED, refer to https://docs.nvidia.com/doca/sdk/doca+installation+guide+for+linux/index.html. Example:

```
$ echo "[doca]
name=DOCA Online Repo
baseurl=https://linux.mellanox.com/public/repo/doca/${VERSION}/${OS}/x86_64/
enabled=1
gpgcheck=0" > /etc/yum.repos.d/doca.repo

$ sudo dnf clean all
$ sudo dnf -y install doca-ofed
```

For another OS, see list here: https://developer.nvidia.com/doca-downloads?deployment_platform=Host-Server&deployment_package=DOCA-Host&target_os=Linux&Architecture=x86_64&Profile=doca-ofed

The `mlnx-ofa_kernel*` packages should be installed at this point. These packages are required to rebuild and/or install Intel® EFS with MOFED support.

Optionally, run the command `export MOFED_PATH=<path to mofed kernel src build directory>` so the Intel® EFS installer knows where to find the MOFED package source when re-building the kernel module. If `MOFED_PATH` is not set in the environment, the installer will set it by automatically detecting MOFED drivers and software installed on the system

NOTE

The `<path to mofed kernel src build directory>` is the directory where the `Module.symvers` file is present after the MOFED package is built.

If `Module.symvers` is not present, the MOFED package must be reinstalled.

The common path for `MOFED_PATH` is:

- On most Distros:

```
export MOFED_PATH=/usr/src/ofa_kernel/x86_64/$(uname -r)/
```

NOTE

The Intel® EFS kernel module must be rebuilt after updating the MOFED kernel module. Refer to Reinstall Software After a Kernel Update about how to rebuild.

4.6 Download the Intel® Ethernet Fabric Suite Software

If the OS you installed did not include the Intel® EFS packages, download the software package from the Download Drivers and Software using the following procedures.

1. Using a web browser, go to https://www.intel.com/content/www/us/en/search.html?ws=text#q=%22host%20software%22&sort=relevancy&f:@tabfilter=[Downloads].

   NOTE

   Some PDF readers may encode the URL incorrectly. If the opened webpage does not contain the Intel® EFS packages, copy the URL and paste it in a web browser.

2. From the results, select Intel® Ethernet Fabric Suite FS Package, Intel® Ethernet Fabric Suite Basic Package, or Intel® Basic Software for InfiniBand Package.
3. Click the Version drop-down arrow to select the latest version.
4. Click the **Download** button for your target OS.
5. Review the Intel Software License Agreement.
6. Click **Accept**.

   The zipped software is saved to your computer.

4.7 Unpack the Tar File

You unpack the tar file using the following procedure.

1. Open an SSH client session and log into the host where the package is being installed.
2. Copy the tar file to the `/root` directory.
3. Change directory to `/root`.

```
cd /root
```

4. Unpack the tar file.
   - For Intel® EFS-Basic, use:

```
tar xvfz IntelEth-Basic.DISTRO.VERSION.tgz
```

- For Intel® EFS-FS, use:

```
tar xvfz IntelEth-FS.DISTRO.VERSION.tgz
```

- For Intel® Basic-IB, use:

```
tar xvfz Intel-Basic-IB.DISTRO.VERSION.tgz
```

NOTE

`DISTRO` refers to the distribution and CPU (e.g., `RHEL86-x86_64`).

5.0 Install the Intel® Ethernet Fabric Suite Software

This section provides information and procedures to install the Intel® Ethernet Fabric Suite Software on the Management Node or on a host node in the fabric.

You install the software using one of the following methods:

- TUI menus (recommended)
- CLI commands
- Linux Distribution Software packages provided by Intel

Intel recommends that you install the Intel® EFS software on the Management Node using the Install TUI, and then use FastFabric to configure the Management Node.

NOTE

Without proper configuration on the Management Node, some tools or applications may not work. For example, MPI applications may require password-less SSH, and some FastFabric functions depend on proper SNMP setup. It is crucial to configure the Management Node with FastFabric TUI or CLI commands after Intel® EFS software installation.

Once the Management Node has been configured, the Basic software can be installed on all the remaining hosts using either the FastFabric TUI or a provisioning or diskless boot mechanism.

NOTE

If you are using a provisioning system, consult the documentation that comes with the provisioning system.

Use the Software Installation Checklists Appendix to track your installation of the software.

5.1 Before You Begin

Before starting the installation, perform the following:

- Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for the list of compatible operating systems and required OS Installation Prerequisites.
- Be sure you have completed all Pre-Installation Requirements.
- Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting INTEL_GPU_DIRECT, if you want Intel GPU support.
- Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting NVIDIA_GPU_DIRECT, if you want NVIDIA GPUDirect support.

- Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting MOFED_PATH, if you want NVIDIA Mellanox OFED support.
- You have downloaded and extracted the software package per Installation Getting Started.
- If you are using a customized installation through the Install CLI command, prepare your command line options. Refer to Using the INSTALL Command Line Options.

5.1.1 Using the INSTALL Command Line Options

The `./INSTALL` command for the Basic and FS installations are issued from the following directories:

- Intel Basic directory: `IntelEth-Basic.`*`DISTRO.VERSION`*
- Intel FS directory: `IntelEth-FS.`*`DISTRO.VERSION`*
- Intel Basic IB directory: `Intel-Basic-IB.`*`DISTRO.VERSION`*

Syntax

```
./INSTALL [-v|-vv] -R osver [-a|-n|-U|-u|-s|-O|-N|-i comp|
-e comp] [-G] [-E comp] [-D comp] [--user-space] [--without-depcheck]
[--rebuild] [--force] [--answer keyword=value]
```

or

```
./INSTALL -C
```

or

```
./INSTALL -V
```

Options

No option selected	Displays the Intel® EFS Software TUI.
`-a`	Installs all components and drivers with the default options.
`-n`	Installs all components and drivers with the default options, but does not change the autostart options.
`-U`	Upgrades/reinstalls all presently installed components and drivers with the default options, and does not change the autostart options.
`-i` *`comp`*	Installs the given component with the default options. This option can appear multiple times on a command line.

IMPORTANT

Using this command to upgrade or downgrade an individual component against the existing FS will update all previously installed components to the version of the individual component being installed.

`--user-space`	Skips kernel space components during installation.
`--without-depcheck`	Disables the check of OS dependencies.
`--rebuild`	Forces a rebuild of kernel module srpms.
`--force`	Forces the installation, even if the distributions do not match. Use of this option can result in undefined behaviors.
`-O`	Keeps the current modified rpm configuration file.
`-N`	Uses a new default rpm configuration file.
`-R` *osver*	Forces install for specific OS kernel version, rather than running kernel.
`-u`	Uninstalls all components and drivers with the default options.
`-s`	Enables autostart for all installed software.
`-e` *comp*	Uninstalls the given component with the default options. This option can appear multiple times on a command line.
`-E` *comp*	Enables autostart of given component. This option can appear with `-D` or multiple times on a command line.

NOTE

To control which installed software is configured for autostart, combine this option with `-a`, `-n`, `-i`, `-e`, and `-U` options.

`-D` *comp*	Disables autostart of the given component. This option can appear with `-E` or multiple times on a command line.

NOTE

To control which installed software is configured for autostart, combine this option with `-a`, `-n`, `-i`, `-e`, and `-U` options.

`-G`	Installs GPU support components (must have GPU drivers installed).

Also, either `NVIDIA_GPU_DIRECT=<DIR>` or `INTEL_GPU_DIRECT=<DIR>` may be in env. If neither is set, INSTALL will automatically detect GPU drivers and software, and install the appropriate IEFS components for that type of GPU.

NOTE

Refer to the Intel® Ethernet Fabric Suite release notes for more information on the GPU software versions and compatibility.

`-v` — Provides verbose logging. Logs to the `/var/log/iefs.log` file.

`-vv` — Provides very verbose debug logging. Logs to the `/var/log/iefs.log` file.

`-C` — Shows the list of supported component names.

`-V` — Outputs the version number of the software.

`--answer keyword=value` — Provides an answer to a question that might occur during the operation. Answers to questions that are not asked are ignored. Invalid answers result in prompting for interactive installations or use of the default for non-interactive installations.

Possible Questions:

`PFC_MODE`	PFC mode (0-Off, 1-Software DCB Willing, 2-Software DCB Unwilling, 3-Firmware DCB Willing).
`ARPTABLE_TUNNING`	Adjusts kernel ARP table size for large fabrics.
`ROCE_ON`	RoCE RDMA transport.
`LIMITS_SEL`	Resource Limits Selector.

Examples

- Start the `INSTALL` TUI and manually walk through the steps for a unique install.

```
./INSTALL
```

Refer to Install Using the TUI Menus for more info on navigating the TUI.

- Automatically install all components and update autostart options for any services required by components.

```
./INSTALL -a
```

This will install the entire CPU-enabled suite.

- Automatically install all components (choosing GPU variations if available) and update autostart options for any services required by components.

```
./INSTALL -a -G
```

This will install the entire GPU-enabled suite.

NOTE

To install with GPU support, ensure you have completed either Download and Install Intel GPU Software (Optional) or Download and Install NVIDIA CUDA Software (Optional) as well as optionally exporting either `INTEL_GPU_DIRECT` or `NVIDIA_GPU_DIRECT`.

- Automatically upgrade all components (choosing GPU variations if available and requested) and leave autostart options as is for any services required by components.

```
./INSTALL -U [-G]
```

This will upgrade to a newer version of the product.

- List the package version.

```
./INSTALL -V
```

This will provide an overall version for the package.

- Automatically install only specific components related to PSM3 and the RV Kernel Module.

```
./INSTALL -i psm3 -i eth_module
```

Useful for only installing the psm3 library and rv kernel module.

- List the valid component names that may be used in `-I`, `-e`, `-E`, or `-D` options.

```
./INSTALL -C [-G]
```

Useful for getting a quick list of components that are included in the package.

Other Information

Supported Component (`comp`) Names:

```
eth_tools, psm3, eth_module, fastfabric, eth_roce, openmpi_gcc_ofi, mpisrc,
delta_debug
```

Supported Component (`comp`) Name Aliases:

```
eth, mpi, psm_mpi, tools
```

Additional component names allowed for -E and -D options:

```
snmp
```

Components for -G (GPU) Installations:

When `INTEL_GPU_DIRECT` is set or Intel® GPU software and drivers are detected:

```
eth_tools, psm3, eth_module, fastfabric, eth_roce, openmpi_gcc_ofi, mpisrc,
delta_debug
```

When `NVIDIA_GPU_DIRECT` is set or NVIDIA GPU software and drivers are detected:

```
eth_tools psm3 eth_module fastfabric eth_roce openmpi_gcc_ofi
openmpi_gcc_cuda_ofi mpisrc delta_debug
```

5.2 Install Using the TUI Menus

You can install both the Intel® Ethernet Fabric Suite Software Basic and FS software packages using the Intel® EFS Software menu.

Procedures

Perform the following steps to install the Intel® EFS Software:

CAUTION

Do not interrupt an operation mid-process. Some operations may take a few minutes to complete.

NOTE

Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support.

NOTE

Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support.

NOTE

Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.

