



Extreme Fabric Automation Deployment Guide

Version 2.4.1

9036926-01 Rev AA
March 2021



Copyright © 2021 Extreme Networks, Inc. All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: www.extremenetworks.com/company/legal/trademarks

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses. End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Table of Contents

Preface.....	5
Text Conventions.....	5
Documentation and Training.....	6
Getting Help.....	7
Subscribe to Product Announcements.....	7
Providing Feedback.....	7
About this Document.....	9
What's New in this Document.....	9
EFA Deployment Preparation.....	10
Supported Platforms and Deployment Models.....	11
EFA Requirements.....	12
General requirements.....	12
High-availability requirements.....	13
EFA Port Requirements.....	15
EFA Installation Modes.....	16
EFA Installation on TPVM.....	16
Overview.....	16
Requirements.....	16
Supported deployments.....	17
EFA Installation for Single-Node Deployments.....	18
Install EFA in a Single-Node Deployment.....	18
Install EFA on TPVM in a Single-Node Deployment	19
Deploy the OVA for EFA.....	21
EFA Installation for Multi-Node Deployments.....	23
EFA Deployment for High Availability	23
Overview.....	23
Services in high-availability mode.....	25
Install EFA in a Multi-Node Deployment.....	26
Install EFA on TPVM in a Multi-Node Deployment	27
EFA Upgrade.....	30
Supported Upgrade Paths.....	30
Upgrade EFA in a Single-node Deployment.....	31
Upgrade EFA on TPVM in a Single-node Deployment.....	32
Upgrade EFA from a Single-node to a Multi-node Deployment.....	33
Upgrade EFA in a Multi-node Deployment.....	34
Upgrade EFA on TPVM in a Multi-node Deployment.....	35
Upgrade EFA on an OVA.....	37
Upgrade SLX-OS, TPVM, and EFA Together.....	37
Upgrade Ubuntu on the EFA Host.....	38

Recover from an Upgrade Failure.....	39
Replace a Node in a Multi-node Deployment.....	40
Replace a Node in a Multi-node TPVM Deployment.....	41
EFA Uninstallation.....	42
Uninstall EFA in a Single-Node Deployment.....	42
Uninstall EFA on TPVM in a Single-Node Deployment.....	42
Uninstall EFA in a Multi-node Deployment.....	42
Uninstall EFA on TPVM in a Multi-Node Deployment.....	43
Redundant Management Network.....	44
Redundant Management Network Overview.....	44
Linux Bonding.....	44
Supported Ports.....	45
No Redundancy Period.....	45
Standby Port Rate Throughput.....	45
Enable Redundant Management.....	45
Redundant Management Data Path.....	47



Preface

This section describes the text conventions used in this document, where you can find additional information, and how you can provide feedback to us.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as ExtremeSwitching switches or SLX routers, the product is referred to as *the switch* or *the router*.

Table 1: Notes and warnings






Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
Words in <i>italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic text</i>	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member</i> [<i>member</i> ...].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and software compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other resources](#) such as white papers, data sheets, and case studies

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit www.extremenetworks.com/education/.

Getting Help

If you require assistance, contact Extreme Networks using one of the following methods:

Extreme Portal

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

The Hub

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

Call GTAC

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. For the support phone number in your country, visit: www.extremenetworks.com/support/contact

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.

Providing Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.

- Improvements that would help you find relevant information in the document.
- Broken links or usability issues.

If you would like to provide feedback, you can do so in three ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.



About this Document


[What's New in this Document](#) on page 9

What's New in this Document

The following table describes information added to this guide for the 2.4.1 software release.

Table 4: Summary of changes

Feature	Description	Described in
EFA 2.4.1 support	The following tables are updated to include 2.4.1 or 2.4.x: <ul style="list-style-type: none">• Bare Metal Deployment Models• OVA Deployment Models• TPVM Deployment Models• EFA, Neutron, and SLX-OS Compatibility• Supported Upgrade Paths	Supported Platforms and Deployment Models on page 11 Supported Upgrade Paths on page 30



EFA Deployment Preparation

[Supported Platforms and Deployment Models](#) on page 11

[EFA Requirements](#) on page 12

[EFA Port Requirements](#) on page 15

[EFA Installation Modes](#) on page 16

[EFA Installation on TPVM](#) on page 16

Supported Platforms and Deployment Models

Support includes bare metal, OVA, and TPVM deployment models, supported TPVM versions, supported SLX-OS software versions, and supported SLX devices.

Table 5: Bare Metal Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
EFA 2.2.0, 2.3.x, and 2.4.x	External server (bare metal)	More than 24	Yes	16.04 and 18.04	<ul style="list-style-type: none"> CPU: 4 cores Storage: 50 GB RAM: 8 GB

Table 6: OVA Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
EFA 2.2.0, 2.3.x, and 2.4.x (Secure mode)	External server (OVA)	More than 24	Yes	18.04	<ul style="list-style-type: none"> CPU: 4 cores Storage: 50 GB RAM: 8 GB

Table 7: TPVM Deployment Models

EFA Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
2.2.0, 2.3.x	<ul style="list-style-type: none"> SLX 9150 SLX 9250 	Up to 24	Yes	18.04	20.1.2x
2.3.x	<ul style="list-style-type: none"> SLX 9150 SLX 9250 SLX 9740 	Up to 24	Yes	18.04	20.2.2a
2.4.x	<ul style="list-style-type: none"> SLX 9150 SLX 9250 SLX 9740 	Up to 24	Yes	18.04	20.2.2b

Table 8: TPVM Software Support

TPVM Version	SLX-OS 20.1.2x	SLX-OS 20.2.1x	SLX-OS 20.2.2a	SLX-OS 20.2.2b	SLX-OS 20.2.3	SLX-OS 20.2.3x	Ubuntu Version	EFA Version
4.0.0-5	Yes	No	No	No	No	No	18.04	2.2.0
4.0.0-7, 4.0.1, 4.1.1	Yes	Yes	Yes	No	No	No	18.04	2.2.0, 2.3.x
4.1.2	No	No	No	Yes	No	No	18.04	2.4.0

Table 8: TPVM Software Support (continued)

TPVM Version	SLX-OS 20.1.2x	SLX-OS 20.2.1x	SLX-OS 20.2.2a	SLX-OS 20.2.2b	SLX-OS 20.2.3	SLX-OS 20.2.3x	Ubuntu Version	EFA Version
4.2.2	No	No	No	No	Yes	No	18.04	2.4.0
4.2.3	No	No	No	No	No	Yes	18.04	2.4.x

Table 9: IP Fabric Topology Matrix

Device	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9150	20.1.x, 20.2.x	✓				✓
SLX 9250	20.1.x, 20.2.x	✓	✓			✓
SLX 9540	20.1.x, 20.2.x	✓			✓	
SLX 9640	20.1.x, 20.2.x				✓	
SLX 9740	20.2.x		✓	✓	✓	✓

Table 10: EFA, Neutron, and SLX-OS Compatibility

EFA Version	Neutron Version	SLX-OS Version
2.3.0	2.3.0_19	20.1.2d
2.3.1	2.3.1_02	20.1.2e
2.3.2	2.3.2_02	20.1.2e, 20.2.2a
2.4.x	3.0.0-23	20.2.2b, 20.2.3, 20.2.3a, 20.2.3b

EFA Requirements

Review this topic for requirements for host names, NTP, user privileges, DNS configuration, passwordless SSH, and IP addresses.

General requirements

- **Host names:** Host names must be unique and consist of alphanumeric characters. Hyphens are the only special characters allowed. No other special characters are allowed by Kubernetes for cluster formation or by the K3s service.
- **NTP:** The server on which EFA is installed must use NTP or be synchronized to the correct time and timezone. Having the correct time and timezone ensures the following:
 - Self-signed certificates have valid start and expiration times.
 - EFA logs have the correct time stamp.
 - The K3s service starts without errors.

