

# Extreme Fabric Automation Deployment Guide

3.1.0

9037627-00 Rev AB  
January 2024



Copyright © 2024 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.

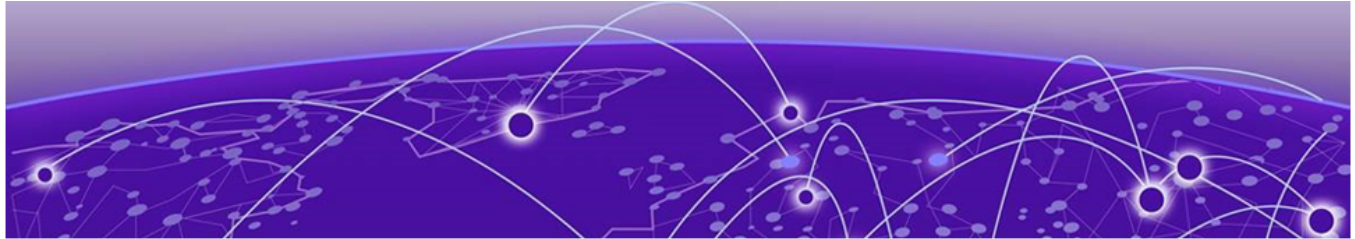
All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

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

---

<b>Preface.....</b>	<b>6</b>
Text Conventions.....	6
Documentation and Training.....	7
Help and Support.....	8
Subscribe to Product Announcements.....	8
Send Feedback.....	9
<b>About this Document.....</b>	<b>10</b>
What's New in this Document.....	10
<b>EFA Deployment Preparation.....</b>	<b>11</b>
Supported Platforms and Deployment Models for Fabric Manager.....	12
Supported Platforms and Deployment Models for Visibility Manager.....	15
EFA Requirements.....	16
General requirements.....	16
High-availability requirements.....	17
EFA Port Requirements.....	19
EFA Installation on TPVM.....	20
Overview.....	20
Requirements.....	21
Supported deployments.....	21
<b>EFA Installation for Single-Node Deployments.....</b>	<b>22</b>
EFA Installer Improvements for Server-Based Deployment.....	22
Install EFA in a Single-Node Deployment.....	23
Install EFA on TPVM in a Single-Node Deployment .....	24
Deploy the OVA.....	26
Configure OVA using Postboot Menu.....	27
<b>EFA Installation for Multi-Node Deployments.....</b>	<b>30</b>
EFA Deployment for High Availability .....	30
Overview.....	30
Install EFA in a Multi-Node Deployment.....	33
Install EFA on TPVM in a Multi-Node Deployment .....	35
<b>EFA Upgrade.....</b>	<b>38</b>
Upgrade EFA in a Single-node Deployment.....	38
Upgrade EFA on TPVM in a Single-node Deployment.....	39
Upgrade EFA from a Single-Node to a Multi-Node Deployment.....	41
Upgrade EFA in a Multi-Node Deployment.....	43
Upgrade EFA on TPVM in a Multi-Node Deployment.....	44
Upgrade EFA on an OVA.....	46
Upgrading SLX-OS, TPVM, and EFA Together.....	47
Requirement for SCP connections.....	47

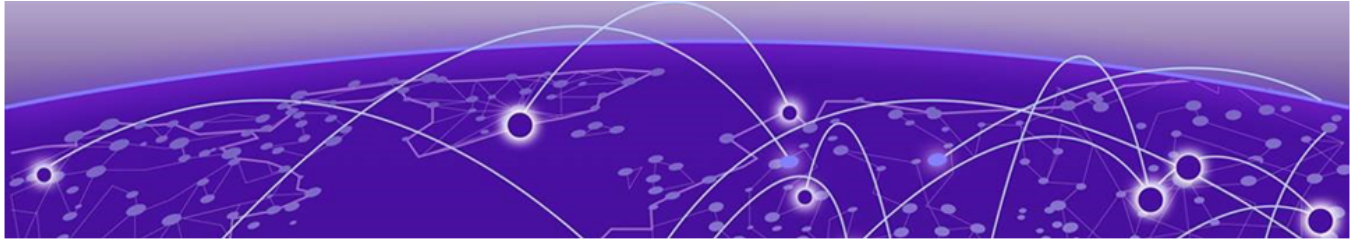
Upgrade EFA, SLX-OS, and TPVM Method 1.....	48
Upgrade EFA, SLX-OS, and TPVM Method 2.....	49
Recover from an Upgrade Failure.....	52
Rollback.....	53
Maintain TPVM Versions After a Rollback in a Multi-Node Deployment.....	53
Rollback SLX.....	54
Rollback EFA.....	54
<b>Node Replacement.....</b>	<b>56</b>
Replace a Node in a Multi-node Deployment.....	56
Replace a Node in a Multi-node TPVM Deployment.....	57
<b>EFA Uninstallation.....</b>	<b>59</b>
Uninstall EFA in a Single-Node or Multi-Node Deployment.....	59
Uninstall EFA on TPVM in a Single-Node and Multi-Node Deployment.....	59
<b>TPVM Upgrade from EFA.....</b>	<b>60</b>
TPVM Complete Package Upgrade.....	60
Assumptions and Limitations.....	60
TPVM Upgrade Workflow Dependencies.....	61
TPVM Upgrade Workflow.....	63
TPVM Upgrade Workflow States.....	65
TPVM Incremental Upgrade.....	66
<b>Upgrade Ubuntu.....</b>	<b>71</b>
Upgrade Ubuntu on the EFA Host - Single Node.....	71
Upgrade Ubuntu on the EFA Host - Multi-Node.....	72
<b>Redundant Management Network.....</b>	<b>74</b>
Redundant Management Network Overview.....	74
Linux Bonding.....	74
Supported Ports.....	75
No Redundancy Period.....	75
Standby Port Rate Throughput.....	75
Enable Redundant Management.....	75
Redundant Management Data Path.....	77
<b>Flexible EFA Deployment for TPVM.....</b>	<b>79</b>
Flexible EFA Deployment Overview.....	79
SLX CLI.....	79
EFA Deployment.....	80
Listing EFA Packages.....	80
Input Parameters on Single-Node Install or Upgrade.....	80
Minimum Required Commands.....	80
Management IP Networks.....	80
Input Parameters on Multi-Node Install or Upgrade.....	81
Deployment Type.....	81
Virtual IPv6.....	81
Management IP Networks.....	81
Build Upgrade and Replacement.....	82
Single CLI for HA Ping-target Parameter.....	82
EFA Installer Improvements for TPVM-Based Deployment.....	83
Upgrades and Service State.....	83

Enable or Disable Services.....84

No Graphics Mode.....84

EFA Deployment with Rollback.....85

Rollback the EFA Upgrade.....85



# Preface

---

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.






## 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> .
<b>Key</b> names	Key names are written in boldface, for example <b>Ctrl</b> or <b>Esc</b> . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press <b>Ctrl+Alt+Del</b>
<i>Words in 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.
<b>NEW!</b>	New information. In a PDF, this is searchable text.

**Table 3: Command syntax**

Convention	Description
<b>bold</b> text	Bold text indicates command names, keywords, and command options.
<i>italic</i> text	Italic text indicates variable content.
[ ]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ <b>x</b>   <b>y</b>   <b>z</b> }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
<b>x</b>   <b>y</b>	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[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/](http://www.extremenetworks.com/education/).

---

## Help and Support

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](http://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.

## Send Feedback

---

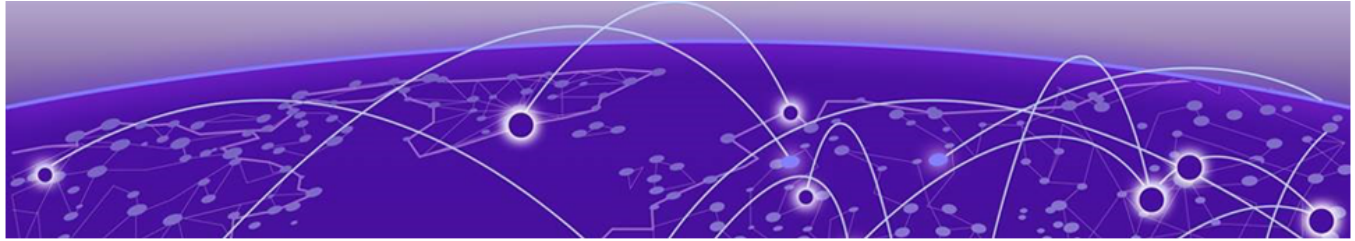
The Information Development team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, do either of the following:

- Access the feedback form at <https://www.extremenetworks.com/documentation-feedback/>.
- Email us at [documentation@extremenetworks.com](mailto: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 10

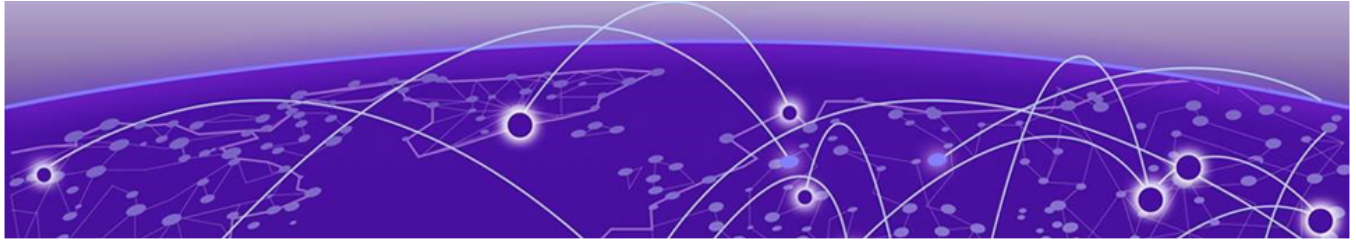
## What's New in this Document

---

The following topics are updated for the ExtremeCloud Orchestrator 3.1.0 software release.

**Table 4: Summary of changes**

Description	Topic
New topic describes EFA installer support for a packet and fabric suites.	<a href="#">EFA Installer Improvements for Server-Based Deployment</a> on page 22
New topic describes EFA installer support for TPVM deployment	<a href="#">EFA Installer Improvements for TPVM-Based Deployment</a> on page 83
New topic describes single CLI commands supported on SLX for HA deployment.	<a href="#">Single CLI for HA Ping-target Parameter</a> on page 82
Updated topics for Northbound IPv6 support in EFA.	<ul style="list-style-type: none"><li>• <a href="#">Install EFA in a Single-Node Deployment</a> on page 23</li><li>• <a href="#">Install EFA on TPVM in a Single-Node Deployment</a> on page 24</li><li>• <a href="#">Configure OVA using Postboot Menu</a> on page 27</li><li>• <a href="#">Install EFA in a Multi-Node Deployment</a> on page 33</li><li>• <a href="#">Upgrade EFA in a Single-node Deployment</a> on page 38</li><li>• <a href="#">Upgrade EFA from a Single-Node to a Multi-Node Deployment</a> on page 41</li><li>• <a href="#">Single CLI for HA Ping-target Parameter</a> on page 82</li><li>• <a href="#">Upgrade EFA on TPVM in a Multi-Node Deployment</a> on page 44</li></ul>



# EFA Deployment Preparation

---

[Supported Platforms and Deployment Models for Fabric Manager](#) on page 12

[Supported Platforms and Deployment Models for Visibility Manager](#) on page 15

[EFA Requirements](#) on page 16

[EFA Port Requirements](#) on page 19

[EFA Installation on TPVM](#) on page 20

## Supported Platforms and Deployment Models for Fabric Manager

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**

EFA Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Virtual Machine
2.7.x, 3.0.0, and 3.1.0	External server (bare metal)	More than 24	Yes	16.04, 18.04, and 20.04	<ul style="list-style-type: none"><li>• CPU: 4 cores</li><li>• Storage: 64 GB</li><li>• RAM: 8 GB</li></ul>

**Table 6: OVA Deployment Models**

EFA Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Virtual Machine
2.7.x, 3.0.0, and 3.1.0	External server (OVA)	More than 24	Yes	18.04	<ul style="list-style-type: none"><li>• CPU: 4 cores</li><li>• Storage: 64 GB</li><li>• RAM: 8 GB</li></ul>

**Table 7: TPVM Deployment Models**

EFA Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
2.7.x	<ul style="list-style-type: none"><li>• SLX 9150</li><li>• SLX 9250</li><li>• SLX 9740</li><li>• Extreme 8520</li><li>• Extreme 8720</li></ul>	Up to 24	Yes	18.04	20.4.1
3.0.x	<ul style="list-style-type: none"><li>• SLX 9150</li><li>• SLX 9250</li><li>• SLX 9740</li><li>• Extreme 8520</li><li>• Extreme 8720</li></ul>	Up to 24	Yes	18.04	20.4.2
3.1.x	<ul style="list-style-type: none"><li>• SLX 9150</li><li>• SLX 9250</li><li>• SLX 9740</li></ul>	Up to 24	Yes	18.04	20.4.2

**Table 7: TPVM Deployment Models (continued)**

EFA Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
	<ul style="list-style-type: none"> <li>Extreme 8520</li> <li>Extreme 8720</li> </ul>				

**Table 8: TPVM Software Support**

TPVM Version	SLX-OS 20.2.3 d/e/f	SLX-OS 20.3.2	SLX-OS 20.3.2 a	SLX-OS 20.3.2 b	SLX-OS 20.3.2 c	SLX-OS 20.3.2 d	SLX-OS 20.3.4/4a	SLX-OS 20.4.1	SLX-OS 20.4.1 b	SLX-OS 20.4.2/2a	Ubuntu Version	EFA Version
4.2.4	Yes	No	No	No	No	No	No	No	No	No	18.04	2.4.x
4.2.5	No	Yes	Yes	No	No	No	No	No	No	No	18.04	2.4.x, 2.5.0
4.2.5	No	No	No	Yes	No	No	No	No	No	No	18.04	2.5.1, 2.5.2
4.2.5	No	No	No	No	Yes	No	No	No	No	No	18.04	2.5.3
4.3.0	No	No	No	No	No	Yes	No	No	No	No	18.04	2.5.4, 2.5.5
4.4.0	No	No	No	No	No	No	Yes	No	No	No	18.04	2.6.0, 2.6.1
4.5.0	No	No	No	No	No	No	No	Yes	No	No	18.04	2.7.0
4.5.1	No	No	No	No	No	No	No	No	Yes	No	18.04	2.7.2
4.5.3	No	No	No	No	No	No	No	No	No	Yes (only with 20.4.2)	18.04	3.0.0
4.5.6	No	No	No	No	No	No	No	No	No	Yes	20.04	3.1.0

**Note**

The seamless TPVM upgrade feature is not available in SLX 20.2.3f.