Step	Task/Prompt	Action
1.	At the command prompt, change directory to the location of the installation software package:	• For Basic, type the following and press **Enter**: **`cd IntelEth-Basic.DISTRO.VERSION`** • For FS, type the following and press **Enter**: **`cd IntelEth-FS.DISTRO.VERSION`** • For Basic-IB, type the following and press **Enter**: **`cd Intel-Basic-IB.DISTRO.VERSION`** where `DISTRO` refers to the distribution and CPU (e.g., `RHEL86-x86_64`).
2.	At the command prompt, start the install script.	Type `./INSTALL` and press **Enter**. *Notes:* • To install FS with GPU support, use `./INSTALL -G` • To install FS with different root directory, use `chroot`. • When the kernel version in the `chroot` environment is different from the host's kernel version, use `./INSTALL -R` to force the FS installation with the target OS kernel version.
3.	Select `1) Install/Uninstall Software`.	Type **1**.
4.	Review the items to be installed.	Accept the defaults (No action required). Type **N** to go to the next page. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Install` or `Don't Install`.
5.	Start the installation.	Type **P** to perform the actions. *Note:* This may take a few minutes.
6.	Preparing OFA *VERSION* release for Install... Rebuild OFA SRPMs (a=all, p=prompt per SRPM, n=only as needed?) [n]:	Press **Enter** to accept the default. *Note:* The system will display prompts that require your response throughout the installation.
7.	For each system prompt...	Accept the defaults by pressing **Enter** to continue. *Note:* Some of the default processes may take a few minutes to complete.
8.	When the Intel® EFS Autostart Menu displays, review the items.	Intel recommends leaving all of the Autostart selections set to the default values. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Enable` or `Disable`.
9.	Run the Intel® EFS Autostart operations.	Type **P**.
10.	For each system prompt, "Hit any key to continue..."	Press any key. *Note:* When the installation completes, you are returned to the main menu.
11.	Exit out of the TUI to the command prompt.	Type **X**.
12.	Reboot the server.	Type `reboot` and press **Enter**. ***Caution:*** Do not interrupt the reboot process. Depending on your operating system, the reboot may take a few minutes.
13.	Verify the installation was successful.	Type `iefsconfig -V` and press **Enter**.
	End Task	

5.3 Install Using CLI Commands

You can install either the Intel® Ethernet Fabric Suite Software Basic, Basic-IB, or FS software packages using the `./INSTALL` command.

The `./INSTALL` command has many options including installing single components, and enabling and disabling autostart of components. This section provides instructions for the default installation, but you can append the install command with specific options for a more customized installation.

NOTE

Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support.

NOTE

Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support.

NOTE

Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.

Procedures

Perform the following steps to install the default Intel® EFS Software configuration:

Step	Task/Prompt	Action
1.	At the command prompt, change directory to the location of the installation software package:	• For Basic, type the following and press **Enter**: **`cd IntelEth-Basic.`*`DISTRO.VERSION`*** • For FS, type the following and press **Enter**: **`cd IntelEth-FS.`*`DISTRO.VERSION`*** • For Basic-IB, type the following and press **Enter**: **`cd IntelBasic-IB.`*`DISTRO.VERSION`*** where *`DISTRO`* refers to the distribution and CPU (e.g. `RHEL86-x86_64`).
2.	At the command prompt, start the install script.	Type `./INSTALL -n` and press **Enter**. *Note:* • To install FS with GPU support, use ./INSTALL -n -G • To install FS with different root directory, use `chroot`. • When the kernel version in the `chroot` environment is different from the host's kernel version, use argument `-R` to force the FS installation with the target OS kernel version.
3.	At the command prompt, reboot the server.	Type `reboot` and press **Enter**.
	End Task	

5.4 Install Using Linux Distribution Software Packages Provided by Intel

The Intel® Ethernet Fabric Suite Software packages contains the OS-specific repository for installing the Intel® EFS software.

This section provides the instructions for installing using the package repository.

Intel introduced virtual packages to facilitate installation.

- A virtual package prefixed with `meta_` is a meta-package for an FS component in the INSTALL script. Installing a meta-package will install the corresponding component.
- A virtual package prefixed with `node_` is an alias package for a typical FS installation on an HPC node.

Default Installation Options

This installation method will install Intel® EFS packages with default options. To install with different options, set the following system environment variables in advance of the installation.

- Variable Name: ETH_ARPTABLE_TUNING

  Values:

  1 - Enable adjust kernel ARP table for large fabric (default)

  0 - Disable adjust kernel ARP table for large fabric
- Variable Name: ETH_ROCE_ON

  Values:

  1 - Enable RoCE on supported NICs (default)

  0 - Disable RoCE on supported NICs
- Variable Name: ETH_LIMITS_CONF

  Values:

  1 - Enable adjusted Memory Limit configuration (default)

  0 - Disable adjusted Memory Limit configuration

Alternatively, you can use `ethsystemconfig` to change the options after installation. Refer to *Intel® Ethernet Fabric Suite Host Software User Guide* for the details.

5.4.1 Repositories Included in the IEFS Package

The IEFS package contains the following repositories:

- IEFS_PKGS – Contains all software needed to be installed on the compute node, management node, or service node, such as storage node.
- IEFS_PKGS_ONEAPI-ZE – Contains all software needed to be installed on the node that includes Intel® GPU cards.
- IEFS_PKGS_CUDA – Contains all software needed to be installed on the node that includes NVIDIA cards.

NOTE

These three repositories cannot coexist on any node. Ensure that only one exists or is enabled.

Interoperate with the INSTALL Script

Intel recommends that you **do not** mix yum/zypper repository-based install with script-based install. Doing so may cause unexpected behaviors. However, you can switch from one install mechanism to another.

- Switching from script-based install to yum/zypper-based install:

  You can switch to the yum/zypper-based install at any time. No special actions are required.

  To switch to a yum/zupper-based install: If Intel® EFS is already partially or fully installed with the script, the yum/zypper-based install will identify the installed packages and skip them during installation.

- Switching from yum/zypper-based install to script-based install:

  The meta and alias packages of the yum/zypper-based install introduce extra dependencies on Intel® EFS packages. This could impact the script-based install because the code directly uses the rpm command for installation, which is sensitive to package dependencies.

  To switch to script-based install: Since you must remove the meta and alias packages first, Intel has improved the INSTALL script to handle this. Using `INSTALL -a`, or `-U`, or `-n` will switch to script-based installation. `INSTALL -u` will remove all packages include the meta and alias packages. Alternatively, you can manually remove the meta and alias packages with the yum/zypper command prior to starting the script-based install.

5.4.2 Repository Deployment into the Environment

The Intel® EFS tar packages contain the repository used to install the Intel® EFS software. It also includes a helper script called `ethcreaterepo` that checks and rebuilds kernel packages, creates the local repository, and recommends packages to install on each compute, management, and service nodes.

Intel recommends using this script to create the local repository ensuring that the correct version kernel packages and the proper repository is created for the GPU support requirement.

For example, if an IEFS_PKGS_CUDA repository already exists and you want to replace it with a repository for IEFS_PKGS, the script will back up and remove the IEFS_PKGS_CUDA to prevent the GPU version packages from installing unintentionally.

Once a local repository has been successfully created, you can transfer it to an enterprise repository, based on organization needs, to allow sharing it among nodes.

The following shows the usage information for `ethcreaterepo`.

```
Usage: ethcreaterepo [-G]
       ethcreaterepo -i
       ethcreaterepo --help

Create a local repo for Intel Fabric Suite packages.
```

```
Options:
  -G     create a repo with GPU Direct support
         NOTE: Need to set either NVIDIA_GPU_DIRECT or INTEL_GPU_DIRECT
               and directory must contain Module.symvers or set var equal to 'CURRENT_KERNEL'
  -i     display information about the repo it will create

  --help produce full help text

Examples:
  ethcreaterepo
  ethcreaterepo -G
  ethcreaterepo -i
```

Once the script executes successfully, it will list the packages for installation.

The following example shows the output for RHEL:

```
Repo IntelEth-FS was successfully created.

Please use the following component metapackages to install Intel software

  meta_eth_module                   : Intel Meta Package for Kernel Module
  meta_eth_module_dkms              : Intel Meta Package for Kernel Module  (DKMS version)
  meta_eth_module_userspace         : Intel Meta Package for Kernel Module  (user space only)
  meta_eth_tools                    : Intel Meta Package for Basic Tools
  meta_fastfabric                   : Intel Meta Package for FastFabric
  meta_mpisrc                       : Intel Meta Package for MPI Source
  meta_openmpi_gcc_ofi              : Intel Meta Package for OpenMPI (ofi,gcc)
  meta_psm3                         : Intel Meta Package for PSM3

To facilitate installation, Intel provides the following aliases for common component combinations:

  node_eth_mgmt             : Useful for management node. Includes all components.
  node_eth_mgmt_userspace   : Useful for container. Same as node_eth_mgmt except it's using user space version components.
  node_eth_mgmt_dkms        : DKMS version node_eth_mgmt. Requires DKMS pre-installed.
  node_eth_compute          : Useful for compute and login node. Includes all components except management (fastfabric)
  node_eth_compute_userspace : Useful for container. Same as node_eth_compute except it's using user space version components.
  node_eth_compute_dkms     : DKMS version node_eth_compute. Requires DKMS pre-installed.
  node_basic                : Provides basic features for a node. Includes all components except fastfabric and mpi components.
  node_basic_userspace      : Useful for container. Same as node_basic except it's using user space version components.
  node_basic_dkms           : DKMS version node_basic. Requires DKMS pre-installed.

Please run iefsconfig to config Intel Ethernet software after finished the installation
```

5.4.3 Using IEFS Repository on Linux OS

You can install the Intel® Ethernet Fabric Suite Software packages on Red Hat Enterprise Linux (RHEL) or SUSE Linux Enterprise Server (SLES) using the OS distribution included in the IEFS package repository and its dependencies.

Assumptions

- You have your software packages ready for installation.

  Refer to *Intel® Ethernet Fabric Suite Software Release Notes* for supported OS packages.

NOTE

Be sure you have completed Download and Install Intel GPU Software (Optional), including exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support.

NOTE

Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support.

NOTE

Be sure you have completed Download and Install DOCA OFED Software (Optional), including exporting `MOFED_PATH`, if you want Mellanox OFED support.

Procedure

Perform the following steps to install the default Intel® Ethernet Fabric Suite Software configuration:

Step	Task/Prompt	Action
Set Up the IEFS Repository		
1.	Create the local repository.	At the command prompt, type: `ethcreaterepo`.
2.	Create the local repository on nodes that need GPU support. *Note:* If you install GPU-supported packages on nodes without the necessary GPU cards, you may see performance degradation.	At the command prompt, type: `ethcreaterepo -G`. *Note:* After execution, the recommended install commands are provided.
3.	On each node, install the Intel® EFS Software. For a list of packages in specific Intel® EFS components, refer to Intel® EFS Software Components to Packages Mapping.	Type `yum install <alias>` under RHEL, or `zypper install <alias>` under SLES where <alias> is the recommended alias pkg (based on node type). *Note:* Alternatively, you can insert the install command (based on node type) in a provision script.
Configure RoCE (Ethernet only)		
4.	At the command prompt, start `iefsconfig`.	Type `iefsconfig`.
5.	Select `2) Reconfigure Eth RoCE (Ethernet only).`	Type **2**.
6.	`Eth RoCE is for Ethernet only, are you sure you want to continue? [y]:`	Press **Enter**
7.	`Enable RoCE RDMA transport (ROCE_ON)? [y]:`	Press **Enter**.
8.	`Resource Limits Selector (0-7) [5]:`	Press **Enter** or type another number depending on the fabric size and applications run on the fabric.

9.	For each interface, configure willing mode Priority Flow Control.
	`Configure interface <dev> now? [y]:`
	`Flow Control config, recommend willing mode Priority Flow Control (PFC)...`
	`Turn off Link Level Flow Control? [y]:`
	`PFC mode (0-Off, 1-Software DCB Willing, 2-Software DCB Unwilling, 3-Firmware DCB Willing)? [1]:`
		Type **y** to configure a interface.
		Press **Enter** to turn off Link Level Flow Control.
		Press **Enter** to set willing mode software DCB.

10.	Reboot the server.	Type `reboot` and press **Enter**. ***Caution:*** Do not interrupt the reboot process. Depending on your operating system, the reboot may take a few minutes.
11.	Verify the installation was successful.	Type `iefsconfig -V` and press **Enter**.
	End Task	

Next Steps

- If you are ready to configure your IEFS software, go to Configuring the Software.

5.5 Reinstall Software After a Kernel Update

If a kernel update occurs due to a distro upgrade, remove Intel® EFS-FS if it is already installed, then download and install the latest Intel® EFS-FS package for the new distro. Intel recommends that you use the following command to uninstall Intel® EFS-FS:

```
iefsconfig -u
```

If a kernel update comes from installing a patch for the current distro, such as installing a security patch, you may need to reinstall the Intel® Ethernet Fabric Suite Software to ensure the Intel® EFS kernel modules are rebuilt against the new kernel. The following provides instructions for rebuilding and reinstalling Intel® EFS kernel modules.

NOTE

You only need to reinstall the component Intel® EFS Eth module.