You can edit `/etc/systemd/timesyncd.conf` to select NTP servers in the `[Time]` section of the configuration file. The `NTP=` option takes a space-separated list of host names or IP addresses.

NTP suggests selecting as many servers as is feasible, but at least 3. Select from the pool of publicly available servers or your company's internal NTP servers. For example:

```
[Time]
NTP=0.pool.ntp.org 1.pool.ntp.org 2.pool.ntp.org 3.pool.ntp.org
```

You can use the following commands to access `timesyncd.conf` and to synchronize your changes.

```
# sudo vim /etc/systemd/timesyncd.conf
# sudo service systemd-timesyncd restart
# systemctl status systemd-timesyncd
# sudo timedatectl set-timezone <your_time_zone>
```

- **NTP:** All devices that EFA manages must use NTP to ensure easy audit trails and logging from EFA.
- **NTP:** The EFA installer allows a maximum drift of 10 seconds across nodes. If the difference is more than 10 seconds, the installer prompts you to synchronize clocks.
- **User privileges:** The user who installs EFA must be a root user or have `sudoers` privileges to ensure components are installed correctly. Installation fails if this requirement is not met.
- **DNS:** DNS configuration on the nodes must be valid or the `/etc/resolv.conf` file must be empty to ensure that the DNS resolution of Kubernetes functions correctly.
 - Ensure that `nslookup` returns the correct host name based on the IP address. For example, `nslookup node1`.
 - Ensure that the DNS servers listed in the `/etc/resolv.conf` file can resolve to the addresses of all the nodes. For example, `dig <node_hostname> +short` should return the correct IP addresses assigned to the hosts.
- **TPVM:** With the 4.0.x releases of TPVM, you can configure DNS, NTP, and LDAP as part of deploying TPVM. For more information, see "Guest OS for TPVM" in the *Extreme SLX-OS Management Configuration Guide*.
- **Netplan:** Refer to [Netplan configuration examples](#) for network configuration using Netplan.

High-availability requirements

- **OS:** All nodes in the high-availability cluster must have the same version of the operating system. For more information about supported operating systems, see [Supported Platforms and Deployment Models](#) on page 11.
- **Host names:** High-availability host names must be unique.
- **IP addresses:**
 - High-availability deployments require an extra IP address: virtual IP, cluster IP, or host IP. Ensure that this extra address is an unallocated IP address in the same subnet as the nodes that will form the cluster.
 - All nodes in the cluster must have an IP address in the same subnet as the virtual IP address.
- **SSH:** (For SLX-OS 20.2.3 and later with TPVM version 4.2.2 and later) Before installing EFA, configure SSH passwordless access between TPVM users. You can use the SLX command line and the following commands.
 - To configure a trusted peer: **device# tpvm config trusted-peer add <peer-tpvm-ip> sudo-user <tpvm-sudo-user> password <sudo-userpassword>**.
 - To display trusted peer information: **device# show tpvm config trusted-peer**.

- To remove a trusted peer: **device# tpvm config trusted-peer remove <peer-tpvm-ip> sudo-user <tpvm-sudo-user> password <sudo-userpassword>**.

**Note**

This SSH configuration applies only for the root user. There is no option for other users.

- SSH:** (For SLX-OS releases earlier than 20.2.3) Before installing EFA, configure passwordless SSH between the nodes that will form the cluster. The following is an example of configuring passwordless SSH from a remote host for two TPVMs.

In the example, the script takes in two parameters, which are the IP addresses of the TPVMs. The example assumes the availability of the public key from the remote host and the RSA keypair.

**Note**

Modify this script to suit your requirements.

```
#!/bin/bash
TPVM1_IP="$1"
TPVM2_IP="$2"
TPVM_USER="extreme"
SSH_OPTION="-o StrictHostKeyChecking=no"

echo "Setting up passwordless ssh login from this host to TPVMs..."

MY_PUB_KEY=`cat ~/.ssh/id_rsa.pub`

ssh $SSH_OPTION $TPVM_USER@$TPVM1_IP "bash -c \"echo $MY_PUB_KEY >>
/home/$TPVM_USER/.ssh/authorized_keys\""

ssh $SSH_OPTION $TPVM_USER@$TPVM2_IP "bash -c \"echo $MY_PUB_KEY >>
/home/$TPVM_USER/.ssh/authorized_keys\""

echo "Generating ssh keypairs for root on TPVMs..."

ssh $SSH_OPTION $TPVM_USER@$TPVM1_IP "bash -c \"sudo ssh-keygen -b 4096 -t rsa -q
-N '' -f /root/.ssh/id_rsa <<< y >/dev/null\""

# This could have been a mkdir -p /root/.ssh so that root's .ssh dir is present.

ssh $SSH_OPTION $TPVM_USER@$TPVM2_IP "bash -c \"sudo ssh-keygen -b 4096 -t rsa -q
-N '' -f /root/.ssh/id_rsa <<< y >/dev/null\""

echo "Setting up passwordless ssh login between TPVMs..."

TPVM1_ROOT_PUB_KEY=`ssh $SSH_OPTION $TPVM_USER@$TPVM1_IP "bash -c \"sudo
cat /root/.ssh/id_rsa.pub\""`

#TPVM2_ROOT_PUB_KEY=`ssh $SSH_OPTION $TPVM_USER@$TPVM2_IP "bash -c \"sudo
cat /root/.ssh/id_rsa.pub\""`

echo "Exchanging ssh public keys for root between TPVMs..."

#ssh $SSH_OPTION $TPVM_USER@$TPVM1_IP "bash -c \"sudo sh -c 'echo
$TPVM2_ROOT_PUB_KEY >> /root/.ssh/authorized_keys\"\""

ssh $SSH_OPTION $TPVM_USER@$TPVM2_IP "bash -c \"sudo sh -c 'echo
$TPVM1_ROOT_PUB_KEY >> /root/.ssh/authorized_keys\"\""

echo "Adding TPVM IPs for root between TPVMs as known hosts to skip first time login
prompts..."
```

```
#ssh $SSH_OPTION $TPVM_USER@$TPVM1_IP "bash -c \"sudo sh -c 'ssh-keyscan -H
$TPVM2_IP >> /root/.ssh/known_hosts' 2>/dev/null\""

ssh $SSH_OPTION $TPVM_USER@$TPVM2_IP "bash -c \"sudo sh -c 'ssh-keyscan -H
$TPVM1_IP >> /root/.ssh/known_hosts' 2>/dev/null\""

echo "Completed passwordless ssh login between TPVMs."
```

EFA Port Requirements

The following tables identify ports that must be available and not used by other services. EFA installation fails if a required port is not available.

Table 11: General port requirements

Port	Service
80	EFA HTTP requests
162	EFA SNMP notifications
443	EFA HTTPs requests
514	Syslog service
3306	MariaDB port
6443	K3S
6514	Secure syslog service
8078	Monitoring service
8079	Host authentication
10010	Containerd
30085	OpenStack service
30672	Rabbitmq
31672	Rabbitmq management

Table 12: Port requirements for high availability

Port	Service
53	Node local DNS for Kubernetes
4567	Galera cluster replication port
4568	Galera incremental state transfer
24007	GlusterFS daemon
24008	GlusterFS management
49152 through 49251	GlusterFS bricks

EFA Installation Modes

You can install EFA in secure mode or non-secure (standard) mode.

You can choose one of these modes when you install EFA:

- **Secure mode:** Traffic to EFA uses the HTTPS protocol. All non-HTTP requests are redirected to the secure port. Traffic out of EFA (toward northbound interfaces) uses TLS.
- **Standard mode:** Traffic to EFA uses the HTTP protocol. Traffic toward northbound interfaces also uses HTTP.