**Table 9: IP Fabric Topology Matrix**

Device	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9150	20.2.x, 20.3.x, 20.4.x	✓				✓
SLX 9250	20.2.x, 20.3.x, 20.4.x	✓	✓	✓		✓
SLX 9540	20.2.x, 20.3.x, 20.4.x	✓			✓	

**Table 9: IP Fabric Topology Matrix (continued)**

Device	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9640	20.2.x, 20.3.x, 20.4.x				✓	
SLX 9740	20.2.x, 20.3.x, 20.4.x		✓	✓	✓	✓
Extreme 8720	20.3.x, 20.4.x	✓	✓	✓	✓	✓
Extreme 8520	20.3.x, 20.4.x	✓			✓	✓

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

EFA Version	Neutron Version	SLX-OS Version
2.5.4, 2.5.5	3.11-04	20.3.2d

## Supported Platforms and Deployment Models for Visibility Manager

Support includes bare metal, OVA, and supported devices and software.

**Table 11: Bare Metal Deployment Models**

EFA Version	Deployment	Ubuntu Version	Virtual Machine
3.1.0	External server (bare metal)	18.04 and 20.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul> Recommended: <ul style="list-style-type: none"> <li>• CPU: 16 cores</li> <li>• Storage: 200 GB</li> <li>• RAM: 32 GB</li> </ul>

**Table 12: OVA Deployment Models**

EFA Version	Deployment	Ubuntu Version	Virtual Machine
3.1.0	External server (OVA)	18.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 13: Supported Devices and Software**

Device	Supported Software
Extreme 9920	Extreme 9920 software with the NPB application <ul style="list-style-type: none"> <li>• 21.1.2.x</li> </ul>
Extreme Routing MLX Series	<ul style="list-style-type: none"> <li>• NetIron 6.3.00 patches</li> </ul>
Extreme Switching SLX 9140	<ul style="list-style-type: none"> <li>• SLX-OS 18s.1.03 patches</li> </ul>
Extreme Switching SLX 9240	<ul style="list-style-type: none"> <li>• SLX-OS 18s.1.03 patches</li> </ul>



### Note

- Upgrade from XVM (Extreme Visibility Manager) to EFA is not supported.
- In the Visibility Manager deployment of EFA, non-NPB devices (SLX) are not supported. There will not be any notifications provided to the users for the discovery failure.
- EFA supports only fixed set of special characters for names. Any additional characters configured in MLX or SLX will be reconciled in EFA and can be edited or deleted. Any configuration name must start with alphabet and can contain " a-z A-Z 0-9 \_ -"

## 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 numeric characters and lowercase alphabetic characters. Do not use uppercase alphabetic 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 the same 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
```



#### Note

If you are not using the provided EFA OVA or TPVM, consult with your system administrator for configuring NTP.

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.



#### Note

If you are not using the provided EFA OVA or TPVM, consult with your system administrator for configuring NTP.

- **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 for Fabric Manager](#) on page 12.
- **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.

The script is a sample of passwordless SSH configuration between two nodes (either TPVM or server).

- **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 or the servers for server-based deployments. 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."
```

- **IP Address:**

1. Do not use the following IPv4 or IPv6 address subnets which are either reserved for K3s or not supported:
  - a. 10.42.0.0/16 subnet
  - b. 10.43.0.0/16 subnet
  - c. 169.254.0.0/16 subnet
  - d. fd42::/48 subnet
  - e. fd43::/112 subnet
2. Do not use IPv4 mapped IPv6 addresses.  
 Format: 0:0:0:0:FFFF:w.x.y.z or ::FFFF:w.x.y.z  
 Example: ::ffff:10.10.10.10 or ::ffff:0a0a:0a0a
3. Do Not use IPv6 Link Local addresses.

## 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 14: 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

**Table 14: General port requirements (continued)**

Port	Service
5672	Rabbitmq
15672	Rabbitmq management

**Table 15: 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 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 must 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 for Fabric Manager](#) on page 12.

You can find the EFA package tar.gz file under the `/efaboot` directory of the SLX device. This applies to a fresh install or upgrade of EFA. For an incremental EFA image upgrade, you can copy the EFA tar.gz file to the `/efaboot` directory on the SLX device before the deployment.



#### Note

For the supported commands for packet and fabric suites for a single deployment, see [EFA Installer Improvements for Server-Based Deployment](#) on page 22.

## 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.

Specify the configuration of TPVM under the config mode.

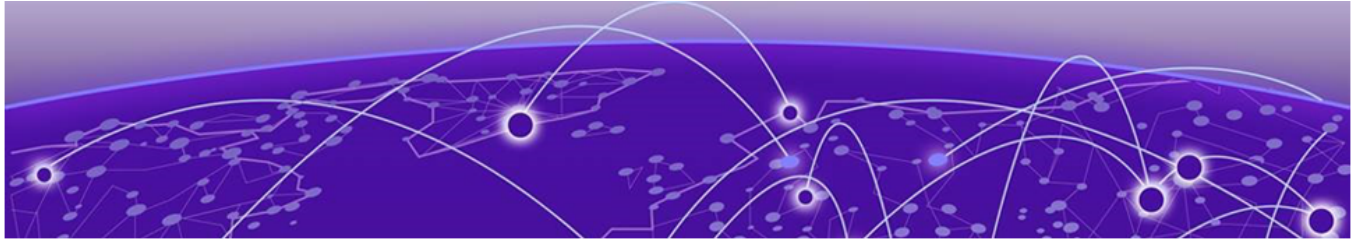
This example configures the NTP, IP, Timezone, Hostname, and DNS configurations.

```
SLX-1# show run tpvm
tpvm TPVM
  auto-boot
  ntp 10.20.53.134
  ntp 10.20.61.191
  dns primary-server 10.31.2.10 secondary-server 10.31.2.11 domain corp.extremenetworks.com
  hostname tpvm
  timezone America/Los_Angeles
  interface management ip 10.20.246.101/20 gw 10.20.240.1
  deploy
!
```

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 24 and [Install EFA on TPVM in a Multi-Node Deployment](#) on page 35.



# EFA Installation for Single-Node Deployments

---

[EFA Installer Improvements for Server-Based Deployment](#) on page 22

[Install EFA in a Single-Node Deployment](#) on page 23

[Install EFA on TPVM in a Single-Node Deployment](#) on page 24

[Deploy the OVA](#) on page 26

[Configure OVA using Postboot Menu](#) on page 27

## EFA Installer Improvements for Server-Based Deployment

---

EFA installer supports packet and fabric suites for server-based single node deployment. The following table provides commands for silent installation:



### Note

In case of failure, the installer automatically collects support-save and unwinds the partial installation.

Check the installer logs in the <Logs directory>/installer directory for any error. For more details, see "Logging and Log Files" in [Extreme Fabric Automation Administration Guide, 3.1.0](#).

Operation	Commands
Installation of packet suite	<code>\$source deployment.sh -i no --deploy-suite packet</code>
Installation of fabric suite	<code>\$source deployment.sh -i no --deploy-suite fabric</code>
Installation of fabric suite with additional management IP	<code>\$source deployment.sh -i no --deploy-suite fabric --sub-intfname intf200 --sub-vlanid 200 --cidr 1.2.3.4/20</code>
Uninstallation	<code>\$source deployment.sh -i no -o undeploy</code>

## Install EFA in a Single-Node Deployment

Install EFA on a single-node server or virtual machine, which is a non-TPVM deployment.

### Before You Begin

Verify that the following minimum virtual machine requirements are met:

- CPU: 4 cores
- Storage: 64 GB



#### Note

Available storage must be at least 30% of the total space available on the disk used

- RAM: 8 GB
- OS: Ubuntu 18.04 or 20.4

Ensure that you have configured NTP according to the [EFA Requirements](#) on page 16.

### About This Task

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



#### Important

Do not use the following IP addresses, which are used by the K3s service:

- 10.42.0.0/16 subnet
- 10.43.0.0/16 subnet
- 169.254.0.0/16 subnet
- fd42::/48 subnet
- fd43::/112 subnet

### Procedure

1. Download the \*.tar.gz image.
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Untar the image.

```
$ tar -xzf efa-3.x.x.tar.gz
```

4. Change to the EFA directory.

```
device# cd efa
```

5. Run the deployment script using the interactive mode.

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

or

Run the deployment using the non-interactive commands shown in the table for installer improvements.

The EFA Installer begins in a series of dialogs.

6. If you selected to install interactively, when prompted, select **Single-node deployment** and **OK**.

**Tip**

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

7. When prompted, select the appropriate suite for the single-node deployment depending on your deployment needs.

**Note**

Fabric must be chosen for managing IP Fabric deployments of SLX devices and packet must be chosen for managing visibility devices.

The `-g no` in the following example is run in a non-interactive mode.

```
root@ubuntu:~/efa# source deployment.sh -g no
Step 1: Checking for EFA Stack...
Please choose: 1 Single-node deployment 2 Multi-node deployment
1
Single-node Deployment
Please choose: 1 Fabric Automation 2 Packet Broker Management
```

8. (Optional) 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.
    - (Optional) IPv6 address in CIDR format. The subnet must not overlap with any IPv6 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.

9. Verify the installation.
  - a. From the EFA command line, run the **efactl status** command to see the status of nodes, pods, and services.
  - b. 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.

## About This Task

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12.

## Procedure

1. Verify that the TPVM is set up for an EFA deployment.
  - a. Verify that the version, SSH keys, and passwordless access configuration are correct for the TPVM via the SLX console or SSH.

For the latest supported version information, see [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12.

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

- b. Verify the versions via the TPVM console or SSH.

```
device# lsb_release -a
```

- c. Verify that NTP is synchronized.

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

- d. 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.

```
device# start-shell
device# scp user@remote-server:~/builds/efa/efa-3.1.0.tar.gz /efaboot/
```

4. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).

5. 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 Address 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
```



### Note

From SLX version 20.4.1 and above, new install or upgrade of EFA on TPVM in a single-node deployment displays the following warning banner on the console:

```
*****
*                               ! ! ! WARNING ! ! !                               *
*   Proceeding with Extreme Fabric Automation deployment                         *
*       1. Do not reboot device(s) or TPVM(s)                                   *
*       2. Do not toggle management port on device(s) or TPVM(s)               *
*       3. Avoid CTRL+C on the installer window                                 *
*****
```

The EFA Installer continues in a series of dialogs.

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

**Tip**

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

7. (Optional) When prompted to configure additional management IP networks, take one of the following steps.
- Select **Yes** and then provide the following information when prompted.
    - 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.
    - An IPv6 address is optional, but an IPv4 address is mandatory.
  - 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 **sudo 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

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 for Fabric Manager](#) on page 12.

There are two OVAs for the users to choose from. Use the OVA image for new installations only.



### Important

Do not use the following IP addresses, which are used by the K3s service:

- 10.42.0.0/16 subnet
- 10.43.0.0/16 subnet
- 169.254.0.0/16 subnet
- fd42::/48 subnet
- fd43::/112 subnet

### Procedure

1. Download the `efa-3.1.0.ova` file for fabric manager or the `xco-xvm-3.1.0.ova` file for visibility manager.
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Deploy the OVA on the hypervisor.
4. Use the following credentials and move to [Configure OVA using Postboot Menu](#) on page 27.

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

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

## Configure OVA using Postboot Menu

### About This Task

The Postboot Menu displays the configuration parameters for setting up root password and static IP address. After you confirm the settings, the system prompts you to reboot the VM for the network IP address to take effect.



### Note

The Postboot Menu is displayed for the first time after you boot the VM. It is not displayed for subsequent reboots.

### Procedure

1. After deploying virtual machine from an OVA file, enter credentials.

The following Welcome screen is displayed on the screen:

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation - Welcome to the Extreme Fabric
Automation Setup
=====

Please enter the information as it is requested to continue with
the configuration. Typically a default value is displayed in brackets.
Pressing the [enter] key without entering a new value will use the
bracketed value and proceed to the next item.

If a default value cannot be provided, the prompt will indicate that the item
```

```
is either (Required) or (Optional). The [enter] key may be pressed without
entering data for (Optional) items. A value must be entered for (Required) items.
```

```
At the end of the setup process, the existing settings will be displayed
and opportunity will be provided to correct any errors.
```

```
=====
Press [enter] to begin setup or CTRL-C to exit:
```

2. To configure root password and IP address and, press **Enter**.  
The following Main Menu is displayed:

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation - Modify Settings
=====
```

```
All of the information needed to complete the installation of the
Extreme Fabric Automation has been entered.
Enter 0 or any key other than a valid selection to continue
If you need to make a change, enter the appropriate number from
the choices listed below.
```

- ```
=====
```
1. Set the root user password
  2. Set network settings
  3. Confirm settings and continue
  4. Exit

```
Enter selection :
```

To exit screen, press **CTRL-C**.

3. To configure root user password, press 1.

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation - Root Password Configuration
=====
```

```
The root password is currently set for this appliance.
```

```
Would you like to set a root password (y/n) [y]? y
```

```
Enter new UNIX password:
```

```
Retype new UNIX password:
```

4. After setting the root password, following main Menu is displayed:

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation - Modify Settings
=====
```

```
All of the information needed to complete the installation of the
Extreme Fabric Automation has been entered.
Enter 0 or any key other than a valid selection to continue
If you need to make a change, enter the appropriate number from
the choices listed below.
```

- ```
=====
```
1. Set the root user password
  2. Set network settings
  3. Confirm settings and continue
  4. Exit

```
Enter selection :
```

```
The screen after user pressed "2" for network settings,
```

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation Interface Configuration
=====
```

```
Configure the interface with static IP . Please choose below option
  1. Static
  2. Quit

Enter selection :
```

5. To configure static IP, press 1. Enter the IP address in CIDR format, AA.BB.CC.DD/EE. This also support dualstack IP configuration.