For Intel® EFS-FS user space only installation, such as Intel® EFS in a container, the Intel® EFS kernel module rebuild is not required because no Intel® EFS kernel module is installed.

To allow rebuilding of the kernel module, you must first install all of the OS kernel-related packages (package filenames starting with `kernel`) for the new kernel. Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for the list of packages required for each supported OS in this release.

5.5.1 Reinstall Using TUI Menus

You can reinstall the Intel® EFS Stack component using the Intel® EFS Software TUI menu.

NOTE

Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support.

NOTE

Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support.

NOTE

Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.

New Kernel is Active

If the new kernel is active, follow the instructions described in Install Using the TUI Menus.

New Kernel is Inactive

If the new kernel is not active, perform the following steps:

1. Launch the INSTALL TUI with the argument `-R <kernel_version>`. For example:

```
./INSTALL -R 4.18.0-147.el8.x86_64
```

NOTE

For GPU support, include the GPU argument `-G`:

```
./INSTALL -G -R 4.18.0-147.el8.x86_64
```

2. Continue following the instructions described in Install Using the TUI Menus.

5.5.2 Reinstall Using CLI Commands

You can reinstall the Intel® EFS Stack component using the `./INSTALL` command.

NOTE

Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support.

NOTE

Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support.

NOTE

Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.

New Kernel is Active

If the new kernel is active, follow the instructions described in Install Using CLI Commands, with the following modification:

In step 2, enter the following command to reinstall the component `eth_module`:

```
./INSTALL -i eth_module
```

NOTE

For GPU support, include the GPU argument `-G`:

```
./INSTALL -G -i eth_module
```

New Kernel is Inactive

If the new kernel is not active, follow the instructions described in Install Using CLI Commands, with the following modification:

In step 2, enter the CLI command with the argument `-R <kernel_version>`. For example:

```
./INSTALL -R 4.18.0-147.el8.x86_64 -i eth_module
```

NOTE

For GPU support, include the GPU argument `-G`:

```
./INSTALL -G -R 4.18.0-147.el8.x86_64 -i eth_module
```

5.5.3 Reinstall Using Linux Distribution Software Packages Provided by Intel

You can reinstall the Intel® EFS module using the OS distribution included in the IEFS package repository and its dependencies.

NOTE

Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support.

NOTE

Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support.

NOTE

Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.

New Kernel is Active

If the kernel is active, perform the following steps:

1. Run `ethcreaterepo` to update the Intel® EFS software repository.

   This script will check kernel version and rebuild Intel® EFS kernel modules.
2. If an enterprise repository is used, update it with the new updated local repository.
3. On each node, update the module packages:
   - For RHEL systems (using the yum command): `kmod-iefs-kernel-updates` and `iefs-kernel-updates-devel`

- For SLES systems (using the zypper command): `iefs-kernel-updates-kmp-default` and `iefs-kernel-updates-devel`

New Kernel is Inactive

NOTE

`ethcreaterepo` supports only current active kernel.

If the new kernel is not active, perform the following steps to update Intel® EFS software repository:

1. Change directory to the location of the installation software package:

```
cd IntelEth-FS.<DISTRO>.<VERSION>
```

2. Rebuild the Intel® EFS kernel module:

```
./INSTALL -R <kernel_version> -i eth_module
```

NOTE

For GPU support:

```
./INSTALL -G -R <kernel_version> -i eth_module
```

3. Navigate to the Intel® EFS software repository:

```
cd repos/IEFS_PKGS/RPMS
```

NOTE

For GPU support:

```
cd repos/IEFS_PKGS_ONEAPI-ZE/RPMS
```

or

```
cd repos/IEFS_PKGS_CUDA/RPMS
```

4. Remove Intel® EFS module package in the Intel® EFS software repository:

```
rm -f *kernel*.rpm
```

5. Link the new Intel® EFS module package to the Intel® EFS software repository:
   - For RHEL system:

```
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
kmod-iefs-kernel-updates-<kernel_version>-<build_number>.x86_64.rpm
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
iefs-kernel-updates-devel-<kernel_version>-<build_number>.x86_64.rpm
```

- For SLES system:

```
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
iefs-kernel-updates-kmp-default-<kernel_version>-<build_number>.x86_64.rpm
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
iefs-kernel-updates-devel-<kernel_version>-<build_number>.x86_64.rpm
```

NOTE

For NVIDIA GPU support:

- For RHEL system:

```
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
CUDA/kmod-iefs-kernel-updates-<kernel_version>-
<build_number>cuda.x86_64.rpm
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
CUDA/iefs-kernel-updates-devel-<kernel_version>-
<build_number>cuda.x86_64.rpm
```

- For SLES system:

```
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
CUDA/iefs-kernel-updates-kmp-default-<kernel_version>-
<build_number>cuda.x86_64.rpm
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
CUDA/iefs-kernel-updates-devel-<kernel_version>-
<build_number>cuda.x86_64.rpm
```

For Intel® GPU support:

- For RHEL system:

```
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
ONEAPI-ZE/kmod-iefs-kernel-updates-<kernel_version>-
<build_number>oneapize.x86_64.rpm
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
ONEAPI-ZE/iefs-kernel-updates-devel-<kernel_version>-
<build_number>oneapize.x86_64.rpm
```

- For SLES system:

```
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
ONEAPI-ZE/iefs-kernel-updates-kmp-default-<kernel_version>-
<build_number>oneapize.x86_64.rpm
ln -s ../../../IntelEth-OFA_DELTA.<DISTRO>-x86_64.<VERSION>/RPMS/<distro>/
ONEAPI-ZE/iefs-kernel-updates-devel-<kernel_version>-
<build_number>oneapize.x86_64.rpm
```

6. Rebuild the repository metadata:

```
createrepo --update ..
```

7. If an enterprise repository is used, update it with the new updated local repository.
8. On each node, update the module packages:
   - For RHEL systems (using the yum command): `kmod-iefs-kernel-updates` and `iefs-kernel-updates-devel`

- For SLES systems (using the zypper command): `iefs-kernel-updates-kmp-default` and `iefs-kernel-updates-devel`

5.6 Install Kernel Module with DKMS

With Dynamic Kernel Module Support (DKMS) installed on your system, you can install the Intel® EFS kernel module with DKMS support so that when a kernel update occurs, you do not need to reinstall Intel® Ethernet Fabric Suite Software. The DKMS framework will automatically rebuild the kernel module during the kernel update.

IMPORTANT

The kernel module rebuild may not work when you update to a new, major OS version. In this case, you must download the corresponding IEFS and reinstall it.

Prerequisites

- Install DKMS prior to performing the following steps.

  NOTE

  DKMS is not available in some Linux distributions. You must download and/or install it by yourself. For example, you can install it from the following locations:

  — RHEL: EPEL repo

  — SLES: openSUSE Tumbleweed repo (https://software.opensuse.org/package/dkms)

Procedures

Follow the instructions described in Install Using the TUI Menus or Install Using CLI Commands to install Intel® Ethernet Fabric Suite Software. When the install script detects DKMS, it will install DKMS version packages.

To install using the IEFS repository, follow the instruction described in Install Using Linux Distribution Software Packages Provided by Intel and chose the dkms version package to install.

5.7 Verifying Intel® EFS GPU Software is Installed

To ensure that the GPU capable software is installed, verify the following (note the bold text in sample outputs):

- For PSM, to verify Intel® GPU support, run:

```
$ rpm -qa | grep psm3 | grep oneapize
```

for RHEL or SLES system.

Sample output:

```
libpsm3-fi-11.4.0.0-198oneapize.x86_64
```

To verify NVIDIA GPU support, run:

```
$ rpm -qa | grep psm3 | grep cuda
```

for RHEL or SLES system.

Sample output:

```
libpsm3-fi-11.1.0.0-89cuda.x86_64
```

- For MPI, run:

```
$ /usr/mpi/gcc/openmpi-4.1.1-cuda-ofi/bin/ompi_info -a | grep
"mpi_built_with_cuda_support" | grep true
```

Sample output:

```
MCA opal base: informational "opal_built_with_cuda_support" (current value:
"true", data source: default, level: 4 tuner/basic, type: bool, synonyms:
mpi_built_with_cuda_support)
MCA mpi base: parameter "mpi_built_with_cuda_support" (current value: "true",
data source: default, level: 4 tuner/basic, type: bool, synonym of:
opal_built_with_cuda_support)
```

- For kernel module, run:

```
$ modinfo rv | grep gpu-direct
```

Sample output:

```
version: 11.1.0.0 gpu-direct
```

Part 2: Configuring the Software

6.0 Configuration Getting Started

This section provides instructions and information for getting started with the Intel® Ethernet Fabric Suite configuration.

NOTE

If you are installing the Intel® Basic-IB Software Package, then the following section is not required as Intel® Ethernet Fabric Suite FastFabric is not included in the Intel® Basic-IB Software Package.

You configure the fabric components using one of the following methods:

- Using the FastFabric TUI menus (recommended).
- Using the CLI commands.

Use the Software Installation Checklists to track your configuration and verification of the software.

6.1 Pre-Configuration Requirements

This section provides the information and procedures needed prior to configuring and verifying the fabric software.

6.1.1 Chassis Configuration Prerequisites

Ensure that the following requirements are met prior to configuring the chassis.

1. **(Switch)** Connect each chassis to the management network through its Ethernet management port. For chassis with redundant management, connect both Ethernet management ports.
2. **(Switch)** Set up the netmask and gateway addresses on each chassis.
3. **(Switch)** Assign each chassis a unique IP address, and appropriately configure the Ethernet management port network settings.
4. **(Switch)** For a chassis with redundant management, assign a unique IP address for each Management Module or Management Spine, and configure their Ethernet management port network settings.
5. **(Switch)** Select a unique name for each chassis, Management Module, and Spine. This name should be configured in DNS or `/etc/hosts` as the TCP/IP name for the Ethernet management port.
6. **(Switch)** Configure the administrator password on each chassis.

   NOTE

   For simplicity, Intel recommends that you initially configure the same administrator password on all switches.

7. **(Switch)** Configure flow control on each chassis. Intel recommends that you enable Priority Flow Control (PFC).
8. **(Switch)** To allow FastFabric to work properly, configure each chassis to enable SNMP v2c access from management nodes.

6.2 How to Use the FastFabric TUI

The FastFabric TUI menus are set up for ease of use. The submenus are designed to present operations in the order they would typically be used during an installation.

NOTE

All FastFabric TUI menu, alpha-based options are case-insensitive.

Selecting Menu Items and Performing Operations

1. From the Intel Ethernet FastFabric Tools main menu, select the target menu item (`1-2`).

```
Intel Ethernet FastFabric Tools
Version: X.X.X.X.X

   1) Host Setup
   2) Host Verification/Admin

   X) Exit (or Q)
```

   The target menu is displayed as shown in the following example:

```
FastFabric Ethernet Host Setup Menu
Host File: /etc/eth-tools/hosts
Setup:
0) Edit Management Config File               [ Skip  ]
1) Edit FF Config and Select/Edit Host File  [ Skip  ]
2) Verify Hosts Pingable                     [ Skip  ]
3) Set Up Password-Less SSH/SCP              [ Skip  ]
4) Copy /etc/hosts to All Hosts              [ Skip  ]
5) Show uname -a for All Hosts               [ Skip  ]
6) Install/Upgrade Intel Ethernet Software   [ Skip  ]
7) Configure SNMP                            [ Skip  ]
8) Build Test Apps and Copy to Hosts         [ Skip  ]
9) Reboot Hosts                              [ Skip  ]
Admin:
a) Refresh SSH Known Hosts                   [ Skip  ]
b) Rebuild MPI Library and Tools             [ Skip  ]
c) Run a Command on All Hosts                [ Skip  ]
d) Copy a File to All Hosts                  [ Skip  ]
Review:
e) View ethhostadmin Result Files            [ Skip  ]

P) Perform the Selected Actions N) Select None
X) Return to Previous Menu (or ESC or Q)
```

2. Type the key corresponding to the target menu item (`0-9`, `a-e`) to toggle the `Skip/Perform` selection.

   More than one item may be selected.
3. Type `P` to perform the operations that were selected.