You cannot change a secure installation to a standard installation. Nor can you change a standard installation to a secure installation.



Note

Upgrade from secure mode to standard mode is not a supported path.

EFA Installation on TPVM

TPVM (Third-Party Virtual Machine) is a general server that resides on some Extreme SLX devices. When EFA is deployed on a TPVM, no other applications can be run on that TPVM.

Overview

In a TPVM deployment, EFA is a microservice-based fabric automation engine that leverages the K3S Kubernetes cluster as an underlying infrastructure for the EFA services deployment. You can install or upgrade the EFA application on a TPVM with one SLX-OS command.

The EFA application binary is shipped with the SLX devices, along with the binaries for SLX-OS and the TPVM. Decoupling EFA from SLX-OS allows for upgrades to EFA without a need to upgrade SLX-OS or the TPVM. EFA can be deployed on one of the SLX devices in the fabric to manage the fabric.

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models](#) on page 11.

The EFA package can be found under the `/efaboot` folder on the SLX device. For an incremental EFA image upgrade, you can copy the EFA tar file to the `/efaboot` directory on the SLX device before the deployment.

Requirements

TPVM must be installed and running on the SLX device. You can accomplish these tasks by running the **tpvm deploy** command on the SLX device.

This example configures TPVM on the management (Eth0) interface, allows DHCP to fetch the IP address and gateway, and enables passwordless SSH and sudo access.

```
device# tpvm deploy mgmt dhcp allow-pwless
```


This example configures TPVM on Eth0 with a static IP address and gateway and an administrative password of "mypassword".

```
device# tpvm deploy mgmt ipaddr <ipaddr> gw <ipaddr> admin-pwd mypassword  
confirm mypassword
```

See the *Extreme SLX-OS Command Reference* for more examples of using this command.

Supported deployments

You can install EFA on TPVM in a single-node deployment or in a multi-node deployment for high availability. For more information, see [Install EFA on TPVM in a Single-Node Deployment](#) on page 19 and [Install EFA on TPVM in a Multi-Node Deployment](#) on page 27.



EFA Installation for Single-Node Deployments

[Install EFA in a Single-Node Deployment](#) on page 18

[Install EFA on TPVM in a Single-Node Deployment](#) on page 19

[Deploy the OVA for EFA](#) on page 21

Install EFA in a Single-Node Deployment

You can install EFA in a single-node, non-TPVM deployment.

Before You Begin

Verify the following prerequisites:

- CPU: 4 cores
- Storage: 50 GB
- RAM: 8 GB
- OS: Ubuntu 16.04 or 18.04

Ensure you have configured NTP according to [EFA Requirements](#) on page 12.

About This Task

To install EFA, you must be a root user or have `sudoers` privileges.

The installation process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

Procedure

1. Download the *.tar.gz image.
2. Untar the image.

```
device# tar -xzf efa-v2.x.x.tar.gz
```

3. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
4. Change to the EFA directory.

```
device# cd efa
```

5. Run the deployment script.

```
device# source deployment.sh
```

The EFA Installer begins in a series of dialogs.

- When prompted, select **Single-node deployment** and **OK**.

**Tip**

Use arrow keys to move between options and the space bar to select an option.

- When prompted, select **Enable secure mode** or **Standard mode** and **OK**.
- When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The installation proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

- Verify the installation.
 - On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - Run the **efa status** command for concise status information.

Install EFA on TPVM in a Single-Node Deployment

You can install EFA on an SLX TPVM in a single-node deployment.

Before You Begin

The EFA tar must be available on the `/efaboot` partition of the SLX device. You need root access to the device.

About This Task

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models](#) on page 11.

The installation process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

Procedure

1. Verify that the TPVM is set up for an EFA deployment.

- a. Verify the versions of TPVM and SLX-OS.

For the latest supported version information, see [Supported Platforms and Deployment Models](#) on page 11.

```
device# show tpvm status
device# show version
device# lsb_release -a
```

- b. Verify that the TPVM has an assigned IP address.

```
device# show tpvm ip-address
```

- c. Validate that the SSH keys are uploaded.

```
device# show tpvm status
```

- d. Verify that passwordless access is configured.

```
device# show tpvm status
```

- e. Confirm NTP on the TPVM.

```
device# tpvm config ntp add server <ip>
```

- f. Verify that NTP is synchronized.

```
device# show tpvm config ntp
```

- g. If necessary, log in to TPVM and configure the NTP time zone.

```
device# tpvm config timezone
```

2. Enter SLX Linux mode.

```
device# start-shell
# cd /efaboot
```

3. Copy the EFA tar file to the SLX device.

```
# scp efa-x.x.x.tar.gz
```

4. Deploy EFA on TPVM from the SLX shell.

```
device# efa deploy
Starting "efa deploy", DO NOT hit CTRL+C
Step 1: Checking if TPVM is deployed ...
Step 2: Get IP Addressed assigned to TPVM to deploy EFA
IP Address of the TPVM 10.x.x.x
Step 3: Checking for EFA packages in /efaboot directory
Step 4: Deploying EFA package efa-2.x.x.tar.gz on 10.x.x.x
```

The EFA Installer continues in a series of dialogs.

5. When prompted, select **Single-node deployment** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

6. When prompted, select **Enable secure mode** or **Standard mode** and **OK**.

7. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The installation proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.
8. Verify the installation.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - c. Run the **efa status** command for concise status information.

Deploy the OVA for EFA

Open Virtual Appliance (OVA) is an OVF file packaged with a base image Ubuntu image and installed with EFA.

Before You Begin

- The virtual machine (VM) on which you deploy the OVA requires a network adapter with a valid IP address and DNS. You use the IP address when you configure the SLX devices to forward syslog entries to the VM. The VM needs DNS configuration to resolve the URL during setup and forwarding of events to the notification subscriber.
- The VM must be able to access devices and the notification subscriber.
- For networks without DHCP, you must assign valid, static IP addresses and DNS. Then reboot the VM. Ensure that all services are up and running before running commands.

About This Task

OVA is compatible with VMware ESXi servers and can be deployed with VMware products. For more information about supported Ubuntu versions, see [Supported Platforms and Deployment Models](#) on page 11.

Use the OVA image for new installations only.



Warning

The EFA 2.x.x OVA is not supported for Oracle VirtualBox. The syslog service requires port forwarding for ports 514 and 6514 on UDP. However, the source IP address of the syslog message will be changed from the SLX device to the host IP, which the syslog service ignores.

Procedure

1. Download the `EFA_v2.x.x_<build_number>.ova` file.
2. Run the OVA.
3. Start the VM.

The credentials for the OVA installation are one of the following:

- User name/Password: ubuntu/ubuntu
- User name/Password: root/dca123

When the VM starts, a start-up script checks whether the IP address of the primary interface eth0 has changed since it was last configured. If the IP address has changed, the script updates the EFA profile and configuration files appropriately and reapplies the K3s application deployment template. This operation takes a few minutes to complete. On subsequent VM reboots, if the IP address has not changed, no operation is performed by the start-up script. The logs are located under `/var/log/efa/installer`.



EFA Installation for Multi-Node Deployments

[EFA Deployment for High Availability](#) on page 23

[Install EFA in a Multi-Node Deployment](#) on page 26

[Install EFA on TPVM in a Multi-Node Deployment](#) on page 27

EFA Deployment for High Availability

Beginning with EFA 2.3.0, you can deploy EFA in a two-node cluster for high availability.

Overview

A high-availability cluster is a group of servers that provide continuous up time, or at least minimum down time, for the applications on the servers in the group. If an application on one server fails, another server in the cluster maintains the availability of the application.