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation Interface Configuration Static
=====
Enter the IPv4 address in cidr format(Required): 10.37.138.101/20
Enter the IPv4 gateway address (Required): 10.37.128.1
Enter the IPv6 address in cidr format(Optional): 2620:100:c:fe08::222/64
Enter the IPv6 gateway address (Optional): 2620:100:c:fe08::1
Enter the IPv4 nameserver address (Optional):

These are the correct network settings that will be used to configure.
=====
Address type: Static
IPv4 Address: 10.37.138.215/20
IPv6 Address: 2620:100:c:fe08::222/64
IPv4 Gateway Address: 10.37.128.1
IPv6 Gateway Address: 2620:100:c:fe08::1
=====

Woud you like to accept the current network settings (y/n) [y]? _
```

6. To confirm the settings, press 3.

System prompts you to reboot the VM for the network settings to take effect.

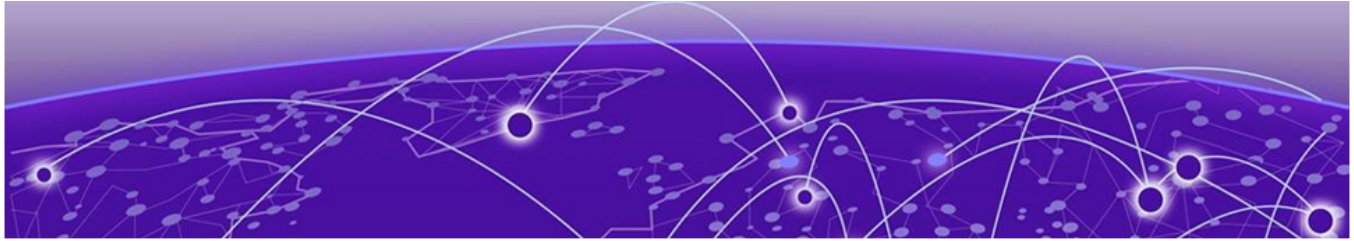
```
=====
Extreme Networks, Inc. - Extreme Fabric Automation - Modify Settings
=====
All of the information needed to complete the installation of the
Extreme Fabric Automation has been entered.
Enter 0 or any key other than a valid selection to continue
If you need to make a change, enter the appropriate number from
the choices listed below.
=====

  1. Set the root user password
  2. Set network settings
  3. Confirm settings and continue
  4. Exit

Enter selection :3

These are the current settings that will be used to configure.
=====
Address type: Static
IP Address: 10.37.138.101/20
Gateway Address: 10.37.128.1
Nameserver Address: 10.37.2.1
=====

Would you like to accept the current network settings and REBOOT (y/n) [y]?
```



# EFA Installation for Multi-Node Deployments

---

[EFA Deployment for High Availability](#) on page 30

[Install EFA in a Multi-Node Deployment](#) on page 33

[Install EFA on TPVM in a Multi-Node Deployment](#) on page 35

## EFA Deployment for High Availability

---

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. The standby node runs processes all the requests when the active node fails.

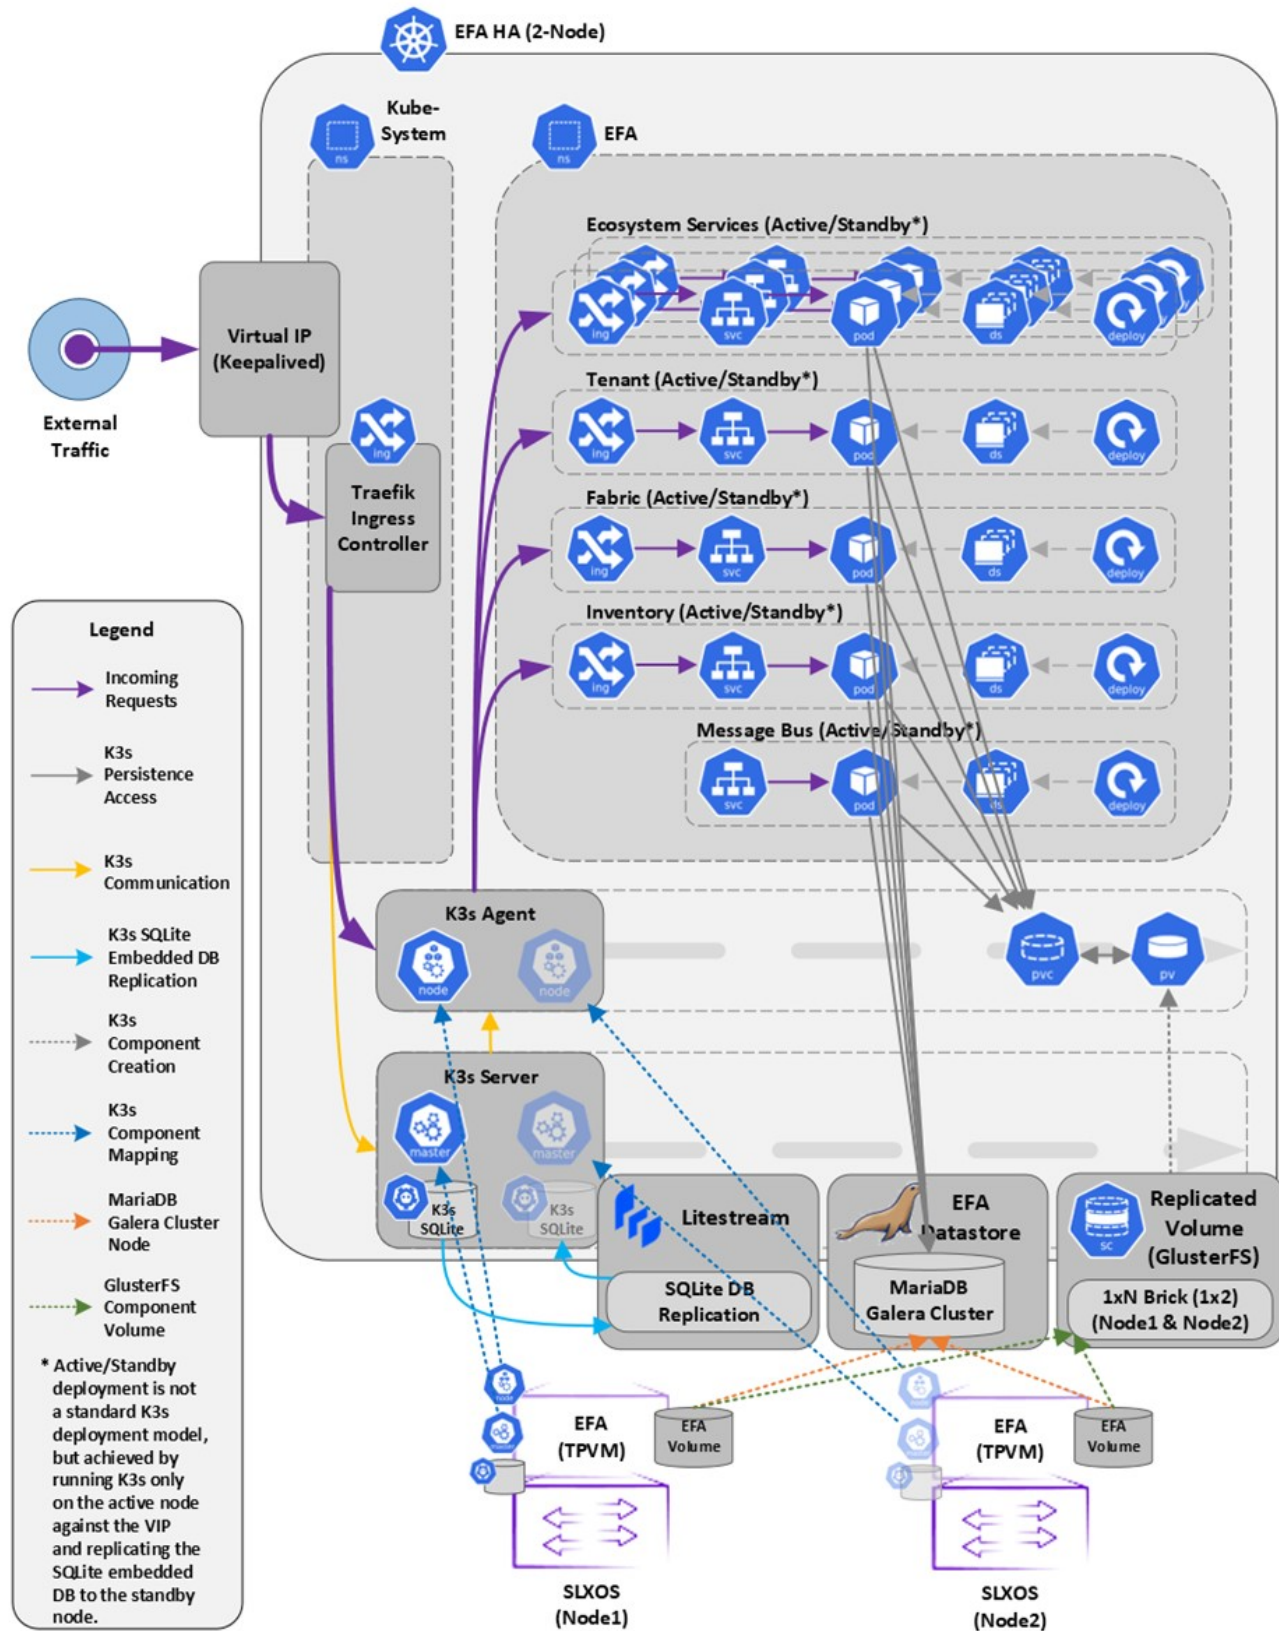


Figure 1: Two-node high-availability deployment

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, 3.1.0](#).

EFA uses the following services to implement an HA deployment:

- Keepalived (VRRP) – It is a program which runs on both nodes. The active node frequently sends VRRP packets to the standby node. If the active node stops sending the packets, keepalived on the standby assume the active role. Thus, the standby node becomes the active node. Each state change runs a keepalived notify script containing logic to ensure EFA's continued operation after a failure. With a two-node cluster, a "split-brain" may occur due to a network partition which leads to two active nodes. When the network recovers, VRRP establishes a single active node that determines the state of EFA.
- K3s server runs on the active node. Kubernetes state is stored in SQLite and is synced in real-time to the standby node using a dedicated daemon, litestream. On a failover, the keepalive notify script on the new active node reconstructs the Kubernetes SQLite DB from the synced state and starts the k3s. K3s runs on one node at a time, not on both nodes and hence the HA cluster looks like a single-node cluster, however, the HA cluster ties itself to the keepalived-managed virtual IP.
- MariaDB and Galera – EFA business states (device, fabric, and tenant registrations and configuration) are stored in a set of databases managed by MariaDB. Both the nodes run a MariaDB server, and the Galera clustering technology is used to keep the business state in sync on both the nodes during normal operation.
- Glusterfs – This is a clustering filesystem used to store EFA's log files, various certificates, and subinterface definitions. A daemon runs on both the nodes which seamlessly syncs several directories.



#### Note

Although Kubernetes run as a single-node cluster tied to the virtual IP, EFA CLIs still operate correctly when they are run from active or standby node. Commands are converted to REST and issued over HTTPS to the ingress controller via the virtual IP tied to the active node.

The **efa status** confirms the following:

- For the active node:
  - All enabled EFA services are reporting Ready
  - Kubernetes state is consistent with all the enabled EFA services (for example, service endpoints exist)
  - The host is a member of the Galera or MariaDB cluster
- For the standby node:
  - It is reachable via SSH from the active node
  - It is a member of the Galera or MariaDB cluster
- For both the nodes:
  - The Galera cluster size is 2 if both the nodes are up, and the cluster size is  $\geq 1$  if the standby node is down.

**Example:**

```
NH-1# show efa status
=====
                        EFA version details
=====
Version : 3.1.0
Build: 109
Time Stamp: 22-10-25:12:45:44
Mode: Secure
Deployment Type: multi-node
Deployment Platform: TPVM
Deployment Suite: Fabric Automation
Virtual IP: 10.20.246.103
Node IPs: 10.20.246.101,10.20.246.102
--- Time Elapsed: 8.512402ms ---

=====
                        EFA Status
=====
+-----+-----+-----+-----+
| Node Name | Role   | Status | IP           |
+-----+-----+-----+-----+
| tpvm2     | active | up     | 10.20.246.102 |
+-----+-----+-----+-----+
| tpvm      | standby | up    | 10.20.246.101 |
+-----+-----+-----+-----+
--- Time Elapsed: 19.168973841s --
```

## Install EFA in a Multi-Node Deployment

You can install EFA in a multi-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 16.

**About This Task**

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

**Note**

EFA Management Interface must have IPv4 address configured. Adding IPv6 address is optional while IPv4 is mandatory during sub interface creation.

**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 the **Fabric suite**. Then, select **OK** as needed to progress.

**Tip**

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

5. When prompted, enter the peer IP address or FQDN of the other node in the cluster.
6. When prompted, enter the virtual IP address for the cluster.
7. (Optional) When prompted, enter the virtual IPv6 address for the cluster.
  - Select **Yes** and then provide the virtual IPv6 addresses.
  - Select **No** to ignore this optional step.
8. (Optional) When prompted to configure additional IP addresses for a health check, take one of the following steps.
  - Select **Yes** and then provide the IPv4 or IPv6 addresses.
  - Select **No** to ignore this optional step.
9. (Optional) 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.
    - An IPv6 address is optional, but an IPv4 address is mandatory.
  - Select **No** to ignore this optional step or when you have finished entering network information.
10. (Optional) When prompted to configure additional management IP network routes, take one of the following steps.
  - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
    - Target network IP address in CIDR format
    - Source IP address for outbound traffic
    - Next-hop or gateway IP address through which access to the destination network is provided
  - Select **No** to ignore this optional step or when you have finished entering route information.

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

11. Verify the installation.
  - a. From the EFA command line, run the **sudo efactl status** command to see the status of nodes, pods, and services.

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

## Install EFA on TPVM in a Multi-Node Deployment

You can install EFA on a TPVM (Third-Party Virtual Machine) in a two-node deployment for high availability.

### Before You Begin

Ensure that the EFA tar file is available on the `/efaboot` partition of the SLX device.

### About This Task

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12.

### Procedure

1. Run the **show tpvm status** command and 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 for Fabric Manager](#) on page 12.
  - b. Verify that the TPVM has an assigned IP address.
  - c. Verify that the SSH keys are uploaded.
  - d. (For SLX-OS releases earlier than 20.2.3) Verify that passwordless access is configured.
  - 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.
  - f. Verify that the NTP is configured on TPVM by running the **show run tpvm** command. If NTP is not configured, configure it using the **tpvm config ntp add server <ip>** command.

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

- g. Verify that NTP is synchronized.
  - h. If necessary, log in to TPVM and configure the NTP time zone from SLX.