NOTES

- If more than one menu item is selected, the operations are performed in the order shown in the menu. This is the typical order desired during fabric setup.
- If you want to perform operations in a different order, you must select the first target menu item, type `P` to perform the operation, then repeat this process for the next menu item operation to be performed, and so on.

4. Type `N` to clear all selected items.
5. Type `X` or press `Esc` or `Q` to exit this menu and return to the Main Menu.

Aborting Operations

While multiple menu items are performing, you have an opportunity to abort individual operations as they come up. After each operation completes and before the next operation begins, you are prompted as shown:

```
Hit any key to continue...
```

- Press `Esc` or `Q` to stop the sequence of operations return to the previous menu.

  Any unperformed operations are still highlighted in the menu. To complete the selected operations, type `P`.
- Press any other key to perform the next selected menu item being performed.

  This prompt is also shown after the last selected item completes, providing an opportunity to review the results before the screen is cleared to display the menu.

Submenu Configuration Files

On each FastFabric submenu, item `0` allows reviewing and editing (using the editor selected by the EDITOR environment variable) of the `mgt_config.xml` file to specify planes in a fabric and SNMP query parameters. Item `1` permits a different file to be selected and edited. It also permits reviewing and editing of the `ethfastfabric.conf` file. The `ethfastfabric.conf` file guides the overall configuration of FastFabric and describes cluster-specific attributes of how FastFabric operates.

At the top of each FastFabric submenu screen beneath the title, the directory and configuration file containing the components on which to operate are shown.

In the following example, the configuration file is noted in bold.

```
FastFabric Ethernet Host Setup Menu
Host File: /etc/eth-tools/hosts
Setup:
0) Edit Management Config File               [ Skip  ]
1) Edit FF Config and Select/Edit Host File  [ Skip  ]
2) Verify Hosts Pingable                     [ Skip  ]
```

NOTE

During the execution of each menu selection, the actual FastFabric command line tool being used is shown. This can be used as an educational aid to learn the command line tools.

The following example snippet shows how the CLI is displayed in the TUI execution.

```
Performing Host Setup: Verify Hosts Pingable
Would you like to verify hosts are ssh-able? [n]:y
Executing: /usr/sbin/ethfindgood -A -f /etc/eth-tools/hostes
```

7.0 Install Host Software on Remaining Servers

This section provides information and procedures to install, configure, and verify the Intel® Ethernet Host Software on the remaining hosts.

NOTE

Intel® Ethernet Fabric Suite FastFabric is also used to install the Intel® Ethernet Fabric Stack Tools on the remaining hosts when using other variations of OFA. In this case, OFA software must be installed on each host prior to installing the Intel® EFS software.

7.1 Before You Begin

Before starting the host installation and configuration, perform the following:

- Gather your information for the configuration files, as needed.

  Note that you can edit the files before you configure the hosts. Configuration files are located under the `/etc/eth-tools` directory. Sample files are installed into `/usr/share/eth-tools/samples` with the suffix `-sample`.

  The following file is used to configure the hosts:

  — `ethfastfabric.conf`: Lists the default settings for most of the FastFabric command line options.

  NOTE

  During setup of password-less SSH, FastFabric provides the opportunity to enter the host root password interactively when needed. Therefore, Intel recommends that you do not place it within the `ethfastfabric.conf` file.

  If you are required to keep the root password for the hosts in the `ethfastfabric.conf` file, Intel recommends that you change the `ethfastfabric.conf` permissions to be 0x600 (root-only access).

  NOTE

  If installing with additional options (for example, with GPU support), be sure to update the `INSTALL` variables in `ethfastfabric.conf` to reflect the additional `INSTALL` arguments.

NOTE

If you intend to install remote nodes with GPU or MOFED support, optionally ensure the related environment variable(s) are already in the remote nodes' environment. This can be accomplished in several ways, but adding the variables to the `~/.bashrc` file for the root user is known to work. However, if the variables are not set in the remote nodes' environment, the INSTALL program will automatically detect GPU drivers and software, and MOFED drivers and software, and set the variables accordingly. Also, be sure to complete the pre-requirements found in Download and Install Intel GPU Software (Optional), Download and Install NVIDIA CUDA Software (Optional), or Download and Install DOCA OFED Software (Optional).

- Intel recommends that a FastFabric topology file is created as `/etc/eth-tools/topology.xml` to describe the intended topology of the fabric. The file can also augment assorted fabric reports with customer-specific information, such as cable labels and additional details about nodes, links, ports, and cables. Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information about topology verification files.

— `hosts`: List of the hosts names (the TCP/IP management network names), except the Management Node from which FastFabric is presently being run. Enter one hostname per line. You can augment a hostname with a list of network interfaces to specify the interfaces joining the fabric. If no network interfaces are defined for a host, all available interfaces on the host will be considered as part of the fabric.

For example:

```
host1
host2
```

NOTE

Do not list the Management Node itself (that is, the node where FastFabric is currently running).

If additional Management Nodes are to be used, they may be listed at this time, and FastFabric can aid in their initial installation and verification.

— `allhosts`: Lists the Management Node's hosts name (the TCP/IP management network name, for example `mgmthost`) and includes the `hosts` file.

For example:

```
mgmthost
include /etc/eth-tools/hosts
```

— `hostverify.sh`: Script to verify the configuration and performance of an individual node. This should be run using `ethverifyhosts`. It can also be run on an individual node directly.

For more details about configuration files, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

For more details about the file format of the configuration files, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

- For the following Host Setup menu item, prepare to answer questions or provide input as follows:
  - **Building Test Apps and Copying to Hosts**

    Choose MPI Directory Selection:

```
Host Setup: Build Test Apps and Copy to Hosts
MPI Directory Selection

Please Select MPI Directory:
0) /usr/mpi/gcc/openmpi-4.1.1-cuda-ofi
1) /usr/mpi/gcc/openmpi-4.1.1-ofi
2) /opt/intel/impi/2019.10.000/intel64
3) Enter Other Directory

X) Return to Previous Menu (or ESC or Q)
```

NOTE

Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for the latest supported MPI Library versions.

- For the following Host Verify menu item, prepare to answer questions or provide input as follows:
  - **Perform Single Host Verification**
    - Prior to using this operation, you must have a copy of the `hostverify.sh` in the directory pointed to by `FF_HOSTVERIFY_DIR`. If the file does not exist in that directory, copy the sample file `/usr/share/eth-tools/samples/hostverify.sh` to the directory pointed to by `FF_HOSTVERIFY_DIR`.

      Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for information on `ethverifyhosts`.
    - Determine the tests to be run and time limit (in minutes) for the tests.

7.2 Install the Host Software on the Remaining Hosts Using the FastFabric TUI Menu

You can configure the host software using the FastFabric Ethernet Host Setup menu.

Assumption

- You are logged into the management node.

Procedures

The following steps provide simplified instructions for first-time configuration of the hosts. For additional details, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

Step	Task/Prompt	Action
Configuring the Hosts		
1.	Access the **FastFabric Ethernet Host Setup menu**.	
	a) If you are not already logged into **FastFabric Ethernet Tools**, at the command prompt...	Type `ethfastfabric` and press **Enter**.
	b) Access the **FastFabric Ethernet Host Setup menu**.	Press **1**.
	c) Select menu items.	• Select items **0 – 3** and **5 – 9**. • Select item **4** if you are using `/etc/hosts` for name resolution (as opposed to using DNS).
	d) Start the operations.	Press **P**. *Note:* Each selected item is preformed in the order of the menu list.
2.	**Edit Management Config File** (menu item 0)	
	a) Edit the `mgt_config.xml` file.	Review the file with a focus on the following: • SnmpPort • SnmpCommunityString If you made any changes, save and close the file. Press any key to continue.
3.	**Edit Configuration and Select/Edit Host File** (menu item 1)	
	a) Edit the `ethfastfabric.conf` file.	Review the file with a focus on the following: • FF_PRODUCT • FF_PACKAGES • FF_INSTALL_OPTIONS • FF_UPGRADE_OPTIONS If you made any changes, save and close the file. Press any key to continue.
	b) Edit the `hosts` configuration file.	Create the file with a list of the hosts names (the TCP/IP management network names), except the Management Node from which FastFabric is presently being run. If you made any changes, save and close the file.
	c) `Do you want to edit/review/change the files? [y]:`	Type **n** and Press **Enter**.
	e) Continue to next step.	Press any key to continue.
4.	**Verify Hosts Pingable** (menu item 2)	
	• If all hosts were found...	Press any key to continue to **Set Up Password-Less SSH/SCP**.
	• If some hosts were not found...	Press **ESC** to exit the menu and review the following list for those hosts that were not found: • Host powered on and booted? • Host connected to management network? • Host management network IP address and network settings consistent with DNS or `/etc/hosts`?

continued...

Step	Task/Prompt	Action
		• Management node connected to the management network? • Management node IP address and network settings correct? • Management network itself up (including switches, routers, and others)? • Correct set of hosts listed in the hosts file? You may need to repeat the previous step to review and edit the file.
5	**Set Up Password-Less SSH/SCP** (menu item 3)	
	a) `Password for root on all hosts:`	Type the password for root on all hosts and press **Enter**.
	b) Continue to next step.	Press any key to continue.
6.	**Copy /etc/hosts to All Hosts** (menu item 4)	
	a) Continue to next step.	Press any key to continue.
7.	**Show uname -a for All Hosts** (menu item 5)	
	a) Continue to next step.	Press any key to continue.
8.	**Install/Upgrade Intel Ethernet Software** (menu item 6)	*Note:* An initial installation uninstalls any existing Intel® EFS software. Initial installs must be performed when installing on a clean system or on a system that has Intel® EFS software installed. To upgrade the fabric, refer to Upgrade the Software on the Remaining Servers.
	a) `Do you want to use ./IntelEth-[Basic|FS].DISTRO.VERSION.tgz? [y]:`	Press **Enter** to accept the default.
	b) `Would you like to do a fresh [i]nstall, an [u]pgrade or [s]kip this step? [u]:`	Type **i** and press **Enter**.
	c) `Are you sure you want to proceed? [n]:`	Type **y** and press **Enter**.
	d) Complete the installation.	Press any key to continue. *Note:* This step may take several minutes to complete.
	• If all hosts install...	Press any key to continue.
	• If any hosts fails to install...	Use the **View ethhostadmin Results Files** menu item to review the result files from the update.
9.	**Configure SNMP** (menu item 7)	
	a) For each prompt...	Provide your selections and press **Enter**.
	b) Continue to next step.	Press any key to continue.
10.	**Build Test Apps and Copy to Hosts** (menu item 8)	
	a) For each prompt...	Provide your selections and press **Enter**.
	b) Continue to next step.	Press any key to continue.
11.	**Reboot Hosts** (menu item 9)	
	a) Continue to next step.	Press any key to continue.
12.	Ensure the hosts fully reboot, as verified through ping over the management network.	Perform Step 4.

continued...

Step	Task/Prompt	Action
Optional Tasks Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information.		
13.	**Refresh SSH Known Hosts** (menu item a)	This menu item refreshes the SSH known hosts list on this server for the Management Network. This may be used to update security for this host if hosts are replaced, reinstalled, renamed, or repaired.
14.	**Rebuild MPI Library and Tools** (menu item b)	This menu item prompts you for selection of which MPI to rebuild, and provides choices as to which available compiler to use.
15.	**Run a Command on All Hosts** (menu item c)	Intel recommends that you run the `date` command on all hosts to verify that the date and time are consistent. If needed, use the **Copy a File to All Hosts** menu item to copy the appropriate files to all hosts to enable and configure NTP.
16.	**Copy a File to All Hosts** (menu item d)	A file on the local host may be specified to be copied to all selected hosts.
Verifying the Host Configuration		
17.	**View ethhostadmin Result Files** (menu item e)	
	a) `About to: vi /root/test.res /root/test.log`	Press any key to review files.
	b) `Would you like to remove test.res test.log test_tmp* and save_tmp in /root ? [n]:`	Press **Enter** to save or type **y** to remove the files.
	End Task	

Next Steps

- To verify the host software is installed and running on the remaining servers, go to Verify the Host Software on the Remaining Servers Using the FastFabric TUI Menu.