In the following diagram, EFA is deployed in the TPVM running on SLX-OS. The two EFA instances are clustered and configured with one IP address, so that clients need to reach only one endpoint. All EFA services are installed on each node. The node on which EFA is installed is the active node and processes all requests. The other node is the standby.

All operations provided by EFA services must be idempotent, meaning they produce the same result for multiple identical requests or operations. For more information, see the "Idempotency" section of the [Extreme Fabric Automation Administration Guide, 2.4.0](#).

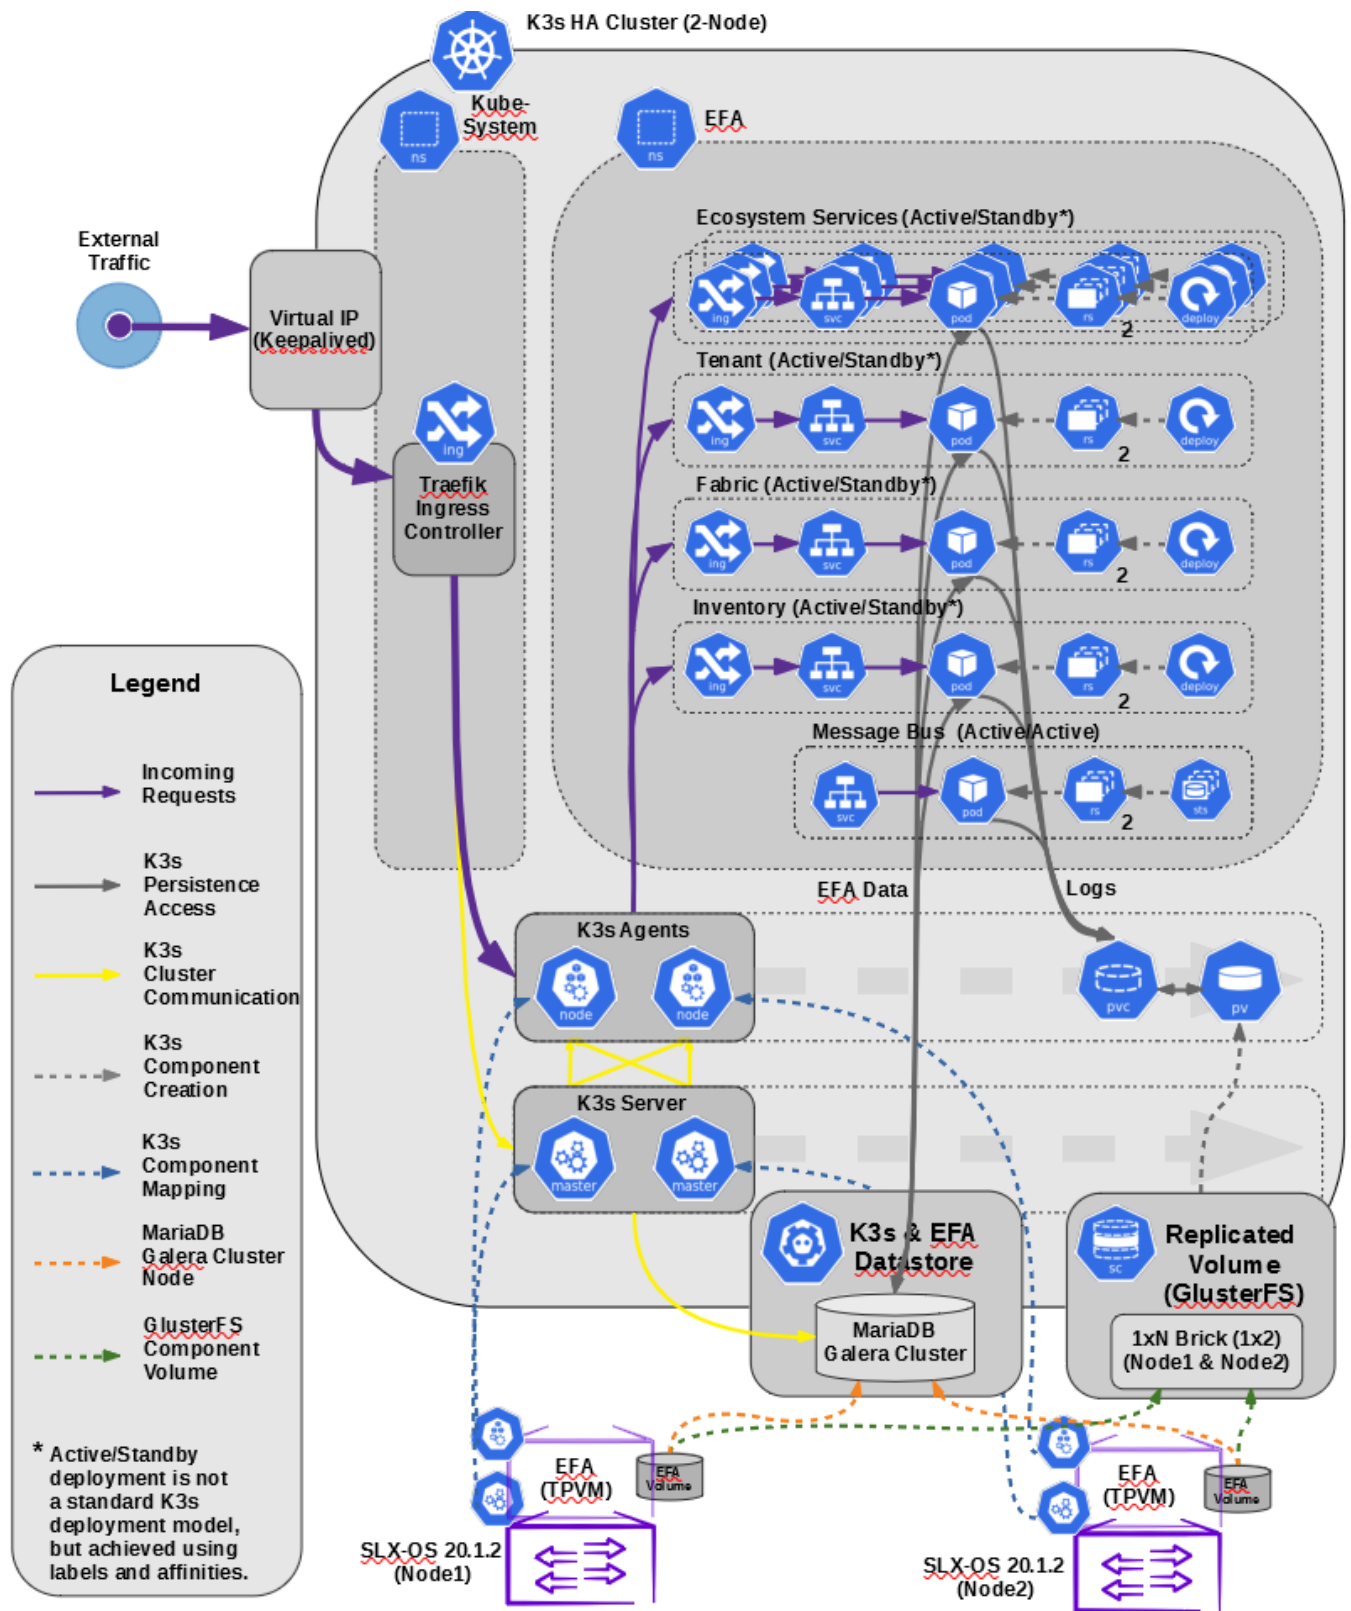


Figure 1: Two-node high-availability deployment

PV: Persistent Volume

A piece of storage in the cluster that was provisioned by an administrator.

PVC: Persistent Volume Claims

A request for storage, similar to how a Pod requests compute resources.