```
device# tpvm config timezone
```

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

```
device# tpvm config host
```

2. Enter SLX Linux mode.

```
device# start-shell
```

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
```



### Note

From SLX version 20.4.1 and above, new install or upgrade of EFA on TPVM in a multi-node deployment displays the following warning banner on the console:

```
*****
*                               ! ! ! WARNING ! ! !                               *
*   Proceeding with Extreme Fabric Automation deployment                         *
*   1. Do not reboot device(s) or TPVM(s)                                       *
*   2. Do not toggle management port on device(s) or TPVM(s)                   *
*   3. Avoid CTRL+C on the installer window                                     *
*****
```

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, 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. (Optional) When prompted, enter the virtual IPv6 address for the cluster.
  - Select **Yes** and then provide the virtual IPv6 addresses.
  - Select **No** to ignore this optional step.
9. (Optional) When prompted to configure additional IP addresses for a health check, take one of the following steps.
  - Select **Yes** and then provide the IPv4 or IPv6 addresses.
  - Select **No** to ignore this optional step.
10. (Optional) 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.
    - An IPv6 address is optional, but an IPv4 address is mandatory.
  - Select **No** to ignore this optional step or when you have finished entering network information.

11. (Optional) When prompted to configure additional management IP network routes, take one of the following steps.

**Note**

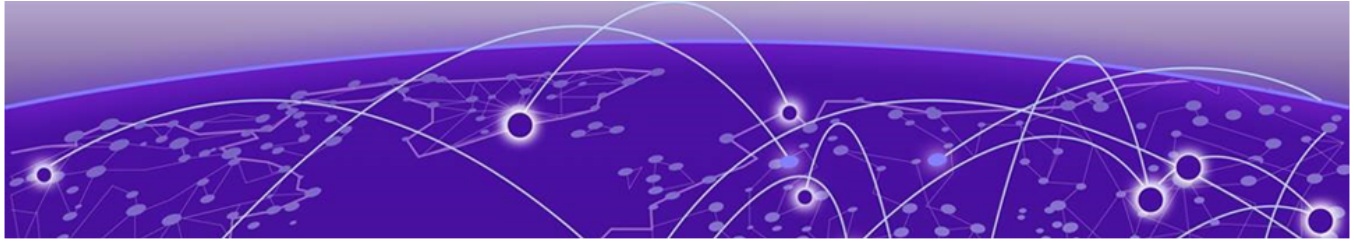
EFA Management Interface must have IPv4 address configured. Adding IPv6 address is optional while IPv4 is mandatory during sub interface creation.

- Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
  - Target network IP address in CIDR format
  - Source IP address for outbound traffic
  - Next-hop or gateway IP address through which access to the destination network is provided
- Select **No** to ignore this optional step or when you have finished entering route information.

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

12. 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

---

- [Upgrade EFA in a Single-node Deployment](#) on page 38
- [Upgrade EFA on TPVM in a Single-node Deployment](#) on page 39
- [Upgrade EFA from a Single-Node to a Multi-Node Deployment](#) on page 41
- [Upgrade EFA in a Multi-Node Deployment](#) on page 43
- [Upgrade EFA on TPVM in a Multi-Node Deployment](#) on page 44
- [Upgrade EFA on an OVA](#) on page 46
- [Upgrading SLX-OS, TPVM, and EFA Together](#) on page 47
- [Recover from an Upgrade Failure](#) on page 52
- [Rollback](#) on page 53

You can upgrade EFA from either of the two previous releases to the latest release.

## Upgrade EFA in a Single-node Deployment

---

Expect the upgrade process to take approximately 8 to 10 minutes, during which EFA services are down.

### About This Task

The upgrade process takes backup of the EFA system before starting the procedure. In case of any failures in upgrade, use this backup to recover the data. For more information, see [Recover from an Upgrade Failure](#) on page 52.

### Procedure

1. Download the image (\*.tar.gz).
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Untar the image.

```
device # tar -xvzf efa-v3.x.x.tar.gz
~/builds/3.1.0/tmp $ tar -xfz efa-3.1.0-58.tar.gz
tar: z: Cannot open: No such file or directory
tar: Error is not recoverable: exiting now
sbr@sbr-virtual-machine
~/builds/3.1.0/tmp $
```

4. Change to the EFA directory.

```
device# cd efa
```

5. Run the following deployment script without any optional parameters.

```
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.
    - An IPv6 address is optional, but an IPv4 address is mandatory.
  - 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. From the EFA command line, run the **sudo efactl status** command to see the status of nodes, pods, and services.
  - b. 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. If more than one EFA version is available in the `/efaboot` directory, you are prompted to select a version during upgrade.

### About This Task

EFA does not support *Non-secure mode*. For more information about the modes, see [#unique\\_33](#).

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 52.

### Procedure

1. Enter SLX Linux mode.

```
device# start-shell
```

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

```
scp <username>@<hostip>:<buildpath>/efa-3.1.0.tar.gz /efaboot
```

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

```
device# efa deploy
```



#### Note

From SLX version 20.4.1 and above, new install or upgrade of EFA on TPVM in a multi-node deployment displays the following warning banner on the console:

```
*****
*                               ! ! ! WARNING ! ! !                               *
*   Proceeding with Extreme Fabric Automation deployment                         *
*   1. Do not reboot device(s) or TPVM(s)                                       *
*   2. Do not toggle management port on device(s) or TPVM(s)                   *
*   3. Avoid CTRL+C on the installer window                                       *
*****
```

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 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.
  - (Optional) IPv6 address in CIDR format. The subnet must not overlap with any IPv6 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.

### 6. 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 **sudo 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 38.

Ensure you have completed the high-availability prerequisites in [EFA Requirements](#) on page 16.

### About This Task

Expect the upgrade process to take approximately 15 - 25 minutes, during which EFA services are down.

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 52.



#### Note

EFA management Interface must have both IPv4 and IPv6 address configured. Adding IPv6 address is optional while IPv4 is mandatory during sub interface creation.

### Procedure

1. Download the image (\*.tar.gz).
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Untar the image.

```
device# tar -xvzf efa-v3.x.x.tar.gz
```

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, 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, enter the virtual IPv6 address for the cluster.
  - Select **Yes** and then provide the virtual IPv6 addresses.
  - Select **No** to ignore this optional step.

10. When prompted to configure additional IP addresses for a health check, take one of the following steps.

- Select **Yes** and then provide the IPv4 or IPv6 addresses.
- Select **No** to ignore this optional step.

11. When prompted to configure additional management IP networks, take one of the following steps.

You can add only one management IP networks during upgrade. After the upgrade, you can add more than one management IP networks.

- 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 4090.
  - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
  - An IPv6 address is optional, but an IPv4 address is mandatory.
- Select **No** to ignore this optional step or when you have finished entering network information.

12. When prompted to configure additional management IP network routes, take one of the following steps.

- Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
  - Target network IP address in CIDR format
  - Source IP address for outbound traffic
  - Next-hop or gateway IP address through which access to the destination network is provided
- Select **No** to ignore this optional step or when you have finished entering route information.

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

13. Verify the upgrade.

- a. (If applicable) 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 **sudo 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

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 automatically backs up the EFA database, so that you can easily recover data if the upgrade fails. You can apply this procedure on either of the nodes.

### Procedure

1. Download the image (\*.tar.gz) to a new sub-folder.
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Untar the image.

```
device# tar -xvzf efa-v3.x.x.tar.gz
```

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, enter the virtual IPv6 address for the cluster.
  - Select **Yes** and then provide the virtual IPv6 addresses.
  - Select **No** to ignore this optional step.
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.



#### Note

EFA management Interface must have both IPv4 and IPv6 address configured. Adding IPv6 address is optional while IPv4 is mandatory during sub interface creation. You can add only one management IP networks during upgrade. After the upgrade, you can add more than one management IP networks.

- 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 4090.
- IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.

**Note**

An IPv6 address is optional, but an IPv4 address is mandatory.

- Select **No** to ignore this optional step or when you have finished entering network information.

10. When prompted to configure additional management IP network routes, take one of the following steps.

An IPv6 address is optional, but an IPv4 address is mandatory.

- Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
  - Target network IP address in CIDR format
  - Source IP address for outbound traffic
  - Next-hop or gateway IP address through which access to the destination network is provided
- Select **No** to ignore this optional step or when you have finished entering route information.

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

11. Verify the upgrade.

- From the EFA command line, run the `sudo efactl status` command to see the status of nodes, pods, and services.
- From the EFA command line, run the `efa status` command for concise status information.

## Upgrade EFA on TPVM in a Multi-Node Deployment

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.

### About This Task

EFA on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12.

By default, EFA is installed in secure mode. For more information about the modes, see [#unique\\_33](#).

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
```



### Note

From SLX version 20.4.1 and above, new install or upgrade of EFA on TPVM in a multi-node deployment displays the following warning banner on the console:

```
*****
*                               ! ! ! WARNING ! ! !                               *
*   Proceeding with Extreme Fabric Automation deployment                         *
*   1. Do not reboot device(s) or TPVM(s)                                       *
*   2. Do not toggle management port on device(s) or TPVM(s)                   *
*   3. Avoid CTRL+C on the installer window                                     *
*****
```

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, enter the peer IP address or FQDN of the other node in the cluster.
6. When prompted, enter the virtual IPv6 address for the cluster.
  - Select **Yes** and then provide the virtual IPv6 addresses.
  - Select **No** to ignore this optional step.
7. When prompted to configure additional IP addresses for a health check, take one of the following steps.
  - Select **Yes** and then provide the IPv4 or IPv6 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.
    - An IPv6 address is optional, but an IPv4 address is mandatory.
  - Select **No** to ignore this optional step or when you have finished entering network information.
9. When prompted to configure additional management IP network routes, take one of the following steps.



#### Note

EFA management Interface must have both IPv4 and IPv6 address configured. Adding IPv6 address is optional while IPv4 is mandatory during sub interface creation.

- Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
  - Target network IP address in CIDR format
  - Source IP address for outbound traffic
  - Next-hop or gateway IP address through which access to the destination network is provided
- Select **No** to ignore this optional step or when you have finished entering route information.

The installation 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 **sudo 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](#) on page 26 for a list of prerequisites.

## Procedure

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

```
# sudo su -
```

3. Copy the OVA build to any location.

```
ubuntu@efa:~$ pwd
/home/ubuntu
ubuntu@efa:~$ scp root@10.20.241.29:/opt/efa_releases/builds/tar/efa-3.1.0.tar.gz .
The authenticity of host '10.20.241.29 (10.20.241.29)' can't be established.
ECDSA key fingerprint is SHA256:Zaj8WMvESDrF0Th2vanA9m085d4FOtzEGCJfMrSqT3U.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.20.241.29' (ECDSA) to the list of known hosts.
root@10.20.241.29's password:
efa-3.1.0.tar.gz                                100% 880MB
90.1MB/s  00:09
ubuntu@efa:~$ ls
efa-3.1.0.tar.gz
ubuntu@efa:~$
```

4. Extract the tar.

```
device# tar -xvzf efa-3.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
```

8. Verify the upgrade:

- a. From the EFA CLI, run the `efactl status` command to view the status of the nodes, pods, and services.
- b. From the EFA CLI, run the `efa status` command to get the concise status information.

## Upgrading SLX-OS, TPVM, and EFA Together

Use the topics in this section to learn how to upgrade an SLX device to the latest supported version of SLX-OS with the latest supported version of TPVM.

### Requirement for SCP connections

The firmware server must support more than 10 unauthenticated SCP connections. You can ensure this requirement by specifying an appropriate value for `'#MaxStartups 10:30:100'` in the `/etc/ssh/sshd_config` file on the firmware server.

The following is an example of an appropriate value:

Full is greater than Start and Start is greater than the number of devices in the fabric.

Run `$ sudo systemctl daemon-reload` to restart the `sshd` service for the changes to the `/etc/ssh/sshd_config` file to take effect..

Run `$ sudo systemctl restart sshd.service` to restart the `sshd` service.

Restarting the `sshd` service does not affect any connected SSH sessions.

## Upgrade EFA, SLX-OS, and TPVM Method 1

Use this upgrade method if the old base version of TPVM is newer than 4.4.0.

### About This Task

This option is the preferred method for upgrading EFA, SLX-OS, and TPVM. For more information about supported versions, see [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12.

In the following procedure, **SLX1** refers to the active EFA node (TPVM1). **SLX2** refers to the standby EFA node (TPVM2).

### Procedure

1. Upgrade EFA to the latest version.

- a. Back up EFA.

```
efa system backup
```

For more information, see "Back up and Restore the EFA System" in the [Extreme Fabric Automation Command Reference, 3.1.0](#).

- b. SCP the backup file to a location outside of TPVM, such as the `/efaboot` partition of SLX-OS where the EFA image is kept.
- c. Copy the EFA image to the `/efaboot` directory on SLX1.
- d. Deploy EFA on any of the SLX.

```
efa deploy
```

- e. When prompted, select **Multi Node Build Upgrade**.



#### Note

If the upgrade process returns `cfg-refreshed`, run a manual Drift and Reconcile on all devices.

2. Upgrade SLX-OS to the latest version.

An SLX-OS upgrade from 20.2.3x to 20.3.2x needs a full install. The procedure performs fabric-wide firmware download by staging the devices in multiple groups with no traffic disruption. Complete the following steps to download firmware on all the devices in the fabric.