7.3 Verify the Host Software on the Remaining Servers Using the FastFabric TUI Menu

You can verify the host software using the FastFabric Ethernet Host Verification/Admin menu.

NOTE

As a result of running this sequence, a `punchlist.csv` file is produced. This file provides a cumulative summary of tests that failed and may be provided to technicians for corrective action. The file can easily be imported into spreadsheets or other tools.

Assumption

- You are logged into the management node.

Procedures

The following steps provide simplified instructions for first-time verification of the hosts. For additional details, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

Step	Task/Prompt	Action
1.	Access the **FastFabric Ethernet Host Verification/ Admin Menu**.	
	a) If you are not already logged into **FastFabric Ethernet Tools**, at the command prompt...	Type `ethfastfabric` and press **Enter**.
	b) Access the **Host Verification/Admin Menu**.	Press **2**.
	c) Select menu items.	Select items **0 – 9**.
	d) Start the operations.	Press **P**. *Note:* Each selected item is preformed in the order of the menu list.
2.	**Edit Management Config and Select Plane** (menu item 0)	
	a) Edit the `mgt_config.xml` file	Review the file with a focus on fabric plane: • Plane — Name — Enable — HostsFile — SwitchesFile If you made any changes, save and close the file.
	b) Select the fabric plane to verify if there are more than one enabled planes	Provide your selection
	c) `Are you sure you want to proceed? [n]:`	Type **y** and Press **Enter**.
	d) Continue to next step	
3.	**Edit Configuration and Select/Edit Host File** (menu item 1)	
	a) Edit the `ethfastfabric.conf` file.	Review the file with a focus on the following: • FF_TOPOLOGY_FILE • FF_DEVIATION_ARGS Review the following parameters that are used for overall fabric health checks: • FF_ANALYSIS_DIR • FF_ALL_ANALYSIS • FF_FABRIC_HEALTH If you made any changes, save and close the file. Press any key to continue.
	b) Create or edit the `allhosts` configuration file.	Create the file with the Management Node's hostname (the TCP/IP management network name, for example `mgmthost`) and include the `hosts` file. *Note:* If you have a cluster with mixed servers or NIC configurations, create a `/etc/eth-tools/*hosts` file for each type of server configuration If you made any changes, save and close the file.
	c) `Do you want to edit/review/change the files? [y]:`	Type **n** and Press **Enter**.

continued...

Step	Task/Prompt	Action
	e) Continue to next step.	Press any key to continue.
4.	**Summary of Fabric Components** (menu item 2)	
	a) After the operation completes...	Review the results against the expected configuration of the cluster. *Note:* If components are missing or degraded, or omitted links are found, they should be corrected. Subsequent steps aid in locating any such links.
	b) Continue to next step.	Press any key to continue.
5.	**Verify Hosts Pingable, SSHable and Active** (menu item 3)	
	a) For each prompt...	Provide the required information and press **Enter**.
	b) After completion of the tests, you are prompted: `Would you like to now use /etc/eth-tools/good as Host File? [y]:`	Press **Enter** to use the file or **n** to discard the file.
6.	**Perform Single Host Verification** (menu item 4)	
	a) `Would you like to edit /root/hostverify.sh and copy to hosts? [y]:`	Review the settings near the top and the list of TESTS selected. *Note:* If you have a cluster with mixed servers or NIC configurations, ensure you add the proper settings for the server configuration (NIC PCIe bus, server memory size, expected single node HPL performance for server, etc) to the `/root/hostverify.sh` file If you made any changes, save and close the file.
	b) For each prompt...	Provide the required information and press **Enter**.
	c) Start the tests.	Press any key to continue.
	d) Review the results file.	Press any key to view the file. Close the file.
	e) Repeat the **Perform Single Host Verification** for each of the host files.	*Note:* If you have a cluster with mixed servers or NIC configurations, Step 5 needs to be repeated for each `*host` file
	f) Continue to next step.	Press any key to continue.
7.	**Verify Eth Fabric Status and Topology** (menu item 5)	
	a) For each prompt...	Provide the required information and press **Enter**.
	a) Continue to next step.	Press any key to continue.
8.	**Verify Hosts Ping via RDMA** (menu item 6)	
	a) Continue to next step.	Press any key to continue.
9.	**Verify PFC via Empirical Test** (menu item 7)	
	a) Continue to next step.	Press any key to continue.
10.	**Refresh SSH Known Hosts** (menu item 8)	
	a) Continue to next step.	Press any key to continue.

continued...

Step	Task/Prompt	Action
11.	**Check MPI Performance** (menu item 9)	*Note:* This test identifies nodes whose performance is not consistent with others in the fabric. It is not intended as a benchmark of fabric latency and bandwidth. This test intentionally uses techniques to reduce test runtime.
	a) For each prompt...	Press **Enter** to select the defaults.
	• If all hosts pass...	Continue to the next step.
	• If any hosts fail...	• Carefully examine the failing hosts to verify the NIC models, PCIe slot used, BIOS settings, and any motherboard or BIOS settings related to devices on PCIe buses or slot speeds. • Also verify that the NIC and any riser cards are properly seated. Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information.
	b) Continue to next step.	Press any key to continue.
Optional Tasks Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information.		
12.	**Check Overall Fabric Health** (menu item a)	This command permits the present fabric configuration to be baselined for use in future fabric health checks. Perform this check after configuring any additional Management Nodes and establishing a healthy fabric through successful execution of all the other tests discussed in this section. If desired, a baseline of an incomplete or unhealthy fabric may be taken for future comparison after making additions or corrections to the fabric. Refer to Configure and Initialize Health Check Tools Using FastFabric CLI Commands for more information.
13.	**Start or Stop Bit Error Rate Cable Test** (menu item b)	This command performs host cable testing. The test allows for starting and stopping an extended Bit Error Rate test.
14.	**Generate All Hosts Problem Report Info** (menu item c)	This command collects configuration and status information from all hosts and generates a single `*.tgz` file that can be sent to an Intel support representative.
15.	**Run a Command on All Hosts** (menu item d)	This command runs the `ethcmdall` command. A Linux shell command may be specified to be executed against all selected hosts. You may also specify a sequence of commands separated by semicolons.
Review the Host Verification		
16.	**View ethhostadmin Result Files** (menu item e)	A file on the local host may be specified to be copied to all selected hosts.
	a) `About to: vi /root/test.res /root/test.log`	Press any key to review files.
	b) `Would you like to remove test.res test.log test_tmp* and save_tmp in /root ? [n]:`	Press **Enter** to save or type **y** to remove the files.
	End Task	

7.4 Install the Host Software on the Remaining Hosts Using the FastFabric CLI Commands

You can configure the host software using the FastFabric Ethernet Host Setup menu.

Assumption

- You are logged into the management node.

Procedures

The following steps provide simplified instructions for first-time configuration of the hosts. For additional details, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

Step	Task/Prompt	Action
1.	(Optional) Edit or review the configuration files.	
	a) Edit the `/etc/eth-tools/ethfastfabric.conf` file.	Review the file with a focus on the following: • *FF_PRODUCT* • *FF_PACKAGES* • *FF_INSTALL_OPTIONS* • *FF_UPGRADE_OPTIONS* If you made any changes, save and close the file.
	b) Create or edit the `/etc/eth-tools/hosts` configuration file.	Create the file with a list of the hosts names (the TCP/IP management network names), except the Management Node from which FastFabric is presently being run. You can augment a hostname with a list of network interfaces to specify the interfaces join the fabric. If no network interfaces defined for a host, all available interfaces on the host will be considered as part of the fabric. If you made any changes, save and close the file.
2.	Verify the hosts are pingable over the management network.	Type **`ethpingall -p`**.
	• If all hosts were found...	Continue to the next step.
	• If some hosts were not found...	Review the following list for those hosts that were not found: • Host powered on and booted? • Host connected to management network? • Host management network IP address and network settings consistent with DNS or `/etc/hosts`? • Management node connected to the management network? • Management node IP address and network settings correct? • Management network itself up (including switches, routers, and others)? • Correct set of hosts listed in the hosts file? You may need to repeat the previous step to review and edit the file.
3.	Set up secure password-less SSH, such that the Management Node can securely log into all the hosts as `root` through the management network, without requiring a password.	Type **`ethsetupssh -S -p -i "" -f hostfile`**.
4.	(Optional) Copy the `/etc/hosts` file on this host to all the other selected hosts. *Note:* DNS is being used, skip this step.	Type **`ethscpall -p -f hostfile /etc/hosts /etc/hosts`**.

continued...

Step	Task/Prompt	Action
5.	(Optional) Copy the `/etc/resolv.conf` file on this host to all the other selected hosts.	Type **`ethscpall -p -f hostfile /etc/resolv.conf /etc/resolv.conf`**.
6.	Show uname -a (OS version) on all the hosts.	Type **`ethcmdall -T 60 -f hostfile 'uname -a'`**.
7.	Install the Intel® Ethernet Host Software on all the hosts. *Note:* An initial installation uninstalls any existing Intel® EFS software. Initial installs must be performed when installing on a clean system or on a system that has Intel® EFS software installed. To upgrade the fabric, refer to Upgrade the Software on the Remaining Servers.	Type **`ethhostadmin -f hostfile -d dir load`**. By default, it looks in the current directory for the `IntelEth-Basic.DISTRO.VERSION.tgz` file.
	• If all hosts install...	Continue to the next step.
	• If any hosts fails to install...	Use the **View ethhostadmin** result files option to review the result files from the update.
8.	Configure SNMP. *Note:* This modifies the `/etc/snmp/snmpd.conf` file on each host.	Type **`ethsetsnmp -f hostfile`**.

9.	Build the MPI sample applications on the Management Node and copy the resulting object files to all the hosts.
	Note: This is in preparation for execution of MPI performance tests and benchmarks in a later step.
	Note: This option is only available when using the Intel® Ethernet Host Software Fabric Suite (FS).
		Type:
		`cp -r -p /usr/src/eth/mpi_apps/ source_dir`
		`cd source_dir`
		`export MPICH_PREFIX=path_to_mpi`
		`make clobber quick`
		`ethscpall -t -p -f hostfile source_dir dest_dir`
		Note: For Intel® GPU support, type the following before building:
		`export MPI_APPS_ONEAPI=y`
		Note: For NVIDIA GPU support, type the following before building:
		`export MPI_APPS_CUDA=y`
		`export CUDA_DIR=cuda_dir`

10.	Reboot all the selected hosts.	Type **`ethhostadmin -f hostfile reboot`**.
11.	Ensure the hosts fully reboot, as verified through ping over the management network.	Perform Step 2.
Verifying the Host Configuration		
15.	View `ethhostadmin` result files.	Type ***`editor result_dir/result_file`***. *Note:* "*`editor`*" indicates the command line editor; for example, vi. The following default files are created: • `test.res` • `test.log`
	End Task	

Next Steps

- To verify the host software is installed and running on the remaining servers, go to Verify the Host Software on the Remaining Servers Using CLI Commands.

7.5 Verify the Host Software on the Remaining Servers Using CLI Commands

You can verify the host software using the FastFabric Ethernet Host Verification/Admin menu.

NOTE

As a result of running this sequence, a `punchlist.csv` file is produced. This file provides a cumulative summary of tests that failed and may be provided to technicians for corrective action. The file can easily be imported into spreadsheets or other tools.

Assumption

- You are logged into the management node.

Procedures

The following steps provide simplified instructions for first-time verification of the hosts for a fabric plane. For additional details, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

Step	Task/Prompt	Action
1.	(Optional) Edit or review the management configuration files.	
	a) Edit the `/etc/eth-tools/mgt_config.xml` file.	Review the file with a focus on fabric plane: • Plane — Name — Enable — HostsFile — SwitchesFile
	b) Create or edit the `HostsFile` and `SwitchesFile` files for each plane	Create the `HostsFile` file with the Management Node's hosts name (the TCP/IP management network name, for example `mgmthost`) and the Ethernet device names (the ethernet interfaces which connect to the fabric, for example `mgmthost:eth2`), and include a hosts file that contains all computer node's hostnames and ethernet device names. Create the `SwitchesFile` file that includes the management network names of all the switches in the fabric. If you made any changes, save and close the file.
2.	(Optional) Edit or review the configuration files.	
	a) Edit the `/etc/eth-tools/ethfastfabric.conf` file.	Review the file with a focus on the following: • *FF_TOPOLOGY_FILE* • *FF_DEVIATION_ARGS* Review the following parameters which are used for overall fabric health checks: • *FF_ANALYSIS_DIR* • *FF_ALL_ANALYSIS* • *FF_FABRIC_HEALTH* If you made any changes, save and close the file. Press any key to continue.

continued...