Brick

The basic unit of storage in GlusterFS, represented by an export directory on a server in the trusted storage pool.

SC: Storage Class

A description of the “classes” of storage in a Kubernetes realm.

SVC: Kubernetes Service

A logical set of Pods and a policy by which to access them.

ING: Kubernetes Ingress

A collection of routing rules that govern how external users access services running in a Kubernetes cluster.

RS: Kubernetes Replica Sets

Ensures how many replicas of a Pod should be running.

K3s

Manages the life cycle of the EFA services in failover or fallback scenarios.

Traefik

An embedded ingress controller (load balancer) packaged with K3s.

GlusterFS

A high-availability replicated volume that maintains the persistent storage for the K3s cluster, the EFA database, and EFA logs.

MariaDB

A database service deployed outside of the K3s cluster in active-standby mode.

RabbitMQ

A messaging service deployed in the cluster in active-active mode.

Services in high-availability mode

EFA services running on K3s are in active-active mode or active-standby mode, depending on the design of the service.

Table 13: EFA service modes

Service	Mode
Authentication	active-standby
RBAC	active-standby
Tenant	active-standby
Fabric	active-standby
Inventory	active-standby

Table 13: EFA service modes (continued)

Service	Mode
RASlog	active-active
Notification	active-standby
System	active-active
Hyper-V	active-standby
OpenStack	active-standby
vCenter	active-standby
api-docs	active-active
rabbitmq	active-active

Install EFA in a Multi-Node Deployment

Beginning with EFA 2.3.0, you can install EFA in a two-node cluster for high availability.

Before You Begin

Ensure passwordless SSH login is enabled between the two servers. For more information, see [EFA Requirements](#) on page 12.

About This Task

To install EFA, you must be a root user or have `sudoers` privileges.

The installation process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

Procedure

1. Untar the tarball on the primary server.

```
device# tar -xzf efa-vX.X.X-X.tar.gz
```

2. Change to the EFA directory.

```
device# cd efa
```

3. Run the installation script.

```
device# source deployment.sh
```

The EFA Installer begins in a series of dialogs.

4. When prompted, select **Multi-node deployment** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

5. When prompted, select **Enable secure mode** or **Standard mode** and **OK**.
6. When prompted, enter the peer IP address or FQDN of the other node in the cluster.
7. When prompted, enter the virtual IP address for the cluster.

8. When prompted to configure additional IP addresses for a health check, take one of the following steps.
 - Select **Yes** and then provide the IP addresses.
 - Select **No** to ignore this optional step.
9. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The installation proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

10. Verify the installation.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - c. From the EFA command line, run the **efa status** command for concise status information.

Install EFA on TPVM in a Multi-Node Deployment

Beginning with EFA 2.3.0, you can install EFA on a TPVM (Third-Party Virtual Machine) in a two-node deployment for high availability.

Before You Begin

The EFA tar file must be available on the `/efaboot` partition of the SLX device. You need root access to the device.

About This Task

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models](#) on page 11.

The installation process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

Installing EFA on a TPVM in a two-node deployment takes approximately 20 minutes.

Procedure

1. Verify that the TPVM is set up for an EFA deployment.

- a. Verify the versions of TPVM and SLX-OS.

For the latest supported version information, see [Supported Platforms and Deployment Models](#) on page 11.

```
device# show tpvm status
device# show version
```

- b. Verify that the TPVM has an assigned IP address.

```
device# show tpvm ip-address
```

- c. Verify that the SSH keys are uploaded.

```
device# show tpvm status
```

- d. (For SLX-OS releases earlier than 20.2.3) Verify that passwordless access is configured.

```
device# show tpvm status
```

- e. (For SLX-OS 20.2.3 and later with TPVM version 4.2.2 and later) Verify that passwordless access is configured for the peer.

```
device# show tpvm config trusted-peer
```

- f. Confirm NTP on the TPVM.

```
device# tpvm config ntp add server <ip>
```

- g. Verify that NTP is synchronized.

```
device# show tpvm config ntp
```

- h. If necessary, log in to TPVM and configure the NTP time zone.

```
device# tpvm config timezone
```

- i. If necessary, configure unique TPVM host names.

```
device# tpvm config host
```

2. Enter SLX Linux mode.

```
device# start-shell
# cd /efaboot
```

3. Copy the EFA tar file to the SLX device.

```
# scp efa-x.x.x.tar.gz
```

4. Deploy EFA on TPVM from the SLX shell.

```
device# efa deploy
Starting "efa deploy", DO NOT hit CTRL+C
Step 1: Checking if TPVM is deployed ...
Step 2: Get IP Addressed assigned to TPVM to deploy EFA
IP Address of the TPVM 10.x.x.x
Step 3: Checking for EFA packages in /efaboot directory
Step 4: Deploying EFA package efa-2.x.x.tar.gz on 10.x.x.x
```

The EFA Installer begins in a series of dialogs.

5. When prompted, select **Multi-node deployment** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

6. When prompted, select **Enable secure mode** or **Standard mode** and **OK**.
7. When prompted, enter the peer IP address or FQDN of the other node in the cluster.
8. When prompted, enter the virtual IP address for the cluster.
9. When prompted to configure additional IP addresses for a health check, take one of the following steps.
 - Select **Yes** and then provide the IP addresses.
 - Select **No** to ignore this optional step.
10. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The installation proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

11. Verify the installation.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - c. From the EFA command line, run the **efa status** command for concise status information.



EFA Upgrade

- [Supported Upgrade Paths on page 30](#)
- [Upgrade EFA in a Single-node Deployment on page 31](#)
- [Upgrade EFA on TPVM in a Single-node Deployment on page 32](#)
- [Upgrade EFA from a Single-node to a Multi-node Deployment on page 33](#)
- [Upgrade EFA in a Multi-node Deployment on page 34](#)
- [Upgrade EFA on TPVM in a Multi-node Deployment on page 35](#)
- [Upgrade EFA on an OVA on page 37](#)
- [Upgrade SLX-OS, TPVM, and EFA Together on page 37](#)
- [Upgrade Ubuntu on the EFA Host on page 38](#)
- [Recover from an Upgrade Failure on page 39](#)
- [Replace a Node in a Multi-node Deployment on page 40](#)
- [Replace a Node in a Multi-node TPVM Deployment on page 41](#)

Supported Upgrade Paths

Upgrade support varies by software version.

Table 14: Supported upgrade paths

From	To EFA 2.2.0	To EFA 2.3.x	To 2.4.x
EFA 2.2.0	N/A	Yes See the various upgrade scenarios in EFA Upgrade .	Not directly. You must upgrade to 2.3.x and then to 2.4.x.
EFA 2.3.x to 2.3.x to 2.4.x (Upgrade from any 2.3.x release to another 2.3.x release or from any 2.3.x release to 2.4.x)	N/A	Yes	Yes



Note

Upgrade from secure mode to standard mode is not a supported path.

Upgrade EFA in a Single-node Deployment

Expect the upgrade process to take approximately 8 minutes, during which EFA services are down but users or automated systems can continue to make calls into EFA.

About This Task

The upgrade process backs up the EFA system, so that you can easily recover data if the upgrade fails. For more information, see [Recover from an Upgrade Failure](#) on page 39.

Procedure

1. Download the image (*.tar.gz) to a new sub-folder.
2. Untar the image.

```
device # tar -xzf efa-v2.x.x.tar.gz
```

3. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
4. Change to the EFA directory.

```
device# cd efa
```

5. Run the deployment script.

```
device# source deployment.sh
```

6. When prompted, select **Upgrade or Redeploy**.
7. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no % or / characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The upgrade proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

8. Verify the upgrade.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - c. Run the **efa status** command for concise status information.