- a. From the EFA command line on TPVM1 (the active node), upgrade SLX2 to the latest SLX-OS version.

```
efa inventory firmware-host register --ip <fw-host-ip>
--protocol scp --username <username> --password <password>
```

- b. From the EFA command line on SLX1, upgrade SLX-OS from 20.2.3x to 20.3.2b.

```
efa inventory device firmware-download prepare add --fabric <fabric name>
--firmware-host <fw-host-ip> --firmware-directory <fw-path>
```

```
efa inventory device firmware-download prepare list --fabric <fabric name>

efa inventory device firmware-download execute --fabric <fabric name>

efa inventory device firmware-download show --fabric <fabric name>
```

3. From the EFA command line, upgrade TPVM1 (SLX1) and TPVM2 (SLX2) to the latest TPVM version using TPVM incremental upgrade image.

For more details, refer [TPVM Incremental Upgrade](#) on page 66.

- a. Back up EFA.

```
efa system backup
```

- a. Verify the TPVM status on SLX1 and SLX2. Ensure both TPVMs are in running state.

```
device# show tpvm status
```

- b. From the active EFA command line run the following command to upgrade TPVM1 and TPVM2.

This is applicable for SLX version 20.4.1 and EFA version 3.0.0 and above.

```
efa inventory device tpvm-upgrade execute --ip<SLX1-IP>,<SLX2-IP>,  
--firmware-host <fw-host-ip> --tpvm-image <path-for-tpvm_inc_upg.deb>
```

- c. From the EFA command line, verify the TPVM upgrade process.

```
efa inventory device tpvm-upgrade show --ip <SLX1-IP>,<SLX2-IP>
```

- d. When the status of the upgrade is complete, perform the following (from the EFA command line) on both nodes.

Run **efa status** to verify that both nodes are up.

Run **sudo efactl status** to verify that all pods on the active node are in Running state.

Run **sudo efactl db-status** to verify that the MariaDB is active (running)

- e. If there is a “System restart required” message in “efa inventory device tpvm-upgrade show” or on TPVM consoles after the upgrade of TPVMs, reboot the TPVM2 (standby) first, and wait for TPVM2 to come up. This step ensures that the services are running with “efactl status” followed by the reboot of TPVM1 (active).

## Upgrade EFA, SLX-OS, and TPVM Method 2

Use this upgrade method if the TPVM base version is older than 4.4.0.

### About This Task

This option is the preferred method for upgrading EFA, SLX-OS, and TPVM. For more information about supported versions, see [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12.

In the following procedure, **SLX1** refers to the active EFA node (TPVM1). **SLX2** refers to the standby EFA node (TPVM2).

## Procedure

1. Upgrade EFA to the latest version.

- a. Back up EFA.

```
efa system backup
```

For more information, see "Back up and Restore the EFA System" in the [Extreme Fabric Automation Administration Guide, 3.1.0](#).

- b. SCP the backup file to a location outside of TPVM, such as the /efaboot partition of SLX-OS where the EFA image is kept.
  - c. Copy the EFA image to the /efaboot directory on SLX1.
  - d. Deploy EFA on SLX1.

```
efa deploy
```

- e. When prompted, select **Multi Node Build Upgrade**.



### Note

If the upgrade process returns `cfg-refreshed`, run a manual DRC on all devices.

2. Upgrade SLX-OS to the latest version.

An SLX-OS upgrade from 20.2.3x to 20.3.2x needs a full install. The procedure performs fabric-wide firmware download by staging the devices in multiple groups with no traffic disruption. Complete the following steps to download firmware on all the devices in the fabric.

- a. From the EFA command line on TPVM1 (the active node), upgrade SLX2 to the latest SLX-OS version.

```
efa inventory firmware-host register --ip <fw-host-ip>
--protocol scp --username <username> --password <password>
```

- b. From the EFA command line on SLX1, upgrade SLX-OS from 20.2.3x to 20.3.2b.

```
efa inventory device firmware-download prepare add --fabric <fabric name>
--firmware-host <fw-host-ip> --firmware-directory <fw-path>

efa inventory device firmware-download prepare list --fabric <fabric name>

efa inventory device firmware-download execute --fabric <fabric name>

efa inventory device firmware-download show --fabric <fabric name>
```

3. From the EFA command line, upgrade TPVM2 (SLX2) to the latest TPVM version.



### Note

Ensure that you upgrade TPVM2 first because it is in the standby node of EFA.

For more information, see [TPVM Complete Package Upgrade](#) on page 60

- a. Back up EFA.

```
efa system backup
```

- b. Verify the trusted-peer configuration on SLX1 and SLX2.

```
device# show tpvm config trusted-peer
U37-55-172# show tpvm config trusted-peer
root@10.20.55.175
```

- c. If a trusted-peer is present on at least one node, from the EFA command line on TPVM1, run the following command to upgrade TPVM2.

```
efa inventory device tpvm-upgrade execute <SLX2-IP>
--firmware-host <fw-host-ip> --tpvm-image <path-for-tpvm.deb>
```

- d. If a trusted-peer is not present on either node, from the EFA command line on TPVM1, run the following command upgrade TPVM2.

```
efa inventory device tpvm-upgrade execute <SLX2-IP>
--firmware-host <fw-host-ip> --tpvm-image <path-for-tpvm.deb>
--trusted-peer-sudo-user <username> --trusted-peer-password <password>
```

- e. From the EFA command line, verify the TPVM upgrade process.

```
efa inventory device tpvm-upgrade show --ip <SLX2-IP>
```

- f. When the status of the upgrade is complete, perform the following (from the EFA command line) on both nodes.

Run **efa status** to verify that both nodes are up.

Run **sudo efactl status** to verify that all pods on the active node are in Running state.

Run **sudo efactl db-status** to verify that the MariaDB is active (running)

- g. (Optional) Verify the TPVM status on SLX2.

```
device# show tpvm status
```



#### Note

- With SLX-OS 20.3.2a and later, TPVM configuration is persisted in subsequent SLX-OS upgrades, so you do not need to configure TPVM in these upgrades. When an SLX device that hosts TPVM is upgraded to 20.3.2a, the existing TPVM continues to run, and all TPVM parameters that were configured with the **tpvm config** command are converted to TPVM **config block** commands. An exception is the "trusted-peer" configuration, which you must manually redo after the upgrade, unless you provide trusted peer parameters when you run **efa inventory device tpvm-upgrade execute**.

For information about TPVM configuration block and migration, see the *Extreme SLXOS Management Configuration Guide*.

- Do not run the **efa inventory device tpvm-upgrade execute** command if the TPVM upgrade is in progress.

#### 4. Upgrade TPVM1 (SLX1) to the latest TPVM version.

- a. From the SLX-OS command line on SLX1, stop and start TPVM to force a failover.

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

- b. When EFA synchronizes after the failover, view the output of the following commands to ensure that both nodes are in their proper state.

Run **efa status** to verify that both nodes are up.

Run **sudo efactl status** to verify that all pods on the active node are in Running state.

Run **sudo efactl db-status** to verify that the MariaDB is active (running)

- c. From the EFA command line on TPVM2 (the active EFA), upgrade TPVM.

```
efa inventory device tpvm-upgrade execute --ip <slx-hosting-stby-efa>
--firmware-host <firmware-host-ip> --tpvm-image <image-path-on-host>
```

- d. From the EFA command line, verify the TPVM upgrade process.

```
efa inventory device tpvm-upgrade show --ip <SLX1-IP>
```

- e. If the upgrade process shows a failure, take the following steps.

Run **device# show run tpvm** to verify whether the trusted-peer on the SLX device is configured with the correct IP address.

If the IP address is incorrect, correct it manually and repeat the upgrade process starting with step 4.c in [Upgrade EFA, SLX-OS, and TPVM Method 2](#) on page 49.

- f. When the upgrade is complete, perform the following (from the EFA command line) on both nodes.

Run **efa status** to verify that both nodes are up.

Run **sudo efactl status** to verify that all pods on the active node are in Running state.

Run **sudo efactl db-status** to verify that the MariaDB is active (running).

- g. (Optional) Verify the TPVM status on SLX1.

```
device# show tpvm status
```

## Recover from an Upgrade Failure

You have the option to recover from an upgrade failure by rerunning the upgrade, or perform a fresh installation, and then restore the system from a backup.

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 41
- [Upgrade EFA on TPVM in a Single-node Deployment](#) on page 39
- [Upgrade EFA in a Single-node Deployment](#) on page 38
- [Upgrade EFA in a Multi-Node Deployment](#) on page 43
- [Upgrade EFA on TPVM in a Multi-Node Deployment](#) on page 44

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 Single-Node and Multi-Node Deployment](#) on page 59
    - [Uninstall EFA in a Single-Node or Multi-Node Deployment](#) on page 59
  - b. Follow the steps for the type of installation you need.
    - [Install EFA on TPVM in a Multi-Node Deployment](#) on page 35
    - [Install EFA in a Multi-Node Deployment](#) on page 33
    - [Install EFA on TPVM in a Single-Node Deployment](#) on page 24
    - [Install EFA in a Single-Node Deployment](#) on page 23
  - c. Restore the EFA backup.

```
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 Command Reference, 3.1.0](#).

## Rollback

Initiate a rollback when there is a deployment failure to ensure data consistency. You can rollback a particular component based on the error or faulty component.

### Maintain TPVM Versions After a Rollback in a Multi-Node Deployment

Both nodes in a multi-node deployment should have the same version of TPVM after an upgrade.

#### About This Task

This procedure addresses a scenario in which TPVM2 (on SLX2) was upgraded, but TPVM1 (on SLX1) was rolled back to a previous version because of a failure during upgrade. To maintain the same version of TPVM on both nodes, you must downgrade, or roll back, TPVM2.

In this procedure, SLX1 and TPVM1 refer to the standby EFA node. SLX2 and TPVM2 refer to the active EFA node. This procedure references TPVM versions 4.2.4 and 4.2.5 for clarity in examples.

#### Procedure

1. From the SLX-OS command line on SLX2, stop and start TPVM to force a failover.

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

2. When EFA synchronizes after the failover, view the output of the following commands (run from TPVM) to ensure that both nodes are in their proper state.
  - a. Run **efa status** to verify that both nodes are up.

- b. Run **efactl status** to verify that all pods on the active node are in Running state.
  - c. Run **efactl db-status** to verify that the MariaDB is active (running).
3. From the EFA command line on TPVM1 (the active EFA), upgrade TPVM.

```
efa inventory device tpvm-upgrade execute --ip <slx-hosting-stby-efa>
--firmware-host <firmware-host-ip> --tpvm-image <image-path-on-host>
```

4. From the EFA command line, verify the TPVM upgrade process.

```
efa inventory device tpvm-upgrade show --ip <SLX2-IP>
```

5. If the upgrade process (step 3) fails, take the following steps.

- a. Delete the TPVM on both SLX devices.

```
device# tpvm uninstall force
```

In the sample scenario, you are deleting version 4.2.5 from the upgraded device and deleting version 4.2.4 from the device on which the TPVM was rolled back.

- b. Install the earlier version of the TPVM on both devices.

In the sample scenario, you are installing version 4.2.4 on both devices, so that both devices have the same version of TPVM.

- c. Install EFA on the TPVM.

For more information, see [Install EFA on TPVM in a Multi-Node Deployment](#) on page 35.

## Rollback SLX

Initiate a rollback when there is a deployment failure.

### Procedure

Run the following commands to download the previous installed version.

```
efa inventory firmware-host register --ip <fw-host-ip> --protocol scp --username
<username> --password <password>
efa inventory device firmware-download prepare add --fabric <fabric name> --firmware-host
<fw-host-ip> --firmware-directory <fw-path>
efa inventory device firmware-download prepare list --fabric <fabric name>
efa inventory device firmware-download execute --fabric <fabric name>
efa inventory device firmware-download show --fabric <fabric name>
```

## Rollback EFA

Initiate a rollback when there is a deployment failure.

### Procedure

1. Unwind the partial installation or undeploy the failed EFA instance.

```
no efa deploy
```

2. Copy the EFA instance.

```
efa deploy
```

3. Use the system backups available in the /apps/efa\_logs/backup/ directory or copy the required backup files to the /apps/efa\_logs/backup/ directory.

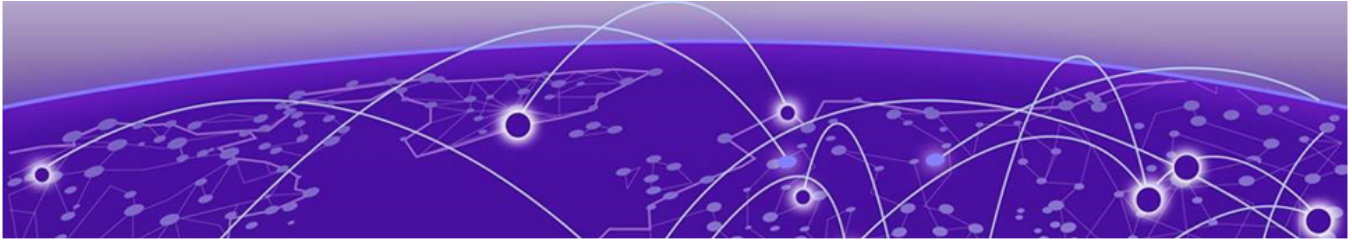
```
extreme@tpvm1:~$ scp root@10.20.48.170:/home/user/EFA-3.1.0-
GA-2022-10-20T07-08-43.921.tar /apps/efa_logs/backup/
```

```
The authenticity of host '10.20.48.170 (10.20.48.170)' can't be established.  
ECDSA key fingerprint is SHA256:rQYa5NjeFWtLvCCUzjELs+9jd/6E+hBeEeHIYdFBs2I.  
Are you sure you want to continue connecting (yes/no)? yes  
Warning: Permanently added '10.20.48.170' (ECDSA) to the list of known hosts.  
root@10.20.48.170's password:  
EFA-3.1.0-GA-2022-10-20T07-08-43.921.tar 100%  
732KB 18.6MB/s 00:00  
extreme@tpvm-71:~$
```

4. Login as an 'extreme' user to the EFA system.
5. Restore the EFA configuration.