Step	Task/Prompt	Action
3.	Provide a brief summary of the counts of components in the fabric plane, including how many switch chips, hosts, and links are in the fabric plane.	Type **`ethfabricinfo -p plane`** where *plane* is the fabric plane name defined in `mgt_config.xml`.
	a) After the operation completes...	Review the results against the expected configuration of the cluster. *Note:* If components are missing or degraded, or omitted links are found, they should be corrected. Subsequent steps aid in locating any such links.
4.	(Optional) Verify each host is pingable.	Type **`ethpingall -p -f hostfile`** where *hostfile* is the fabric plane `HostsFile` defined in `mgt_config.xml`.
	• If all hosts were found...	Continue to the next step.
	• If some hosts were not found...	Review the following list for those hosts that were not found: • Host powered on and booted? • Host connected to management network? • Host management network IP address and network settings consistent with DNS or `/etc/hosts`? • Management node connected to the management network? • Management node IP address and network settings correct? • Management network itself up (including switches, routers, and others)? • Correct set of hosts listed in the hosts file? You may need to repeat the previous step to review and edit the file.
5.	Verify each host is pingable, SSH-enabled, and active on the Intel® Ethernet Fabric and produces a list of good hosts meeting all criteria.	Type **`ethfindgood -f hostfile -p plane`**.
	a) Review the `ethsorthosts` files.	The following files are created in `ethsorthosts` order with all duplicates removed in the `CONFIG_DIR/` directory: • good_*plane* • alive_*plane* • running_*plane* • active_*plane* • bad_*plane*
6.	Perform a single host test on all hosts.	Type **`ethverifyhosts -k -c -u hostverify.res -T timelimit -f hostfile test`**. *Note:* If you have a cluster with mixed servers or NIC configurations, ensure you add the proper settings for the server configuration (NIC PCIe bus, server memory size, expected single node HPL performance for server, etc) to the `/root/hostverify.sh` file. Repeat the this step for each of the `*host` files.
7.	Verify Eth Fabric status and topology.	Type **`ethlinkanalysis -f hostfile -U -x snapshot_suffix all verifyall > $FF_RESULT_DIR/linkanalysis.res 2>&1`**. *Note:* The results can be seen in the `$FF_RESULT_DIR/linkanalysis.res` file. A punch list of issues is appended to the `$FF_RESULT_DIR/punchlist.csv` file.

continued...

Step	Task/Prompt	Action
8.	Verify that RDMA is properly configured and running on all the hosts.	Type **ethhostadmin -f *hostfile* rping**.
	• If successful...	Continue to next step.
	• If not successful...	Verify that the management host has RDMA installed and configured.
9.	Verify that PFC is properly configured and running on all the switches and hosts.	Type **ethhostadmin -f *hostfile* pfctest**.
	• If successful...	Continue to next step.
	• If not successful...	Verify that both switches and hosts have proper PFC configured.
10.	Refresh the SSH `known_hosts` file on the Management Node to include all the hosts.	Type **ethsetupssh -p -U -f *hostfile***.
11.	Perform a quick check of PCIe and MPI performance through end-to-end latency and bandwidth tests. *Note:* This test identifies nodes whose performance is not consistent with others in the fabric. It is not intended as a benchmark of fabric latency and bandwidth. This test intentionally uses techniques to reduce test runtime.	Type **ethhostadmin -f *hostfile* mpiperfdeviation**.
	• If all hosts pass...	Continue to the next step.
	• If any hosts fail...	• Carefully examine the failing hosts to verify the NIC models, PCIe slot used, BIOS settings, and any motherboard or BIOS settings related to devices on PCIe buses or slot speeds. • Also verify that the NIC and any riser cards are properly seated. Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information.
Optional Tasks		
12.	Baseline the present fabric configuration for use in future fabric health checks. *Note:* This should be performed after configuring any additional Management Nodes.	Type **ethallanalysis -f *hostfile* -b**.
13.	Perform host cable testing. *Notes:* • The test allows for starting and stopping an extended Bit Error Rate test. • Intel recommends that you run this test for 20-60 minutes for a thorough test. While the test is running, monitor the fabric for signal integrity or stability errors `ethreport`. Once the desired test time has elapsed, stop the test using the command:	Type **ethcabletest -f *hostfile* start**.
	a) To stop the test...	Type **ethcabletest -f *hostfile* stop**.

continued...

Step	Task/Prompt	Action
Review the Host Verification		
14.	View `ethhostadmin` result files.	Type ***`editor result_dir/result_file`***. *Note:* "*`editor`*" indicates the command line editor; for example, vi. The following default files are created: • `test.res` • `test.log`
	End Task	

8.0 Perform Initial Health Check

The Health Check tool performs the initial fabric verification and creates a baseline of the hardware and software configuration. Once a good baseline has been established, you use the tools to compare the present fabric against the baseline and check its health. Baselines are rerun when changes occur such as fabric upgrades, hardware replacements or changes, and software configuration changes.

8.1 Before You Begin

Before starting the health check configuration, perform the following:

- Be sure you have the latest *Intel® Ethernet Fabric Suite Software Release Notes* for reference.
- Gather your information for the configuration files, as needed.

  Note that you can edit the files before you configure the health check. Configuration files are located under the `/etc/eth-tools` directory. Sample files are installed into `/usr/share/eth-tools/samples` with the suffix `-sample`.

  The following files is used to configure the health check:

  — `ethfastfabric.conf`: Lists the default settings for most of the FastFabric command line options.

    Review and update the following parameters as needed:

    - FF_ANALYSIS_DIR

      This parameter should be updated to reflect the type of SM (esm or hsm).
    - FF_ALL_ANALYSIS
    - FF_FABRIC_HEALTH

  For more details about configuration files, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

  For more details about the file format of the configuration files, refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

8.2 Configure and Initialize Health Check Tools Using FastFabric TUI Menu

The health check tools may be run on one or more Management Nodes within the cluster. You set up and use the Intel® Ethernet Fabric Suite FastFabric health check tools using the FastFabric TUI Menu.

For more information about health check tools, see the detailed discussion in the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

Assumption

- You are logged into the target management node.

Procedures

Repeat the following steps on *each* Management Node that will use the health check tools.

Step	Task/Prompt	Action
1.	Access the **FastFabric Ethernet Host Verification/ Admin Menu**.	
	a) If you are not already logged into **FastFabric Ethernet Tools**, at the command prompt...	Type `ethfastfabric` and press **Enter**.
	b) Access the **Host Verification/Admin Menu**.	Press **2**.
2.	**Edit Management Config and Select Plane** (menu item 0)	Select item **0** and press **P**.
	a) Review and edit the `mgt_config.xml` file	If you made any changes, save and close the file.
	b) Select the fabric plane to check if there are more than one enabled planes	Provide your selection.
	c) `Are you sure you want to proceed? [n]:`	Type **y** and Press **Enter**.
3.	**Edit Configuration and Select/Edit Host File** (menu item 1)	Select item **1** and press **P**.
	a) Edit `ethfastfabric.conf` file.	Review the following parameters: • FF_ANALYSIS_DIR • FF_ALL_ANALYSIS • FF_FABRIC_HEALTH
4.	**Check Overall Fabric Health** (menu item a)	Select item **a** and press **P**.
	a) `Performing Host Admin: Check Overall Fabric Health Baseline present configuration? [n]:`	Press **Enter**.
5.	Check the results.	
	a) If no errors were encountered...	Continue to next step.
	b) If any errors are encountered...	Perform the following: 1. Resolve the errors. 2. Rerun steps 4 and 5 until a clean run occurs.
6.	Create a cluster configuration baseline.	
	a) **Check Overall Fabric Health** (menu item a)	Select item **a** and press **P**.
	b) `Performing Host Admin: Check Overall Fabric Health Baseline present configuration? [n]:`	Press **y** and press **Enter**.
7.	If required, schedule regular runs of `ethallanalysis` through cron or other mechanisms.	Refer to the Linux OS documentation for more information on cron. Also refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information about `ethallanalysis` and its automated use.
	End Task	

8.3 Configure and Initialize Health Check Tools Using FastFabric CLI Commands

The health check tools may be run on one or more Management Nodes within the cluster. You set up and use the Intel® Ethernet Fabric Suite FastFabric health check tools using CLI commands.

For more information, see the detailed discussion in the *Intel® Ethernet Fabric Suite FastFabric User Guide*.

Assumption

- You are logged into the target management node.

Procedures

Repeat the following steps on *each* Management Node that will use the health check tools.

Step	Task/Prompt	Action
1.	Edit `/etc/eth-tools/mgt_config.xml` file.	Review and edit the defined fabric planes and their *HostsFile* and *SwitchesFile*.
2.	Edit `/etc/eth-tools/ethfastfabric.conf` file.	Review the following parameters: • FF_ANALYSIS_DIR • FF_ALL_ANALYSIS • FF_FABRIC_HEALTH
3.	Perform a health check for a fabric plane	Type **`ethallanalysis -p plane -e`**.
4.	Check the results.	
	a) If no errors were encountered...	Continue to next step.
	b) If any errors are encountered...	Perform the following: 1. Resolve the errors. 2. Rerun steps 3 and 4 until a clean run occurs.
5.	Create a cluster configuration baseline. *Note:* This may also be done using the FastFabric TUI menu by selecting **Check Overall Fabric Health** and answering **y** to the question: `Baseline present configuration? [n]:`	Type **`ethallanalysis -p plane -b`**.
6.	If required, schedule regular runs of `ethallanalysis` through cron or other mechanisms.	Refer to the Linux OS documentation for more information on cron. Also refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information about `ethallanalysis` and its automated use.
	End Task	

9.0 Installation Verification and Additional Settings

This section provides instructions for verifying that the software has been properly installed and configured, the Intel® Ethernet Fabric drivers are loaded, and the fabric is active and ready to use. Information on NICs and performance tuning is also provided.

9.1 ARP Neighbor Table Setup for Large Clusters

On large clusters or subnets, the ARP neighbor table may overflow and produce a neighbor table overflow message to `/var/log/messages` along with other effects such as ping failing. The Intel® Ethernet Fabric Suite includes a script, enabled by default during installation, that automatically tunes the ARP Neighbor Table when invoked. The script, `eth-arptbl-tuneup`, is run once by iefs.service when a node starts or restarts, but it can also be run manually.

To run `eth-arptbl-tuneup` manually, it must be executed at the root. The syntax for the parameter options is as follows:

- `eth-arptbl-tuneup start` - Adjust kernel ARP table size
- `eth-arptbl-tuneup stop` - Restore previous configuration
- `eth-arptbl-tuneup status` - Check if original table size was changed
- `eth-arptbl-tuneup restart` - Stop then start
- `eth-arptbl-tuneup force-reload` - Stop then start
- `eth-arptbl-tuneup --help` - Usage information for the script

There are two ARP/Neighbor tables in the kernel, one for IPv4 networks and one for IPv6 networks. The operating system uses these tables for the particular network in use. There are three threshold parameters for each table: `gc_thresh1`, `gc_thresh2`, and `gc_thresh3`. You can check the present threshold level 1 by entering the command:

```
cat /proc/sys/net/ipv4/neigh/default/gc_thresh1
```

Repeat the command for `gc_thresh2` and `gc_thresh3`.

Part 3: Upgrading the Software

10.0 Upgrade Getting Started

This section provides instructions and information for getting started with the Intel® Ethernet Fabric Suite upgrade.