Upgrade EFA on TPVM in a Single-node Deployment

You can upgrade EFA on TPVM (Third-Party Virtual Machine) from the SLX device.

Before You Begin

The EFA tar file must be available on the `/efaboot` partition of the SLX device. You need root access to the device. If more than one EFA version is available in the `/efaboot` directory, you are prompted to select a version during upgrade.

About This Task

The upgrade process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

The upgrade process backs up the EFA system, so that you can easily recover data if the upgrade fails. For more information, see [Recover from an Upgrade Failure](#) on page 39.

Procedure

1. Enter SLX Linux mode.

```
device# start-shell
# cd /efaboot
```

2. Copy the EFA tar file to the SLX device.

```
# scp efa-x.x.x.tar.gz
```

3. Deploy EFA on the TPVM from the SLX device.

```
device# efa deploy
```

The EFA Installer begins in a series of dialogs.

4. When prompted, select **Single-node deployment** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

5. When prompted, select **Enable secure mode** or **Standard mode** and **OK**.
6. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The upgrade proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

7. Verify the upgrade.
 - a. On the SLX device, run the `show efa status` command to see details of the installation and the state of services.
 - b. From the EFA command line, run the `efactl status` command to see the status of nodes, pods, and services.
 - c. From the EFA command line, run the `efa status` command for concise status information.

Upgrade EFA from a Single-node to a Multi-node Deployment

You can upgrade a single-node deployment of EFA to a multi-node deployment.

Before You Begin

Ensure the single node is running EFA 2.3.0 or later. Upgrade if necessary. For more information, see [Upgrade EFA in a Single-node Deployment](#) on page 31.

Ensure you have completed the high-availability prerequisites in [EFA Requirements](#) on page 12.

About This Task

Expect the upgrade process to take approximately 8 minutes, during which EFA services are down but users or automated systems can continue to make calls into EFA.

The upgrade process backs up the EFA system, so that you can easily recover data if the upgrade fails. For more information, see [Recover from an Upgrade Failure](#) on page 39.

The upgrade process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

Procedure

1. Download the image (*.tar.gz) to a new sub-folder.
2. Untar the image.

```
device# tar -xzf efa-v2.x.x.tar.gz
```

3. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
4. Change to the EFA directory.

```
device# cd efa
```

5. Run the deployment script.

```
device# source deployment.sh
```

The EFA Installer begins in a series of dialogs.

6. When prompted, select **Multi-node deployment** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

7. When prompted, select **Enable secure mode** or **Standard mode** and **OK**.
8. When prompted, enter the peer IP address or FQDN of the other node in the cluster.
9. When prompted, enter the virtual IP address for the cluster.

10. When prompted to configure additional IP addresses for a health check, take one of the following steps.
 - Select **Yes** and then provide the IP addresses.
 - Select **No** to ignore this optional step.
11. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The upgrade proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

12. Verify the upgrade.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - c. From the EFA command line, run the **efa status** command for concise status information.

Upgrade EFA in a Multi-node Deployment

With EFA 2.3.0 and later, you can upgrade EFA in a multi-node, high-availability deployment.

About This Task

Expect the upgrade process to take approximately 10 minutes, during which EFA services are down but users or automated systems can continue to make calls into EFA.

The upgrade process also backs up the EFA database, so that you can easily recover data if the upgrade fails.

Procedure

1. Download the image (*.tar.gz) to a new sub-folder.
2. Untar the image.

```
device# tar -xzf efa-v2.x.x.tar.gz
```

3. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
4. Change to the EFA directory.

```
device# cd efa
```

5. Run the deployment script.

```
device# source deployment.sh
```

The EFA Installer begins in a series of dialogs.

6. When prompted, select **Upgrade or Redeploy**.
7. When prompted to configure additional IP addresses for a health check, take one of the following steps.
 - Select **Yes** and then provide the IP addresses.
 - Select **No** to ignore this optional step.
8. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The upgrade proceeds. Messages describe the progress and indicate when EFA is deployed.

9. Verify the upgrade.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
 - c. From the EFA command line, run the **efa status** command for concise status information.

Upgrade EFA on TPVM in a Multi-node Deployment

With EFA 2.3.0 and later, you can upgrade a multi-node deployment of EFA on TPVM (Third-Party Virtual Machine).

Before You Begin

The EFA tar file must be available on the `/efaboot` partition of the SLX device. You need root access to the device.

About This Task

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models](#) on page 11.

The upgrade process prompts you to select *secure mode* or *standard mode*. For more information about these modes, see [EFA Installation Modes](#) on page 16.

Installing EFA on a TPVM in a two-node deployment takes approximately 20 minutes.

Procedure

1. Enter SLX Linux mode.

```
device# start-shell
# cd /efaboot
```

2. Copy the EFA tar file to the SLX device.

```
# scp efa-x.x.x.tar.gz
```

3. Deploy EFA on TPVM from the SLX shell.

```
device# efa deploy
Starting "efa deploy", DO NOT hit CTRL+C
Step 1: Checking if TPVM is deployed ...
Step 2: Get IP Addressed assigned to TPVM to deploy EFA
IP Address of the TPVM 10.x.x.x
Step 3: Checking for EFA packages in /efaboot directory
Step 4: Deploying EFA package efa-2.x.x.tar.gz on 10.x.x.x
```

The EFA Installer begins in a series of dialogs.

4. When prompted, select **Multi-node deployment** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

5. When prompted, select **Enable secure mode** or **Standard mode** and **OK**.
6. When prompted, enter the peer IP address or FQDN of the other node in the cluster.
7. When prompted, enter the virtual IP address for the cluster.
8. When prompted to configure additional IP addresses for a health check, take one of the following steps.
 - Select **Yes** and then provide the IP addresses.
 - Select **No** to ignore this optional step.
9. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no **%** or **/** characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.

The upgrade proceeds. Messages summarize your selections, describe the progress, and indicate when EFA is deployed.

10. Verify the upgrade.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.

- c. From the EFA command line, run the **efa status** command for concise status information.

Upgrade EFA on an OVA

Open Virtual Appliance (OVA) is an OVF file packaged with a base image Ubuntu image and installed with EFA.

About This Task

See [Deploy the OVA for EFA](#) on page 21 for a list of prerequisites.

Procedure

1. Log in to the OVA.
2. Switch to the super-user.

```
# sudo su -
```

3. Copy the new tar file to `/opt/godcapp/`.
4. Extract the tar.

```
device# tar -xvf efa-2.x.x.tar.gz
device# cd efa
```

5. Run the deployment script.

```
device# source deployment.sh
```

6. When prompted, select **Upgrade or Redeploy**.

The upgrade proceeds. Messages describe the progress and indicate when the upgrade is complete.

7. When the upgrade is complete, update your shell's environment.

```
device# source /etc/profile
```

Upgrade SLX-OS, TPVM, and EFA Together

Take the following steps to upgrade an SLX device to the latest supported version of SLX-OS with the latest supported version of TPVM.

Procedure

1. If your system is running a version of EFA earlier than 2.3.x, upgrade to EFA 2.3.x.
For more information, see [Supported Upgrade Paths](#) on page 30.
2. Back up EFA.
For more information, see "Back up and Restore the EFA System" in the [Extreme Fabric Automation Administration Guide, 2.4.0](#).
3. Copy the backup file to a remote location, such as the `/efaboot` of the SLX device or SCP on TPVM.



Note

The TPVM backup process backs up only the database and not the application.

4. Upgrade the firmware on the SLX device.

For more information, see "Device Image Management" in the [Extreme Fabric Automation Administration Guide, 2.4.0](#).