```
efa system restore --backup-tar <file_name>
```

6. Post EFA rollback, if you see any devices in `cfg refresh error` state when you run the **efa fabric show** command, run the **efa inventory device update --ip <device\_ip>** command to correct it.



# Node Replacement

---

[Replace a Node in a Multi-node Deployment](#) on page 56

[Replace a Node in a Multi-node TPVM Deployment](#) on page 57

## 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 2.4.4 or later.
- Ensure you have completed the high-availability prerequisites in [EFA Requirements](#) on page 16.
- Ensure that EFA is not deployed on the replacement node.
- Ensure that the faulty node is shutdown.

### 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 active 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.

```
$ efa status
```

For more information on how to recover SLX configs, refer to the [Extreme Fabric Automation Administration Guide, 3.1.0](#).

## 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 16.

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 active node where EFA is installed.

### Procedure

1. Enter SLX Linux mode and copy the EFA tar file to the SLX device.

```
device# start-shell

scp <username>@<hostip>:<buildpath>/efa-3.1.0.tar.gz
```

2. 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 Address 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.

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 and select **OK**.

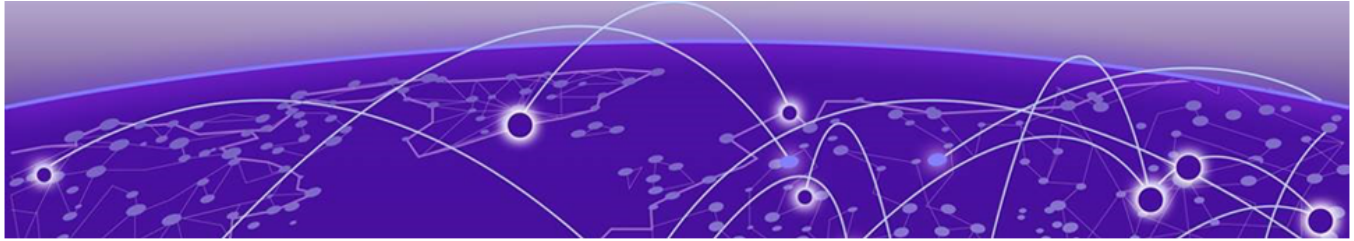
The node replacement proceeds. Messages indicate the progress and when the replacement is complete.

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

```
device# sudo efactl status
```

**Note**

To recover the SLX configuration, see the "Replace a Faulty Device" topic in the [Extreme Fabric Automation Administration Guide, 3.1.0](#).



# EFA Uninstallation

---

[Uninstall EFA in a Single-Node or Multi-Node Deployment](#) on page 59

[Uninstall EFA on TPVM in a Single-Node and Multi-Node Deployment](#) on page 59

## Uninstall EFA in a Single-Node or Multi-Node Deployment

---

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

### Procedure

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

```
source deployment.sh --operation undeploy --interactive no
```

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

## Uninstall EFA on TPVM in a Single-Node and Multi-Node Deployment

---

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

### Procedure

1. (On a Single-Node deployment) From the SLX device console, uninstall EFA.

```
device# no efa deploy
```

- a. When prompted to continue, enter *y*.

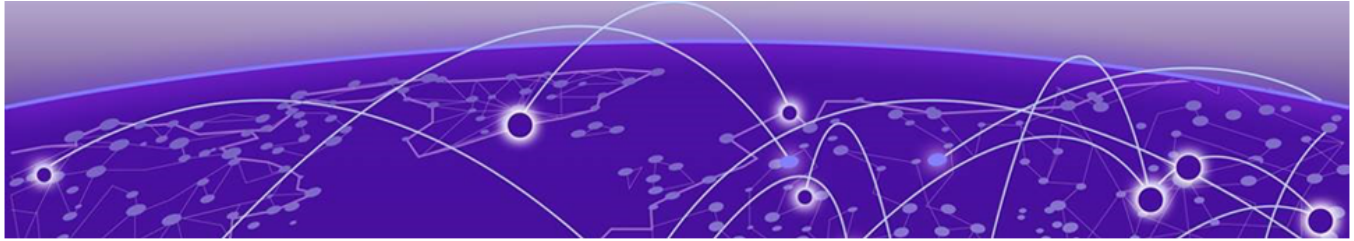
The uninstall process proceeds. A message indicates when the EFA stack is uninstalled.

2. (On a Multi-Node deployment) On the node where EFA is installed, run the deployment script.

```
# source deployment.sh
```

- a. When prompted, select **Remove the current EFA Stack**.

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



# TPVM Upgrade from EFA

---

[TPVM Complete Package Upgrade](#) on page 60

[TPVM Incremental Upgrade](#) on page 66

## TPVM Complete Package Upgrade

---

The TPVM upgrade activity upgrades the device's TPVM image, where a TPVM is installed and running an EFA instance which is managing the device. After the TPVM image is updated, an EFA instance gets reinstalled with the same EFA version as before and rejoined with the active EFA instance.



### Note

For more information about commands and supported parameters, see [Extreme Fabric Automation Command Reference, 3.1.0](#)

## Assumptions and Limitations

- EFA supports SLX-OS 20.3.2a and later. The TPVM upgrade has SLX-OS dependencies for the new SLX **tpvm upgrade** and **tpvm revert** commands, and the TPVM configurations being persisted in the SLX running-config.
- EFA does not support TPVM upgrade on a single-node TPVM deployment.
- TPVM can be deployed on multiple high-availability (HA) nodes, but you can upgrade TPVM only on the standby TPVM node.
- TPVM upgrade is allowed only on the EFA HA nodes that are managing the devices hosting the EFA HA instances.
- The EFA version is reinstalled and remains the same after the TPVM upgrade. An EFA version upgrade must not be performed during the TPVM upgrade. The EFA version upgrade must be done before or after the TPVM upgrade has been performed on both HA nodes as part of the normal EFA upgrade deployment.
- Only one device's TPVM can be upgraded at a time.

## TPVM Upgrade Workflow Dependencies

Before you start the TPVM upgrade, review the TPVM config and registration dependencies.

### *TPVM Configuration Persistence*

The TPVM running-config and operational data (including the TPVM image version and TPVM IP address) from the SLX device are persisted in the EFA DB. The following table describes the TPVM configurations that are persisted in the EFA DB.

TPVM Config	SLX Command Execution Stage	Type	Value	Description
auto-boot	Install only	Boolean	Exists or does not exist	Must always be enabled for an EFA TPVM.
password	Pre-start only	string	An encoded non-clear text password string. If does not exist then default is "password".	Extreme user password will not be clear-text in the running-config. Using encoded password string will still configure the SLX TPVM properly. If no password is set then default "password" is used.
interface management • ip • gw	Pre-start only	string	• "dhcp" or "IPv4 address" • "IPv4 address"	If configured as dhcp, the EFA must still fetch the existing management IP and Gateway IP to validate that the TPVM IP remains the same after the upgrade or node replacement.
interface insight • ipv4 • gw	Pre-start only	string	• "dhcp" or "IPv4 address" • "IPv4 address"	
hostname	Post-start	string	"hostname"	
timezone	Post-start	string	"timezone"	
dns server	Post-start	string	"FQDN" or "IPv4 address"	
ntp server	Post-start	string	"FQDN" or "IPv4 address"	

TPVM Config	SLX Command Execution Stage	Type	Value	Description
ldap <ul style="list-style-type: none"> <li>host</li> <li>port</li> <li>secure</li> <li>basedn</li> <li>rootdn</li> <li>password</li> </ul>	Post-start		<ul style="list-style-type: none"> <li>"FQDN" or "IPv4 address" or "IPv6 address"</li> <li>0-65535</li> <li>Exists or not exists</li> <li>Base domain name</li> <li>Root domain name</li> <li>Root domain name password</li> </ul>	
ldap ca-cert <ul style="list-style-type: none"> <li>protocol</li> <li>user</li> <li>password</li> <li>host</li> <li>directory</li> <li>filename</li> </ul>	Post-start	string	<ul style="list-style-type: none"> <li>"scp"</li> <li>Username</li> <li>Password</li> <li>"IPv4 address"</li> <li>Directory</li> <li>Filename</li> </ul>	The ca-cert for LDAP must be stored on the firmware-host and for EFA to support the node replacement. The ca-cert can also have IPv6 address.
trusted-peer <ul style="list-style-type: none"> <li>ip</li> <li>password</li> <li>sudo-user</li> </ul>	Post-start	string	<ul style="list-style-type: none"> <li>"IPv4 address"</li> <li>Sudo user password</li> <li>Sudo username</li> </ul>	Trusted-peer config typically exist on one of the EFA nodes. Push this config to the correct node after the upgrade.
deploy	Install	Boolean	<ul style="list-style-type: none"> <li>Exist</li> <li>Does not exist</li> </ul>	Installs, starts, and applies the configurations to the TPVM instance.

### *Device Registration Enhancements*

The complete existing TPVM config information will be persisted in the EFA DB when the device is registered, and only during the initial device registration. The **TPVM running-config** config information is read and stored during the deep device discovery phase so that user visible device registration times are not impacted. The TPVM config will only be fetched and stored during the initial device registration and not during subsequent device updates.

### *Timer-based TPVM Config Updates*

A timer is set up to poll daily data for any TPVM config changes and persist any new TPVM config changes for EFA HA peer managed devices.

## TPVM Upgrade Workflow

### Procedure

1. Perform validations on user input for the device IP, firmware-host, and tpvm-image.
  - a. The device IP is a registered device with the minimum supported SLX version and with associated TPVM configuration. It must be one of the EFA HA peers managing the device.
  - b. Ensure that the firmware-host is registered prior to the TPVM upgrade.
  - c. The `tpvm-image` is validated later during the SLX TPVM upgrade.
2. Read the current TPVM configuration and operational data (including TPVM version and IP address) from the device, and then perform the following validations. TPVM configuration will be pushed to the device in the node replacement case.
  - a. If TPVM is neither configured nor installed, then the TPVM configuration persisted in the EFA DB is pushed to the device and TPVM instance is installed. This operation supports the node replacement RMA case.
  - b. If TPVM configuration from the device differs from the persisted EFA configuration, then the device's configuration has priority, and the EFA DB is updated.

### TPVM Configuration Special Handling for RMA Node Replacement Case

- When the TPVM configuration interface management IP is set to "dhcp", the TPVM IP address must remain the same. This is due to a dependency on EFA deployment where the active node is expecting the peer node to be configured with a specific IP address. The peer node IP cannot be changed without restarting EFA HA cluster daemons on the active node.

### TPVM Configuration Special Handling for All Cases

- You must re-apply the trusted-peer configuration on the node where it was applied previously. It exists on only one of the nodes in the EFA HA cluster. The appropriate node is identified and the trusted peer configuration is pushed to the correct node during TPVM upgrade or node replacement.
3. Issue the appropriate SLX command to the device to upgrade or install the TPVM.
    - a. Issue the **tpvm upgrade** command to the device. The device stops and takes a snapshot to roll back in case of failure. The device downloads the TPVM image and upgrade the TPVM instance. The TPVM starts after the upgrade of the TPVM instance, and the existing TPVM configurations are programmed on the running TPVM instance.
    - b. In a node replacement case, the TPVM configuration is pushed to the device in the previous step. The **tpvm deploy** command is issued to the device. No TPVM snapshot is needed because the replacement switch is typically a new switch with no TPVM configured.
  4. Redeploy EFA on the upgraded or installed TPVM node with the current EFA version from the active node. Allow the redeployed peer node to rejoin the EFA HA cluster.
    - a. A new deployment strategy is driven from the current active EFA node to redeploy the current EFA version only on the newly upgraded or installed TPVM peer node. Because the active EFA instance remains operational during the deployment, the peer node can rejoin without disrupting the active EFA instance.

**Example****Example:**

```
(efa:extreme)extreme@node-1:~$ efa inventory device tpvm-upgrade execute --ip
10.20.48.162 --firmware-host 10.31.2.101 \
> --tpvm-image /buildsjc/sre_fusion/Nightly/tpvm/tpvm4.5.6/tpvm4.5.6_221103_2338/dist/
SWBD2900/tpvm_inc_upg-4.5.6-0.amd64.deb
TPVM Upgrade Execute [success]
Monitor TPVM upgrade execution progress using:

    efa inventory device tpvm-upgrade show --ip 10.20.48.162
    efa inventory device tpvm-upgrade show --execution-id a2c07243-bae0-46ea-aa2c-
e932e409d0bd

Please do not execute other commands on the device until process is completed

--- Time Elapsed: 145.914563ms ---
(efa:extreme)extreme@node-1:~$ while [ 1 ] ; do efa inventory device tpvm-upgrade show
--ip 10.20.48.162 ; sleep 120s ; done
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
+-----+
|IP      |Host|Model|Chassis | ASN | Role | Current TPVM| Target TPVM|Update  |
|Status  |Detailed|Failed|Upgrade |      |      | Start Time  | Last Update Time  |
|Address |Name|   Name   |      |      | Version
| Version |State |           |      |      | Status
| State  | Type |           |      |      |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
+-----+
|10.20   |AS2 |3012 |SLX9250 | 64512| Spine| 4.5.3       |      |In      |Device
Validation | None |      |Incremental |2022-11-05 23:52:29 | 2022-11-05 23:52:36 |
|.48.162 |    | -32C |      |      |
|      |    |      |Progress| Started      |
|      |Upgrade | -0700 PDT |      | -0700 PDT      |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
+-----+
TPVM Upgrade Show Details
--- Time Elapsed: 372.428607ms ---