10.1 Upgrade Prerequisites

Prior to upgrading the Intel® EFS software, ensure the following items have been completed:

- Review the *Intel® Ethernet Fabric Suite Software Release Notes* for a list of compatible software.
- Back up the following configuration files, if applicable, in case the upgrade fails:
  - `/etc/eth-tools/*`
  - `/etc/sysconfig/eth-tools/*`
  - `/var/usr/lib/eth-tools/analysis/baseline/*`
  - Refer to the OS documentation for a list of any other OS-specific files that should be included in any backups.
- Add or remove OS packages per *Intel® Ethernet Fabric Suite Software Release Notes*, OS Installation Prerequisites.

10.2 Download the Intel® Ethernet Fabric Suite Software

If the OS you installed did not include the Intel® EFS packages, download the software package from the Download Drivers and Software using the following procedures.

1. Using a web browser, go to https://www.intel.com/content/www/us/en/search.html?ws=text#q=%22host%20software%22&sort=relevancy&f:@tabfilter=[Downloads].

   NOTE

   Some PDF readers may encode the URL incorrectly. If the opened webpage does not contain the Intel® EFS packages, copy the URL and paste it in a web browser.

2. From the results, select Intel® Ethernet Fabric Suite FS Package, Intel® Ethernet Fabric Suite Basic Package, or Intel® Basic Software for InfiniBand Package.
3. Click the Version drop-down arrow to select the latest version.
4. Click the **Download** button for your target OS.
5. Review the Intel Software License Agreement.
6. Click **Accept**.

   The zipped software is saved to your computer.

10.3 Unpack the Tar File

You unpack the tar file using the following procedure.

1. Open an SSH client session and log into the host where the package is being installed.
2. Copy the tar file to the `/root` directory.
3. Change directory to `/root`.

```
cd /root
```

4. Unpack the tar file.
   - For Intel® EFS-Basic, use:

     ```
     tar xvfz IntelEth-Basic.DISTRO.VERSION.tgz
     ```

   - For Intel® EFS-FS, use:

     ```
     tar xvfz IntelEth-FS.DISTRO.VERSION.tgz
     ```

   - For Intel® Basic-IB, use:

     ```
     tar xvfz Intel-Basic-IB.DISTRO.VERSION.tgz
     ```

NOTE

`DISTRO` refers to the distribution and CPU (e.g., `RHEL86-x86_64`).

11.0 Upgrade the Intel® Ethernet Fabric Suite Software

This section provides information and procedures to upgrade to the Intel® Ethernet Fabric Suite Software.

11.1 Before You Begin

Before starting the upgrade, perform the following:

- Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for the list of compatible operating systems.
- Be sure you have completed all Upgrade Prerequisites.
- You have downloaded and extracted the software package per Upgrade Getting Started.
- If your configuration files contain custom changes that you want to carry forward, make a backup copy of it.

NOTE

If you are upgrading nodes which have GPU and/or MOFED support, optionally ensure the related environment variable(s) are already in the environment for all nodes you plan to upgrade. Refer to Download and Install Intel GPU Software (Optional), Download and Install NVIDIA CUDA Software (Optional), and Download and Install DOCA OFED Software (Optional) for more information.

11.2 Upgrade Scenarios

The following upgrades are available:

- To upgrade to a new version of the Intel® Ethernet Fabric Suite Software for compute nodes (Basic), go to Upgrade the Intel® EFS-Basic Software.
- To upgrade to a new version of the Intel® Ethernet Fabric Suite Software for Management Nodes (FS), go to Upgrade the Intel® EFS-FS Software.
- To upgrade a Management Node from Intel® EFS-Basic to Intel® EFS-FS, go to Upgrade from Intel® EFS-Basic to Intel® EFS-FS Software Package.
- To upgrade to a new version of the Intel® Basic Software for InfiniBand for compute nodes (Basic-IB), go to Upgrade the Intel® Basic-IB Software.

11.3 Upgrade the Intel® EFS-Basic Software

You upgrade the Intel® EFS-Basic using the `IntelEth-Basic.`*`DISTRO`*`.`*`VERSION`*`.tgz` package file.

Procedures

Perform the following steps to upgrade the Intel® Ethernet Fabric Suite Software:

Step	Task/Prompt	Action
1.	At the command prompt, change directory to the location of the new Basic software package.	Type the following and press **Enter**: **`cd IntelEth-Basic.`*`DISTRO.VERSION`*** where `DISTRO` refers to the distribution and CPU (e.g., `RHEL86-x86_64`).
2.	At the command prompt, start the install script.	Type `./INSTALL` and press **Enter**.
3.	Select `1) Install/Uninstall Software`.	Type **1**.
4.	Review the items to be upgraded.	Accept the defaults. Type **N** to go to the next page. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Upgrade` or `Don't Install`.
5.	Start the upgrade.	Type **P** to perform the actions.
6.	For each system prompt...	Accept the defaults by pressing **Enter** to continue.
7.	When the Intel Ethernet Autostart Menu displays, review the items.	Intel recommends leaving all of the Autostart selections set to the default values. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Enable` or `Disable`.
8.	Run the Ethernet Autostart operations.	Type **P**.
9.	For each system prompt, "Hit any key to continue..."	Press any key. *Note:* When the installation completes, you are returned to the main menu.
10.	Exit the installation menu.	Type **X**.
11.	Reboot the server.	Type `reboot` and press **Enter**.
	End Task	

11.4 Upgrade the Intel® EFS-FS Software

You can upgrade the Intel® EFS-FS software to a new version using the Intel® EFS Software and FastFabric TUI menus.

Optimally, to upgrade the Intel® EFS-FS software, you perform the following sequence:

1. Upgrade the Management Node.
2. Reboot the Management Node.
3. Upgrade the remaining servers.
4. Reboot the remaining servers.

11.4.1 Upgrade the Software on the Management Nodes

To upgrade the Intel® Ethernet Fabric Suite Software, you first upgrade the Management Nodes using the `IntelEth-FS.`*`DISTRO`*`.`*`VERSION`*`.tgz` package file.

Procedures

Perform the following steps to upgrade the Intel® Ethernet Fabric Suite Software on each Management Node:

Step	Task/Prompt	Action
1.	At the command prompt, change directory to the location of the new IEFS software package.	For FS, type the following and press **Enter**: **`cd IntelEth-FS.`*`DISTRO.VERSION`*** where `DISTRO` refers to the distribution and CPU (e.g., `RHEL86-x86_64`).
2.	At the command prompt, start the install script.	Type `./INSTALL` and press **Enter**.
3.	Select `1) Install/Uninstall Software`.	Type **1**.
4.	Review the items to be upgraded.	Accept the defaults. Type **N** to go to the next page. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Upgrade` or `Don't Install`.
5.	Start the upgrade.	Type **P** to perform the actions.
6.	For each system prompt...	Accept the defaults by pressing **Enter** to continue.
7.	When the Intel Ethernet Autostart Menu displays, review the items.	Intel recommends leaving all of the Autostart selections set to the default values. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Enable` or `Disable`.
8.	Run the Ethernet Autostart operations.	Type **P**.
9.	For each system prompt, "Hit any key to continue..."	Press any key. *Note:* When the installation completes, you are returned to the main menu.
10.	Exit the installation menu.	Type **X**.
11.	Reboot the server.	Type `reboot` and press **Enter**.
	End Task	

Next Steps

- If you are ready to upgrade the fabric, go to Upgrade the Software on the Remaining Servers.

11.4.2 Upgrade the Software on the Remaining Servers

After upgrading the Management Nodes, you upgrade the fabric software on the remaining servers using the FastFabric Ethernet Host Setup menu.

Assumptions

- You have upgraded the Management Nodes.
- You are logged in a Management Node.

Procedures

Perform the following steps to upgrade the servers:

Step	Task/Prompt	Action
Configuring the Hosts		
1.	Access the **FastFabric Ethernet Host Setup menu**.	
	a) If you are not already logged into **Intel Ethernet FastFabric Tools**, at the command prompt...	Type `ethfastfabric` and press **Enter**.
	b) Access the **FastFabric Ethernet Host Setup menu**.	Press **1**.
	c) Select menu items.	Select items **1** (optional) and **5**.
	d) Start the operations.	Press **P**. *Note:* Each selected item is preformed in the order of the menu list.
2.	(Optional) **Edit Configuration and Select/Edit Host File** (menu item 1)	
	a) Edit the `ethfastfabric.conf` file.	Review the file and change as needed. If you made any changes, save and close the file. Press any key to continue.
	b) Edit the `hosts` configuration file.	Review the file and change as needed. If you made any changes, save and close the file.
	c) `Do you want to edit/review/change the files? [y]:`	Type **n** and Press **Enter**.
	e) Continue to next step.	Press any key to continue.
3.	**Install/Upgrade Intel Ethernet Software** (menu item 6)	
	a) `Do you want to use ./IntelEth-[Basic|FS].DISTRO.VERSION.tgz? [y]:`	Press **Enter** to accept the default.
	b) `Would you like to do a fresh [i]nstall, an [u]pgrade or [s]kip this step? [u]:`	Press **Enter** to accept the default.
	c) `Are you sure you want to proceed? [n]:`	Type **y** and press **Enter**.
	d) When the Intel Ethernet Autostart Menu displays, review the items.	Intel recommends leaving all of the Autostart selections set to the default values. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Enable` or `Disable`.
	e) Run the Ethernet Autostart operations.	Type **P**.
	f) For each system prompt, "Hit any key to continue..."	Press any key. *Note:* When the installation completes, you are returned to the main menu.
	g) Complete the installation.	Press any key to continue.
	• If all hosts install...	Press any key to continue.
	• If any hosts fails to install...	Use the **View ethhostadmin Result Files** menu item to review the result files from the update.

continued...

Step	Task/Prompt	Action
Optional Tasks Refer to the *Intel® Ethernet Fabric Suite FastFabric User Guide* for more information.		
4.	**Reboot Hosts** (menu item 9)	**(Linux)** This menu item reboots all the selected hosts and ensures they go down and come back up properly, as verified through ping over the management network. When the hosts come back up, they are running the installed Intel® Ethernet Host Software.
5.	**Run a Command on All Hosts** (menu item c)	**(Linux)** For any other setup operations that need to be performed on all hosts, this menu item executes the specified Linux shell command against all selected hosts. It can also execute a sequence of commands separated by semicolons. *Note:* Check the relevant release notes for the new Intel® Ethernet Host Software release being installed for any additional required steps.
6.	**Copy a File to All Hosts** (menu item d)	This menu item specifies a file on the local host to be copied to all selected hosts. *Note:* When using the Intel® Ethernet Fabric Suite software, you can use FastFabric to upgrade the Intel® Ethernet Fabric Suite software stack on the remaining hosts.
7.	**View ethhostadmin Result Files** (menu item e)	
	a) `About to: vi /root/test.res /root/test.log`	Press any key to review files.
	b) `Would you like to remove test.res test.log test_tmp* and save_tmp in /root ? [n]:`	Press **Enter** to save or type **y** to remove the files.
	End Task	

Next Steps

- To verify the host software has been upgraded and running on the remaining servers, go to Verify the Host Software on the Remaining Servers Using the FastFabric TUI Menu.

11.5 Upgrade from Intel® EFS-Basic to Intel® EFS-FS Software Package

You can upgrade from Intel® EFS-Basic to Intel® EFS-FS in order to install the `FastFabric` software.

Assumptions

- Intel® EFS-Basic is installed on the server.

Procedures

Perform the following steps to upgrade to Intel® EFS-FS:

Step	Task/Prompt	Action
1.	At the command prompt, change directory to the location of the new IEFS software package.	For FS, type the following and press **Enter**: **`cd IntelEth-FS.`*`DISTRO.VERSION`*** where `DISTRO` refers to the distribution and CPU (e.g. `RHEL86-x86_64`).
2.	At the command prompt, start the install script.	Type `./INSTALL` and press **Enter**.
3.	Select `1) Install/Uninstall Software`.	Type **1**.
4.	Review the items to be upgraded. *Note:* Ensure that `FastFabric` show `Install` and all other selections show `Up To Date`.	Type **N** to go to the next page. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Upgrade` or `Don't Install`.
5.	Start the upgrade.	Type **P** to perform the actions.
6.	For each system prompt...	Accept the defaults by pressing **Enter** to continue.
7.	When the Intel Ethernet Autostart Menu displays, review the items.	Intel recommends leaving all of the Autostart selections set to the default values. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Enable` or `Disable`.
8.	Run the Ethernet Autostart operations.	Type **P**.
9.	For each system prompt, "Hit any key to continue..."	Press any key. *Note:* When the installation completes, you are returned to the main menu.
10.	Exit the installation menu.	Type **X**.
11.	Reboot the server.	Type `reboot` and press **Enter**.
	End Task	

11.6 Upgrade the Intel® Basic-IB Software

You upgrade the Intel® Basic-IB using the `Intel-Basic-IB.`*`DISTRO.VERSION`*`.tgz` package file.