5. In a multi-node deployment, perform the following steps on all nodes in the cluster except the node on which you will deploy EFA.
 - a. Stop the TPVM application.


```
device# tpvm stop
```
 - b. Uninstall the TPVM application.


```
device# tpvm uninstall
```
 - c. From the device, remove the old TPVM image from the `/tftpboot/SWBD2900/*` directory. This step requires administrator access.
 - d. Copy the new TPVM build to the `/tftpboot/SWBD2900/*` directory.
 - e. Deploy TPVM with DHCP or static IP addresses.
 - f. Configure NTP on TPVM.


```
device# tpvm config ntp add server <ip-addr>
```
6. Ensure that the EFA build is available in `/efaboot` on the SLX device.
7. Repeat step 5 on the node on which you will deploy EFA (the active node in the cluster).
8. Deploy EFA.

For more information, see [Install EFA on TPVM in a Multi-Node Deployment](#) on page 27 or [Install EFA on TPVM in a Single-Node Deployment](#) on page 19.
9. Verify the deployment.
 - a. On the SLX device, run the `show efa status` command to see details of the installation and the state of services.
 - b. From the EFA command line, run the `efactl status` command to see the status of nodes, pods, and services.
 - c. From the EFA command line, run the `efa status` command for concise status information.

Upgrade Ubuntu on the EFA Host

You can upgrade Ubuntu in single-node and multi-node deployments.

Before You Begin

Ensure that EFA is at release 2.3.0 or later. Upgrade if necessary. For more information, see [EFA Upgrade](#) on page 30.

Ensure that the nodes you want to upgrade are healthy and that EFA services are operating.

About This Task

EFA is supported on Ubuntu 16.04 and 18.04 as described in [Supported Platforms and Deployment Models](#) on page 11. You can upgrade from 16.04 to 18.04 while EFA is installed.



Note

- This process is not supported for deployments of EFA on TPVM.
- This process assumes that the node you are upgrading is connected to the internet. The Ubuntu [Release Notes](#) indicate that there is no offline upgrade option.

Procedure

1. Update the Ubuntu package database.

```
# sudo apt-get update
```

2. Upgrade all Ubuntu packages.

```
# sudo apt-get upgrade
```

3. To upgrade Ubuntu on a single node, take the following steps.

- a. Upgrade the node.

```
# sudo do-release-upgrade
```

EFA is not operational while the upgrade is in progress.

- b. Reboot the system.
- c. Verify that EFA is operational.

4. To upgrade Ubuntu in a two-node cluster, take the following steps.

- a. Upgrade one node in the cluster.

```
# sudo do-release-upgrade
```

If you run the upgrade on the active node, then failover to the standby node occurs. EFA is not operational during the failover.

- b. Upgrade the second node.
- c. Verify that the nodes are at the new version.
- d. Verify that EFA is operational.

Recover from an Upgrade Failure

You can recover from an upgrade failure by rerunning the upgrade or by performing a fresh installation and then restoring the system with the backup taken at the beginning of the upgrade process.

Procedure

1. To rerun the upgrade, follow the steps for the type of upgrade you were attempting.
 - [Upgrade EFA from a Single-node to a Multi-node Deployment](#) on page 33
 - [Upgrade EFA on TPVM in a Single-node Deployment](#) on page 32
 - [Upgrade EFA in a Single-node Deployment](#) on page 31
 - [Upgrade EFA in a Multi-node Deployment](#) on page 34
 - [Upgrade EFA on TPVM in a Multi-node Deployment](#) on page 35
2. To perform a fresh installation and restore the system backup, take the following steps.
 - a. Uninstall EFA to remove any components that might have been installed before the upgrade failed.
 - [Uninstall EFA on TPVM in a Multi-Node Deployment](#) on page 43
 - [Uninstall EFA in a Multi-node Deployment](#) on page 42
 - [Uninstall EFA on TPVM in a Single-Node Deployment](#) on page 42
 - [Uninstall EFA in a Single-Node Deployment](#) on page 42

- b. Follow the steps for the type of installation you need.
 - [Install EFA on TPVM in a Multi-Node Deployment](#) on page 27
 - [Install EFA in a Multi-Node Deployment](#) on page 26
 - [Install EFA on TPVM in a Single-Node Deployment](#) on page 19
 - [Install EFA in a Single-Node Deployment](#) on page 18
- c. Restore the backup.

```
device# efa system restore --backup-tar <filename>.tar.gz
```

For more information about backup tar files, see the "EFA System Backup and Restoration" section of the [Extreme Fabric Automation Administration Guide, 2.4.0](#).

Replace a Node in a Multi-node Deployment

You can use the upgrade process to replace a faulty node in a multi-node deployment.

Before You Begin

Ensure the cluster with the faulty node is running EFA 2.3.0 or later.

Ensure you have completed the high-availability prerequisites in [EFA Requirements](#) on page 12.

Ensure that EFA is not deployed on the replacement node.

About This Task

During this process, the faulty node is decommissioned, the replacement node is provisioned, and the active node is reconfigured to form the cluster.

Perform this procedure on the node where EFA is installed.

Procedure

1. Navigate to the directory where the EFA file (*.tar.gz) is untarred.
2. Run the deployment script.

```
device# source deployment.sh
```

The EFA Installer begins in a series of dialogs.

3. When prompted, select **Multi Node Build Upgrade with Node Replacement** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

4. When prompted, enter the IP address or host name of the replacement peer node.
5. Select **OK**.

The node replacement proceeds. Messages indicate the progress and when the replacement is complete.

6. Verify the status of EFA after the node replacement.

```
device# efactl status
```


Replace a Node in a Multi-node TPVM Deployment

You can use the upgrade process to replace a faulty node in a multi-node TPVM deployment.

Before You Begin

Ensure the cluster with the faulty node is running EFA 2.3.0 or later.

Ensure you have completed the high-availability prerequisites in [EFA Requirements](#) on page 12.

Ensure that EFA is not deployed on the replacement node.

About This Task

During this process, the faulty node is decommissioned, the replacement node is provisioned, and the active node is reconfigured to form the cluster.

Perform this procedure on the node where EFA is installed.

Procedure

1. Enter SLX Linux mode.

```
device# start-shell
# cd /efaboot
```

2. Copy the EFA tar file to the SLX device.

```
# scp efa-x.x.x.tar.gz
```

3. Deploy EFA on TPVM from the SLX shell.

```
device# efa deploy
Starting "efa deploy", DO NOT hit CTRL+C
Step 1: Checking if TPVM is deployed ...
Step 2: Get IP Addressed assigned to TPVM to deploy EFA
IP Address of the TPVM 10.x.x.x
Step 3: Checking for EFA packages in /efaboot directory
Step 4: Deploying EFA package efa-2.x.x.tar.gz on 10.x.x.x
```

The EFA Installer begins in a series of dialogs.

4. When prompted, select **Multi Node Build Upgrade with Node Replacement** and **OK**.



Tip

Use arrow keys to move between options and the space bar to select an option.

5. When prompted, enter the IP address or host name of the replacement peer node and select **OK**.

The node replacement proceeds. Messages indicate the progress and when the replacement is complete.

6. Verify the status of EFA after the node replacement.

```
device# efactl status
```



EFA Uninstallation

[Uninstall EFA in a Single-Node Deployment on page 42](#)

[Uninstall EFA on TPVM in a Single-Node Deployment on page 42](#)

[Uninstall EFA in a Multi-node Deployment on page 42](#)

[Uninstall EFA on TPVM in a Multi-Node Deployment on page 43](#)

Uninstall EFA in a Single-Node Deployment

When EFA is uninstalled, EFA services are stopped and the database and directories are removed.

Procedure

1. On the node where EFA is installed, run the deployment script.

```
# source deployment.sh
```

2. When prompted, select **Remove the current EFA Stack**.

The uninstall process proceeds. A message indicates when the EFA stack is uninstalled.