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
+-----+
|IP      |Host|Model|Chassis |ASN   |Role |Current
|Target TPVM |Update |      |Status |      |
Detailed    |Failed|Upgrade |      |Start Time  | Last Update Time  |
|Address |Name|   Name   |      |      |TPM
Version |Version   |State |      |      |
Status  |State | Type |      |      |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
+-----+
|10.20   |AS2 |3012 |SLX9250 |64512 |Spine
|4.5.6   |4.5.6   |Completed|TPVM Upgrade
|Reboot Required |      |Incremental |2022-11-05 23:52:29 | 2022-11-06 00:01:11 |
|.48.162 |    | -32C |      |      |
|      |    |      |Workflow Completed|
for TPVM Instance|      |Upgrade | -0700 PDT |      | -0700 PDT      |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
+-----+
+-----+
```

## TPVM Upgrade Workflow States

TPVM Upgrade State	Next State	Case	Description
TPVM Upgrade Workflow Started	Device Validation	-Normal Upgrade -Node Replacement	Initial start state for the TPVM upgrade workflow.
Device Validation	<b>Success:</b> TPVM Config Validation <b>Failure:</b> TPVM Upgrade Workflow Finished	-Normal Upgrade -Node Replacement	Ensure that the provided device IP has associated TPVM configurations persisted in the EFA DB, and the device's TPVM IP is one of the EFA peer node IPs.
TPVM Config Validation	<b>-Normal Upgrade:</b> <b>Success:</b> TPVM Upgrade <b>Failure:</b> TPVM Upgrade Workflow Finished <b>-Node Replacement:</b> <b>Success:</b> TPVM Configuration <b>Failure:</b> TPVM Upgrade Workflow Finished	-Normal Upgrade -Node Replacement	Read TPVM config and operational data from the device and determine if it is a normal TPVM Upgrade or a Node Replacement case.  1. If TPVM config and operational data are present on the device and TPVM IP is one of the EFA peers, then it is a normal TPVM upgrade case.  2. If there is no TPVM config present on the device, then it is a node replacement case.  3. If TPVM config and operational data are present on the device and TPVM IP does not match one of the EFA peers, then validation for a normal TPVM upgrade was unsuccessful.  The <b>detailed status</b> column from the <b>tpvm-upgrade show</b> command output shows the nature of the issue and possible remedy.
TPVM Configuration	<b>Success:</b> TPVM Installation <b>Failure:</b> TPVM Upgrade Workflow Finished	-Node Replacement	Device's running-config is programmed using TPVM config data from EFA DB.

TPVM Upgrade State	Next State	Case	Description
TPVM Installation	<b>Success:</b> EFA Deploy Peer and Rejoin <b>Failure:</b> TPVM Upgrade Workflow Finished	-Node Replacement	TPVM install and start is invoked on the device.
TPVM Upgrade	<b>Success:</b> EFA Deploy Peer and Rejoin <b>Failure:</b> TPVM Revert	-Normal Upgrade	TPVM upgrade is invoked on the device.
TPVM Revert	<b>Success:</b> TPVM Upgrade Workflow Finished <b>Failure:</b> TPVM Upgrade Workflow Finished	-Normal Upgrade	On failure of “Upgrading TPVM” or “Deploying EFA for Rejoin”, the TPVM revert is invoked to roll-back the failed TPVM upgrade.
EFA Deploy Peer and Rejoin	<b>Success:</b> TPVM Upgrade Workflow Finished <b>Failure:</b> TPVM Revert	-Normal Upgrade -Node Replacement	On active EFA node, re-deploying of EFA on the peer node for rejoin is invoked.
TPVM Upgrade Workflow Finished	N/A	-Normal Upgrade -Node Replacement	End state for the TPVM upgrade workflow.

## TPVM Incremental Upgrade

### Before You Begin

EFA 3.0.0 and later releases support incremental upgrade. Ensure that the TPVM is running and the SLX version is 20.4.1 or later.

### About This Task

The TPVM incremental upgrade enables you to upgrade both active and standby TPVM nodes. The TPVM incremental upgrade is applicable for both multi-node and single node deployment with EFA running. This upgrade drastically reduces the upgrade time (around 3 minutes) compared to the full upgrade.

EFA automatically determines whether the TPVM upgrade is an incremental upgrade or a full upgrade based on the TPVM image name. If the image name contains `inc_upg`, then EFA determines that the upgrade as an incremental upgrade.



#### Note

- You can perform the incremental upgrade on either one of the TPVMs (active or standby) or on both the TPVMs at same time.
- The **efa inventory device tpvm-upgrade show** command displays the failed state in case of upgrade failures.

The **efa inventory device tpvm-upgrade show** command displays the TPVM information including version, IP address, device IP, TPVM hostname, and SLX version.

- The TPVM upgrade can be restarted in case of EFA restart or inventory service restart.
- TPVM upgrade execution detailed status (Detailed Status) shows whether a reboot is required for the incremental upgrade. If the incremental upgrade output shows that reboot is required for both active and standby nodes, ensure that you reboot the standby TPVM first, and then the active TPVM.

#### Procedure

1. Run the **efa inventory device tpvm-upgrade execute** command.

The CLI fetches either the active TPVM IP address or the standby TPVM IP address or both for a TPVM incremental upgrade. EFA does not allow more than one instance of TPVM incremental upgrade per device.

```
efa inventory device tpvm-upgrade execute --10.24.44.66
```

2. Run the **efa inventory device tpvm-upgrade show** command.

The **efa inventory device tpvm-upgrade show** command shows the device status for a tpvm-upgrade operation.

```
efa inventory device tpvm-upgrade show
```

The following are the examples of a TPVM upgrade show:

- TPVM Upgrade show with one TPVM IP address

```
efa inventory device tpvm-upgrade show --ip 10.24.80.58
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model
| Chassis Name | ASN | Role | Current TPVM Version | Target TPVM Version
| Update State | Status | Detailed Status | Start
Time | Last Update Time | Failed State |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX | 4001
| BR-SLX9640 | 4200000000 | Leaf | 4.5.0 | 4.5.0
| Completed | TPVM Upgrade Workflow Completed | None | 2022-05-19
08:31:11 -0700 PDT | 2022-05-19 08:33:52 -0700 PDT | TPVM Config Validation Failed |
```

```

+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+

```

Failed State: Shows the last failed state, if upgrade fails. On a successful run, this is a null string.

- TPVM Upgrade show with two TPVM IP addresses

```

efa inventory device tpvm-upgrade show --ip 10.24.80.58,10.24.80.56
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| IP Address | Host Name |
Model | Chassis Name | ASN | Role | Current TPVM Version
| Target TPVM Version | Update State | Status | Detailed Status
| Start Time | Last Update Time | Failed State |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX | 4001
| BR-SLX9640 | 42000000000 | Leaf | 4.5.0
| | In Progress | Device Validation Started | None
| 2022-05-26 13:41:14 -0700 PDT | 2022-05-26 13:41:19 -0700 PDT |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.24.80.56 | SLX |
4001 | BR-SLX9640 | 42000000000 | Leaf | 4.5.0
| | In Progress | Device Validation Started | None
| 2022-05-26 13:41:14 -0700 PDT | 2022-05-26 13:41:18 -0700 PDT |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+

```

- TPVM Upgrade show with two TPVM IP addresses and image path

```

efa inventory device tpvm-upgrade show --ip 10.24.80.58,10.24.80.56 --firmware-host
10.31.80.101 --tpvm-image <image_path>
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| IP Address | Host Name |
Model | Chassis Name | ASN | Role | Current TPVM Version
| Target TPVM Version | Update State | Status | Detailed Status
| Start Time | Last Update Time | Failed State |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX | 4001
| BR-SLX9640 | 42000000000 | Leaf | 4.5.0
| | In Progress | Device Validation Started | None
| 2022-05-26 13:41:14 -0700 PDT | 2022-05-26 13:41:19 -0700 PDT |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.24.80.56 | SLX |
4001 | BR-SLX9640 | 42000000000 | Leaf | 4.5.0

```

```

|                               | In Progress | Device Validation Started | None
| 2022-05-26 13:41:14 -0700 PDT | 2022-05-26 13:41:18 -0700 PDT |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+-----+-----+-----+

```

- TPVM Upgrade show with execution ID

```

efa inventory device tpvm-upgrade show --execution-id 670cb89e-d8d1-4213-
ac97-20403458627f
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
| IP Address | Host Name | Model
| Chassis Name | ASN | Role | Current TPVM Version
| Target TPVM Version | Update State | Status | Detailed Status
| Start Time | Last Update Time | Failed State |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.24.80.58 | SLX | 4001 |
BR-SLX9640 | 4200000000 | Leaf | 4.5.0
|                               | In Progress | TPVM Upgrade Started | None
| 2022-05-26 13:41:14 -0700 PDT | 2022-05-26 13:41:58 -0700 PDT |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.24.80.56 | SLX | 4001 |
BR-SLX9640 | 4200000000 | Leaf | 4.5.0
|                               | In Progress | TPVM Upgrade Started | None
| 2022-05-26 13:41:14 -0700 PDT | 2022-05-26 13:41:54 -0700 PDT |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
TPVM Upgrade Show Details
--- Time Elapsed: 124.961824ms ---

```

- TPVM Upgrade show with reboot required information

```

(efa:extreme)extreme@node180:~$ efa inventory device tpvm-upgrade show --ip
10.24.80.56,10.24.80.58
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
| IP Address | Host Name | Model | Chassis Name | ASN
| Role | Current TPVM Version | Target TPVM Version | Update State |
Status | Detailed Status | Failed State |
Upgrade Type | Start Time | Last Update Time |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.24.80.56 | SLX | 4001 | BR-SLX9640 | 4200000000
| Leaf | 4.5.2 | 4.5.2 | Completed | TPVM Upgrade
Workflow Completed | Reboot Required for TPVM Instance | Incremental
Upgrade | 2022-07-28 16:19:46 -0700 PDT | 2022-07-28 16:25:34 -0700 PDT |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+

```

```

+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX      | 4001 | BR-SLX9640 | 4200000000
| Leaf | 4.5.2      | 4.5.2      | Completed | TPVM Upgrade
Workflow Completed | None      |      |      | Incremental
Upgrade | 2022-07-28 16:19:46 -0700 PDT | 2022-07-28 16:26:24 -0700 PDT |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
TPVM Upgrade Show Details
--- Time Elapsed: 156.8739ms ---
(efa:extreme)extreme@node180:~$

```

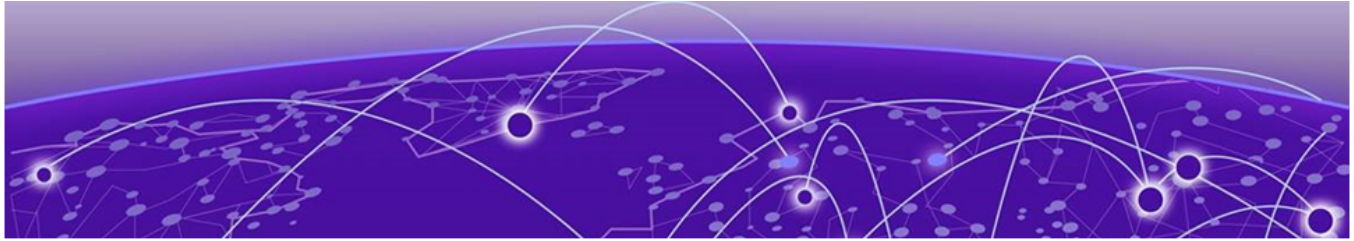
3. Run the **efa inventory device tpvm list** command.

The **efa inventory device tpvm list** command shows the TPVM list information such as TPVM IP address, SLX IP address, TPVM hostname, TPVM version, and SLX firmware version.

```

efa inventory device tpvm list
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Device IP Address | TPVM IP Address | TPVM Hostname |      SLX Firmware Version
| TPVM Version |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.56      | 10.24.80.180 | node180      | 20.4.2slxos20.4.2_220614_1000
| 4.5.0 |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58      | 10.24.80.181 | node181      | 20.4.1 | 4.5.0 |
|      |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+

```



# Upgrade Ubuntu

---

[Upgrade Ubuntu on the EFA Host - Single Node](#) on page 71

[Upgrade Ubuntu on the EFA Host - Multi-Node](#) on page 72

## Upgrade Ubuntu on the EFA Host - Single Node

---

Upgrade Ubuntu in single-node deployments.

### Before You Begin

Ensure that EFA is at release 3.1.0 or later. Upgrade if necessary. For more information, see [EFA Upgrade](#) on page 38.

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, 18.04, and 20.04 as described in [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12. You can upgrade from 16.04 to 18.04 and 18.04 to 20.04 while EFA is installed.



### Note

- This process is not supported for deployments of EFA on TPVM. For TPVM upgrade, see [Upgrading SLX-OS, TPVM, and EFA Together](#) on page 47, [TPVM Complete Package Upgrade](#) on page 60, and [TPVM Incremental Upgrade](#) on page 66.
- 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, and then upgrade all Ubuntu packages for standalone deployments.

```
sudo apt update && sudo apt update -y
```

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

- a. Run the following command to install the packages that have been kept back.

```
sudo apt-get upgrade <package-name>
```

- b. If any package upgrade requires reboot of server, reboot the server.

- c. Upgrade the node.

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

EFA is not operational while the upgrade is in progress.

- d. Reboot the system.
- e. Verify that EFA is operational. Perform the following (from the EFA command line) on both nodes.

Run **efa status** to verify that both nodes are up.

Run **sudo efactl status** to verify that all pods on the active node are in Running state.

Run **sudo efactl db-status** to verify that the MariaDB is active (running)

## Upgrade Ubuntu on the EFA Host - Multi-Node

Upgrade Ubuntu in multi-node deployments.

### Before You Begin

Ensure that EFA is at release 3.1.0 or later. Upgrade if necessary. For more information, see [EFA Upgrade](#) on page 38.

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, 18.04, and 20.04 as described in [Supported Platforms and Deployment Models for Fabric Manager](#) on page 12. You can upgrade from 16.04 to 18.04 and 18.04 to 20.04 while EFA is installed.



#### Note

- This process is not supported for deployments of EFA on TPVM. For TPVM upgrade, see [Upgrading SLX-OS, TPVM, and EFA Together](#) on page 47, [TPVM Complete Package Upgrade](#) on page 60, and [TPVM Incremental Upgrade](#) on page 66.
- 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, and then upgrade all Ubuntu packages for standalone deployments.

```
sudo apt update && sudo apt upgrade -y
```

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

- a. Upgrade one node in the cluster and reboot the system. Preferably, the standby node.

```
# 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.

Both the nodes must have the `sudo apt update` and `sudo apt upgrade -y` step present.

- b. Upgrade the second node using the same procedure as used on the first node of the cluster.
- c. Verify that the nodes are at the new version by executing **uname -a** and **cat /etc/os-release**.