Procedures

Perform the following steps to upgrade the Intel® Basic Software for InfiniBand:

Step	Task/Prompt	Action
1.	At the command prompt, change directory to the location of the new Basic-IB software package.	Type the following and press **Enter**: **`cd Intel-Basic-IB.`*`DISTRO.VERSION`*** where `DISTRO` refers to the distribution and CPU (e.g., `RHEL86-x86_64`).
2.	At the command prompt, start the install script.	Type `./INSTALL` and press **Enter**.
3.	Select `1) Install/Uninstall Software`.	Type **1**.
4.	Review the items to be upgraded.	Accept the defaults.

continued...

Step	Task/Prompt	Action
		Type **N** to go to the next page. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Upgrade` or `Don't Install`.
5.	Start the upgrade.	Type **P** to perform the actions.
6.	For each system prompt...	Accept the defaults by pressing **Enter** to continue.
7.	When the Intel Ethernet Autostart Menu displays, review the items.	Intel recommends leaving all of the Autostart selections set to the default values. *Note:* If you need to change any item, enter the alphanumeric character associated with the item to toggle between `Enable` or `Disable`.
8.	Run the Ethernet Autostart operations.	Type **P**.
9.	For each system prompt, "Hit any key to continue..."	Press any key. *Note:* When the installation completes, you are returned to the main menu.
10.	Exit the installation menu.	Type **X**.
11.	Reboot the server.	Type `reboot` and press **Enter**.
	End Task	

Appendix A Software Installation Checklists

This section includes checklists to help you track tasks during fabric installation, configuration, and upgrade.

Print the appropriate checklists to monitor your progress.

A.1 Intel® EFS-Basic Checklists

A.1.1 Intel® EFS-Basic Pre-Installation Checklist

Step	Description	Complete
1.	Ensure that hardware is installed, cabled, and powered.	
2.	Ensure that a NIC is installed in each server.	
3.	The hardware configuration should be reviewed to ensure everything is installed properly, according to the plan. Refer to the local hardware configuration plan. *Note:* Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support. *Note:* Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support. *Note:* Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.	
4.	Ensure that the required operating system is installed on each server with the following options: • Root user command prompt ends in "#" or "$". *Note:* A space must appear after "#" or "$". • Fancy and colored prompts must be disabled. • TCL and Expect packages are installed on all Fabric Management Nodes. Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for supported operating systems.	
5.	Ensure the capability of remote login as root is enabled. • SSH server is enabled. • All servers are configured with the same root password.	
6.	Ensure that there is a TCP/IP Host Name Resolution. • If using `/etc/hosts`, update the `/etc/hosts` file on the Fabric Management Node. • If using DNS, all Management Network and host names are added to DNS. • The `/etc/resolv.conf` file is configured on Fabric Management Node.	
7.	If using InfiniBand hardware, ensure that the SM (Subnet Manager) is properly configured and running, so that all ports are ACTIVE and have LIDs assigned.	
8.	Ensure that an NTP server is set up.	
9.	Ensure that the sshd config is set up.	

A.1.2 Install the Intel® EFS-Basic Software Checklist

Step	Description	Complete
1.	Download and Extract the Intel® EFS-Basic Software Package. Refer to Installation Getting Started.	
2.	Install the Intel® Ethernet Fabric Suite Software.	

A.1.3 Upgrade the Intel® EFS-Basic Software Checklist

Step	Description	Complete
1.	Complete the steps in the Upgrade Prerequisites.	
2.	Download and Unpack the new Intel® EFS-Basic Software Package per Download the Intel® Ethernet Fabric Suite Software and Unpack the Tar File, respectively.	
3.	Upgrade the Intel® EFS-Basic Software on each compute node per Upgrade the Intel® EFS-Basic Software.	

A.2 Intel® EFS-FS Checklists

A.2.1 Intel® EFS-FS Pre-Installation Checklist

Step	Description	Complete
1.	Ensure that hardware is installed, cabled, and powered.	
2.	Ensure that a NIC is installed in each server.	
3.	The hardware configuration should be reviewed to ensure everything is installed properly, according to the plan. Refer to the local hardware configuration plan. *Note:* Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally)exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support. *Note:* Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support. *Note:* Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.	
4.	Ensure that the required operating system is installed on each server with the following options: • Root user command prompt ends in "`#`" or "`$`". *Note:* A space must appear after "`#`" or "`$`". • Fancy and colored prompts must be disabled. • TCL and Expect packages are installed on all Fabric Management Nodes. Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for supported operating systems.	
5.	Ensure the capability of remote login as root enabled. • SSH server is enabled. • All servers are configured with the same root password.	
6.	Ensure that there is a TCP/IP Host Name Resolution. • If using `/etc/hosts`, update the `/etc/hosts` file on the Fabric Management Node. • If using DNS, all Management Network and host names are added to DNS. • The `/etc/resolv.conf` file is configured on Fabric Management Node.	
7.	If using InfiniBand hardware, ensure that the SM (Subnet Manager) is properly configured and running, so that all ports are ACTIVE and have LIDs assigned.	
8.	Ensure that an NTP server is set up.	
9.	Ensure that the sshd config is set up.	

A.2.2 Install and Configure the Intel® EFS-FS Software Checklist

Step	Description	Complete
1.	Download and extract the Intel® EFS-FS Software Package. Refer to Installation Getting Started.	
2.	Install the Intel® Ethernet Fabric Suite Software.	
3.	Install the Host Software on the Remaining Hosts Using the FastFabric TUI Menu.	
5.	Verify the Host Software on the Remaining Servers Using the FastFabric TUI Menu.	
6.	Configure and Initialize Health Check Tools Using FastFabric CLI Commands.	

A.2.3 Upgrade the Intel® EFS-FS Software Checklist

Step	Description	Complete
1.	Complete the steps in the Upgrade Prerequisites.	
2.	Download and extract the new Intel® EFS-FS Software Packages per Upgrade Getting Started.	
3.	Upgrade the Software on the Management Nodes.	
4.	Upgrade the Software on the Remaining Servers.	
5.	Verify the Host Software on the Remaining Servers Using the FastFabric TUI Menu.	

A.3 Intel® Basic-IB Checklists

A.3.1 Intel® Basic-IB Pre-Installation Checklist

Step	Description	Complete
1.	Ensure that hardware is installed, cabled, and powered.	
2.	Ensure that a NIC is installed in each server.	
3.	The hardware configuration should be reviewed to ensure everything is installed properly, according to the plan. Refer to the local hardware configuration plan. *Note:* Be sure you have completed Download and Install Intel GPU Software (Optional), including optionally exporting `INTEL_GPU_DIRECT`, if you want Intel GPU support. *Note:* Be sure you have completed Download and Install NVIDIA CUDA Software (Optional), including optionally exporting `NVIDIA_GPU_DIRECT`, if you want NVIDIA GPUDirect support. *Note:* Be sure you have completed Download and Install DOCA OFED Software (Optional), including optionally exporting `MOFED_PATH`, if you want Mellanox OFED support.	
4.	Ensure that the required operating system is installed on each server with the following options: • Root user command prompt ends in "`#`" or "`$`". *Note:* A space must appear after "`#`" or "`$`". • Fancy and colored prompts must be disabled. • TCL and Expect packages are installed on all Fabric Management Nodes. Refer to the *Intel® Ethernet Fabric Suite Software Release Notes* for supported operating systems.	
5.	Ensure the capability of remote login as root enabled. • SSH server is enabled. • All servers are configured with the same root password.	
6.	Ensure that there is a TCP/IP Host Name Resolution. • If using `/etc/hosts`, update the `/etc/hosts` file on the Fabric Management Node. • If using DNS, all Management Network and host names are added to DNS. • The `/etc/resolv.conf` file is configured on Fabric Management Node.	
7.	Ensure that the SM (Subnet Manager) is properly configured and running, so that all ports are ACTIVE and have LIDs assigned.	
8.	Ensure that an NTP server is set up.	
9.	Ensure that the sshd config is set up.	

A.3.2 Install the Intel® Basic-IB Software Checklist

Step	Description	Complete
1.	Download and Extract the Intel® Basic-IB Software Package. Refer to Installation Getting Started.	
2.	Install the Intel® Ethernet Fabric Suite Software.	

A.3.3 Upgrade the Intel® Basic-IB Software Checklist

Step	Description	Complete
1.	Complete the steps in the Upgrade Prerequisites.	
2.	Download and Unpack the new Intel® Basic-IB Software Package per Download the Intel® Ethernet Fabric Suite Software and Unpack the Tar File, respectively.	
3.	Upgrade the Intel® Basic-IB Software on each compute node per Upgrade the Intel® Basic-IB Software.	

Appendix B Intel® EFS Software Components to Packages Mapping

To help you choose the proper packages for installation, the following table maps Intel® EFS Software components to their respective packages.

Table 1. Intel® EFS Software Components Mapping

Component	Package	Description
	iefsconfig	Configuration tool and user space initialization scripts for the Intel® Ethernet Fabric Suite FastFabric
eth_module	kmod-iefs-kernel-updates or iefs-kernel-updates-kmp-default	Kernel modules built for the Linux kernel family of processors
	iefs-kernel-updates-devel	Development headers for Intel RV driver interface
	iefs-kernel-updates-dkms	DKMS version updated kernel modules
psm3	libpsm3-fi	PSM3 libfabric provider, or PSM3 provider, is the low-level, user-level communications interface for the Intel® EFS family of products. Higher level communications interfaces, such as MPI, which are implemented to use libfabric, may make use of the PSM3 provider to provide optimized communications for the Intel EFS family of products.
tools (also eth_tools)	eth-tools-basic	Basic tools for fabric management necessary on all compute nodes
fastfabric	eth-tools-fastfabric	Tools for managing the fabric on a management node
	eth-mpi-apps	Applications and source for testing MPI performance in conjunction with eth-fastfabric or as standalone
eth_roce		A virtual component with install and uninstall behaviors so RoCE configuration can be initialized and restored
openmpi_gcc_ofi	openmpi_gcc_ofi	Tools necessary to compile, link, and run Open MPI jobs
openmpi_gcc_cuda_ofi	openmpi_gcc_cuda_ofi	Tools necessary to compile, link, and run Open MPI jobs which use CUDA
mpisrc	openmpi.src	Source RPM for openmpi
snmp		A virtual component with install and uninstall behaviors so snmp configuration can be initialized and restored

Appendix C Intel® EFS Containers Installation

This application note provides basic information for building and running containers on Linux-based computer platforms that incorporate Intel® EFS networking technology.

The basic steps are:

1. Install the latest Intel® EFS-Basic release drivers and libraries from the Intel® support site for your Linux distribution. The INSTALL program will report any missing dependencies; ensure that they are installed before retrying.
2. Decide which container technology is appropriate for your needs, install it, and run it.
    a. For Singularity, the Intel® Ethernet device is already available to a running container. Additional steps are not needed to use the Intel® Ethernet device interface.
    b. For Docker®, the Intel® Ethernet device is already available to a running container. Additional steps are not needed to use the Intel® Ethernet device interface.
3. Install your application and any additional required user space libraries and software into a container image, and run it.

User space libraries should include the user space libraries needed to interface with the Intel® EFS driver. Use the "--user-space" option to the INSTALL program to accomplish this.

For a simple recipe, Intel recommends that you use "like on like," which means running on a standard supported OS distribution, such as Rocky 9.2, with a container image based on the same kernel and OS distribution, with a matching version of the Intel® EFS-Basic release used by each, such version 11.7.0.0.110. Other combinations may work, but there is no support implied.