Uninstall EFA on TPVM in a Single-Node Deployment

When EFA is uninstalled, EFA services are stopped and the database and directories are removed.

Procedure

1. From the SLX device console, uninstall EFA.

```
device# no efa deploy
```

2. When prompted to continue, enter *y*.

The uninstall process proceeds. A message indicates when the EFA stack is uninstalled.

Uninstall EFA in a Multi-node Deployment

When EFA is uninstalled, EFA services are stopped and the database and directories are removed.

Procedure

1. On the node where EFA is installed, run the deployment script.

```
# source deployment.sh
```

2. When prompted, select **Remove the current EFA Stack**.

The uninstall process proceeds. A message indicates when the EFA stack is uninstalled.

Uninstall EFA on TPVM in a Multi-Node Deployment

When EFA is uninstalled, EFA services are stopped and the database and directories are removed.

About This Task

Perform this task from the SLX device console. If both nodes are up, run the commands on the node on which TPVM is running. As a best practice, run the commands on the active node. Otherwise, run the commands on both nodes.

Procedure

1. Stop and then start the TPVM to ensure there are no DNS resolution issues.

```
device# tpvm stop
device# tpvm start
```

2. Uninstall EFA.

```
device# no efa deploy
```

3. When prompted to continue, enter `y`.

The uninstall process proceeds. A message indicates when the EFA stack is uninstalled.



Redundant Management Network

[Redundant Management Network Overview](#) on page 44

[Enable Redundant Management](#) on page 45

[Redundant Management Data Path](#) on page 47

Redundant Management Network Overview

Redundant Management Network provides fault tolerance for the management path. This is done using Linux bonding by pairing the physical management port of the chassis with any one of the physical front panel user ports.




Figure 2: Redundant Management Network Overview

Linux Bonding

The `redundant-management enable` command can be used to pair one of the front panel ports with the conventional `Mgmt 0` port to form a Linux Bonding interface, `bond0` at SLX Linux OS.

- The Linux bond will be in Active/Standby mode. The Physical Management port is the primary and active port. The configured front panel port will be in Standby mode.
 - `mode 1` supported by Linux Bonding with `Mgmt 0` (`eth0`) is the primary port.
 - The front panel port allows traffic through it only if `Mgmt 0` is down. `Mgmt 0` takes over Active port as soon as it recovers.

- If the active primary Mgmt 0 path experiences failure, SLX OS and TPVM OS can be reached through Standby path.

Supported Ports

Any SLX front panel port can be used at native speed and property for Linux Bonding.



Note

- SLX 9640 and SLX 9150 - Preferred ports are 10G/1G port in 1G mode.
- SLX 9640 - Avoid Insight port 0/24.
- SLX 9250 - Breakout mode 4x1G ports are available to allow the Mellanox adapter with 1G transceiver. Because the adapter occupies the whole cage, only the first member port (:1) can be used as redundant management interface.

No Redundancy Period

Redundancy is not supported if the device is reloaded or in ZTP mode.

- After reloading a device, use the **redundant-management enable** command or startup config replay to enable Linux Bonding or redundancy.
- Upon factory arrival, across first power cycle, or due to `write erase` CLI, ZTP mode is set in with factory default configuration.
- Breakout mode 1G ports are not supported in the factory default configuration.

Standby Port Rate Throughput

Since internal path for Standby traffic is Control Plane traffic on PCIe Channel between ASIC and CPU, its function of internal CPU load is totally unrelated and independent of front panel physical port limit and capability.

Enable Redundant Management

Redundant management provides fault tolerance for the management path.

About This Task

Perform this procedure on a supported SLX device. For more information, see [Redundant Management Network Overview](#) on page 44.

Procedure

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface configuration mode.

```
device(config)# interface ethernet 0/32
```

3. Enable Redundant Management.

```
device(conf-if-eth-0/32)# redundant-management enable
```

Example

This example configures Ethernet 0/32 at 10G speed.

```
device# config
device(config)# interface ethernet 0/32
device(conf-if-eth-0/32)# redundant-management enable
device(conf-if-eth-0/32)# no shut
```

This example configures Ethernet 0/32 at 1G speed.

```
device# config
device(config)# interface ethernet 0/32
device(conf-if-eth-0/32)# speed 1000
device(conf-if-eth-0/32)# redundant-management enable
device(conf-if-eth-0/32)# no shut
```

This examples configures Ethernet 0/32 on an SLX 9250 with a Mellanox adapter at 1G speed.

```
device# conf t
device(config)# hardware
device(config-hardware)# connector 0/32
device(config-connector-0/32)# breakout mode 4x1G
device(config-connector-0/32)# end
device# conf t
device(config)# interface ethernet 0/32:1
device(conf-if-eth-0/32:1)# redundant-management enable
device(conf-if-eth-0/32:1)# no shut
```

Example

These examples show interface details when redundant management is enabled.

```
device# show interface management 0

interface Management 0
line-speed actual "1000baseT, Duplex: Full"
oper-status up
ip address "static 10.x.x.x/22"
ip gateway-address 10.x.x.x
ipv6 ipv6-address [ ]
ipv6 ipv6-gateways [ ]
redundant management port 0/32

device# show ip interface brief

Flags: I - Insight Enabled U - Unnumbered interface M - Redundant management port
Interface          IP-Address      Vrf              Status            Protocol
=====          =====
Ethernet 0/1       unassigned      default-vrf      administratively  down
Ethernet 0/2       unassigned      default-vrf      administratively  down
...
Ethernet 0/32 (M) unassigned      mgmt-vrf         administratively  down
...

device# show interface ethernet 0/32

Ethernet 0/32 is admin down, line protocol is down (admin down)
Redundant management mode is enabled
Hardware is Ethernet, address is 609c.9f5a.a35f
Current address is 609c.9f5a.a35f
Pluggable media not present
Description: Insight port
Interface index (ifindex) is 202350592 (0xc0fa000)
MTU 9216 bytes
Maximum Speed : 10G
```

```
LineSpeed Actual : Nil
LineSpeed Configured : Auto, Duplex: Full
Priority Tag disable
Forward LACP PDU: Disable
Route Only: Disabled
Tag-type: 0x8100
Last clearing of show interface counters: 00:01:13
Queueing strategy: fifo
FEC Mode - Disabled
Receive Statistics:
0 packets, 0 bytes
Unicasts: 0, Multicasts: 0, Broadcasts: 0
64-byte pkts: 0, Over 64-byte pkts: 0, Over 127-byte pkts: 0
Over 255-byte pkts: 0, Over 511-byte pkts: 0, Over 1023-byte pkts: 0
Over 1518-byte pkts(Jumbo): 0
Runts: 0, Jabbers: 0, CRC: 0, Overruns: 0
Errors: 0, Discards: 0
Transmit Statistics:
0 packets, 0 bytes
Unicasts: 0, Multicasts: 0, Broadcasts: 0
Underruns: 0
Errors: 0, Discards: 0
Rate info:
Input 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Output 0.000000 Mbits/sec, 0 packets/sec, 0.00% of line-rate
Route-Only Packets Dropped: 0
Time since last interface status change: 00:01:13
```

Redundant Management Data Path

SLX Linux, boots with bond0 interface with Primary Active eth3 (Physical Management 0 Interface). Interface bond0 is subordinate to vBridge (eth0), which serves as Management 0 interface to SLX Linux and all applications on it. This eth0 is connected through Linux Tap to TPVM eth0. TPVM eth0 contains a separate MAC. The IPv4 address is assigned to eth0 through DHCP or static.

At SLX Linux, the logical proxy interface Eth0.15 or Eth0.32.1 is created to represent the front panel port as a Standby member for bond0.




Figure 3: Data path overview