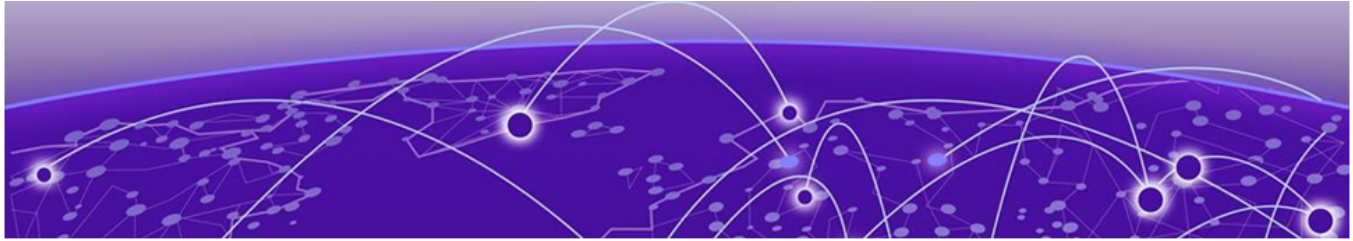
```
device$ uname -a
Linux xco-101-93 4.15.0-194-generic #205-Ubuntu SMP Fri Sep 16 19:49:27 UTC 2022
x86_64 x86_64 x86_64 GNU/Linux
device$ cat /etc/os-release
NAME="Ubuntu"
VERSION="18.04.6 LTS (Bionic Beaver)"
ID=ubuntu
ID_LIKE=debian
PRETTY_NAME="Ubuntu 18.04.6 LTS"
VERSION_ID="18.04"
HOME_URL="https://www.ubuntu.com/"
SUPPORT_URL="https://help.ubuntu.com/"
BUG_REPORT_URL="https://bugs.launchpad.net/ubuntu/"
PRIVACY_POLICY_URL="https://www.ubuntu.com/legal/terms-and-policies/privacy-policy"
VERSION_CODENAME=bionic
UBUNTU_CODENAME=bionic
```

- d. Verify that EFA is operational. Perform the following (from the EFA command line) on both nodes.

Run **efa status** to verify that both nodes are up.

Run **sudo efactl status** to verify that all pods on the active node are in Running state.

Run **sudo efactl db-status** to verify that the MariaDB is active (running)



# Redundant Management Network

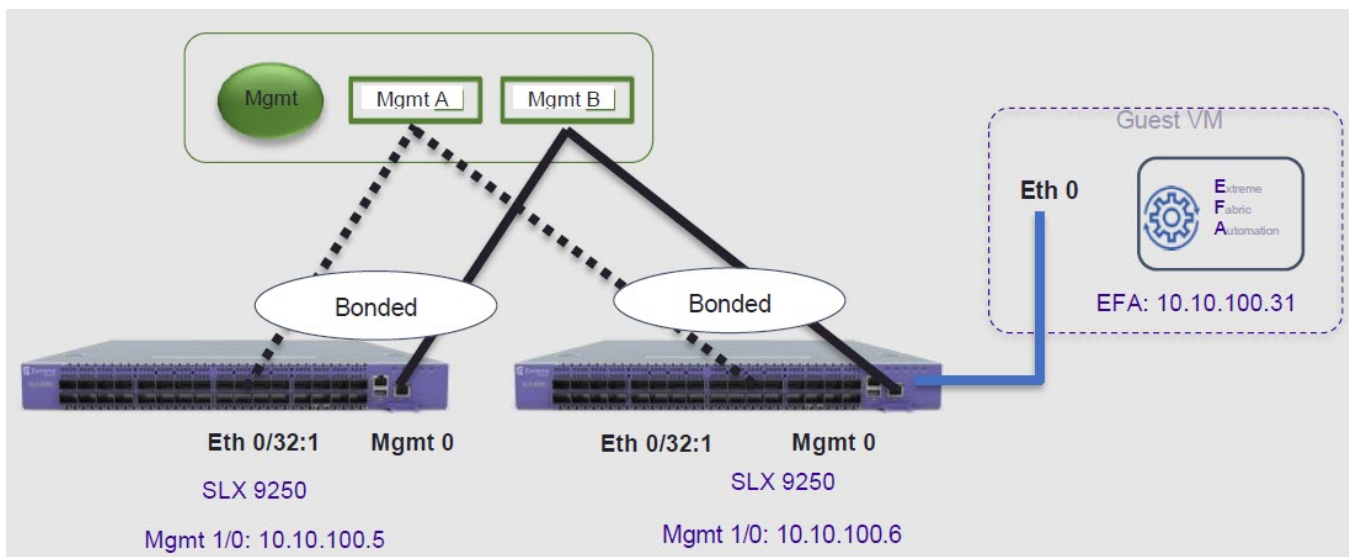
[Redundant Management Network Overview](#) on page 74

[Enable Redundant Management](#) on page 75

[Redundant Management Data Path](#) on page 77

## 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 support 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.
- 8720 has a dual management port. It does not need RME CLI.

## 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 74.

### 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  down
Ethernet 0/2    unassigned      default-vrf     administratively down  down
...
Ethernet 0/32 (M) unassigned      mgmt-vrf        administratively down  down
...
device# show interface ethernet 0/32

Ethernet 0/32 is admin down, line protocol is down (admin down)
```

```

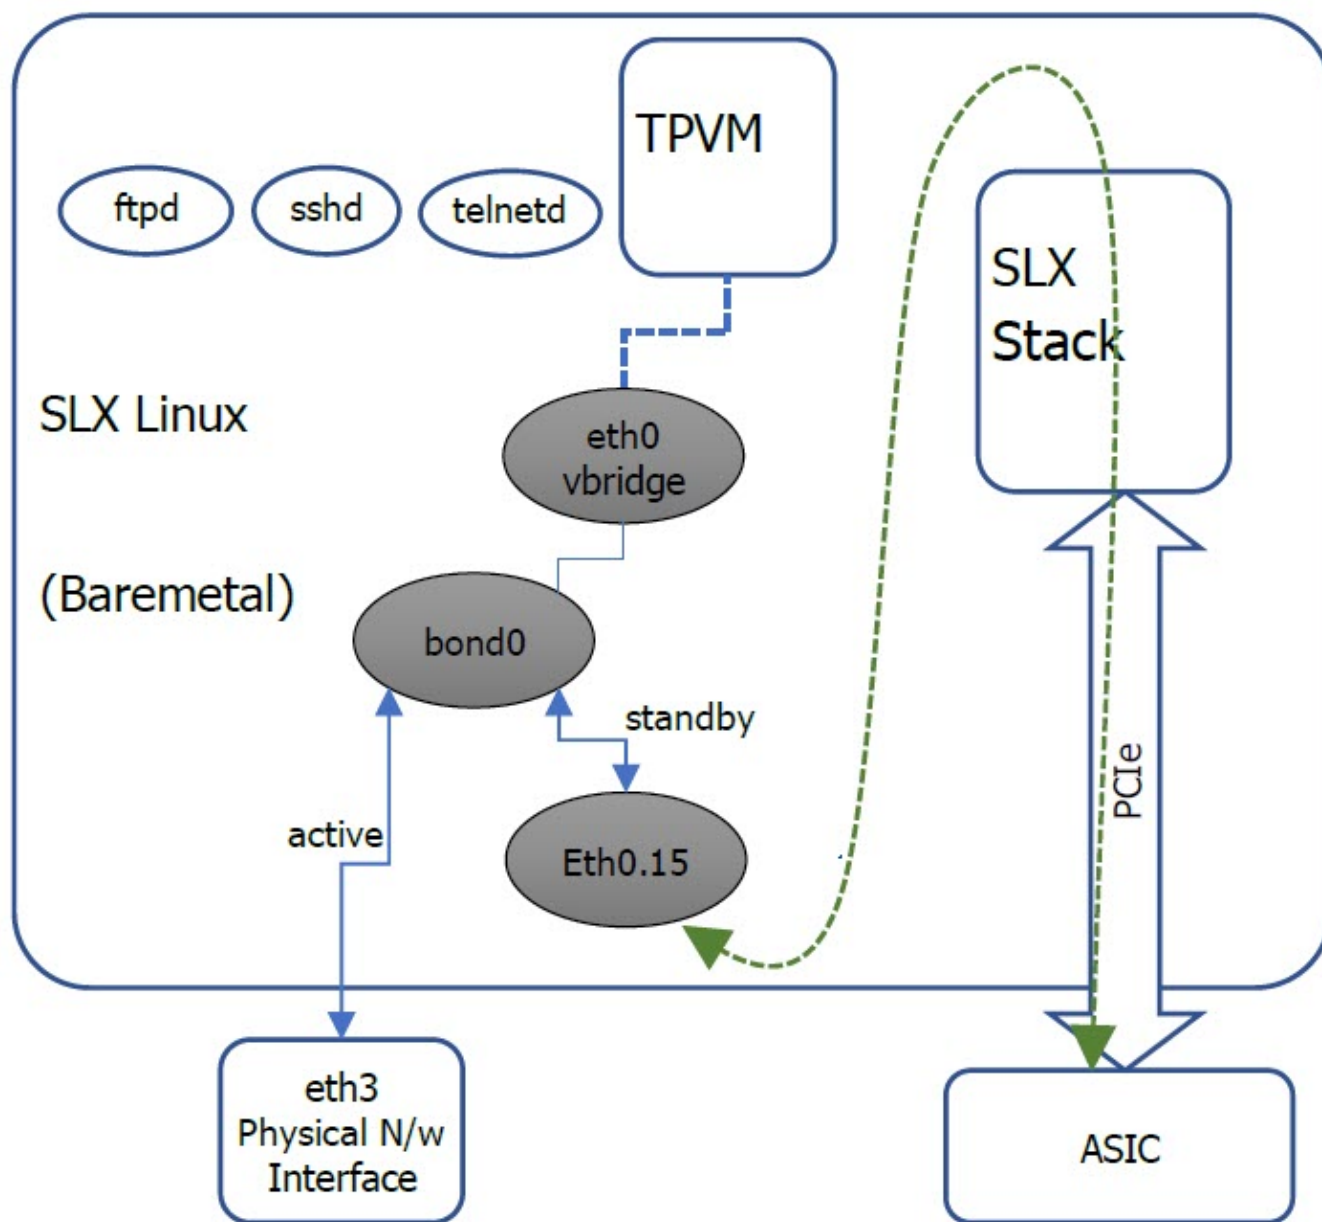
Redundant management mode is enabled
Hardware is Ethernet, address is 609c.9f5a.a35f
Current address is 609c.9f5a.a35f
Pluggable media not present
Description: Insight port
Interface index (ifindex) is 202350592 (0xc0fa000)
MTU 9216 bytes
Maximum Speed : 10G
LineSpeed Actual : Nil
LineSpeed Configured : 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

```

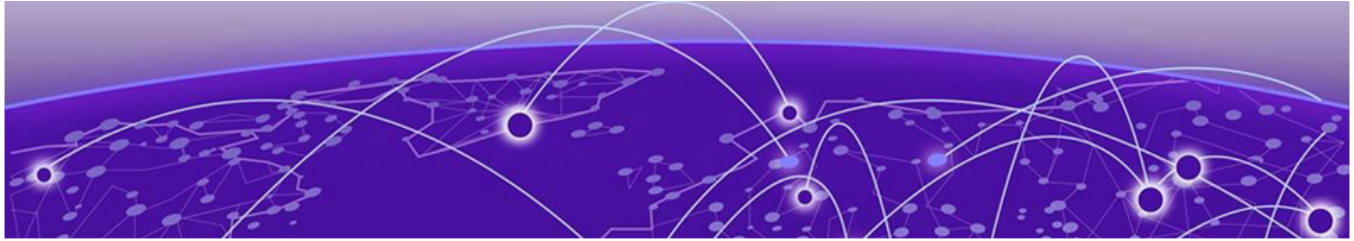
## Redundant Management Data Path

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 though 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**



# Flexible EFA Deployment for TPVM

---

[Flexible EFA Deployment Overview](#) on page 79

[Input Parameters on Single-Node Install or Upgrade](#) on page 80

[Input Parameters on Multi-Node Install or Upgrade](#) on page 81

[EFA Installer Improvements for TPVM-Based Deployment](#) on page 83

[No Graphics Mode](#) on page 84

[EFA Deployment with Rollback](#) on page 85

[Rollback the EFA Upgrade](#) on page 85

## Flexible EFA Deployment Overview

---

Flexible EFA deployment on TPVM removes the need of providing various setup parameters in an interactive way.



### Note

- Ensure that you have valid EFA packages and TPVM installation is in place in the correct path.
- Ensure that the SLX CLI versions are correct
- If you have changed a deployment parameter in the EFA 3.1.0 or above, but the SLX 20.4.2 or above is not updated to reflect the EFA changes, it does not appear in deployment. Use the noninteractive parameters along with the EFA deployment commands.
- When deployment is in progress, do not reboot or toggle management ports on the target devices for the TPVM and EFA deployment. Avoid CTRL+C on the installer window.
- As a best practice, do not use the IPv6 address that is converted from IPv4 address. For example, do not use the IPv6 address `::ffff:a14:f663` which is converted from the IPv4 address.

## SLX CLI

The SLX CLI contains all the required parameters for a TPVM deployment. Using the commands directly on SLX CLI, you can specify the parameters for the TPVM deployment in a single command without responding to prompts.

For example,

```
efa deploy non-interactive single-node package /efaboot/efa-3.1.0.tar.gz ...
```

## EFA Deployment

Using the graphical interface for deploying EFA is the default procedure. You can use the optional parameters to deploy EFA without using the graphical interface.

Following are the minimal required commands to start a deployment:

Deployment type	Commands
Single node install or upgrade	<code>efa deploy non-interactive single-node package /efaboot/efa-3.1.0.tar.gz</code>
Multi Node install or upgrade	<code>efa deploy non-interactive multi-node package /efaboot/efa-3.1.0.tar.gz peer-node 10.20.246.102 vip4 10.20.246.103</code>
Multi Node upgrade with replacement	<code>efa deploy non-interactive multi-node package /efaboot/efa-3.1.0.tar.gz replacement-ip 10.200.246.155</code>

## Listing EFA Packages

You can find the EFA package under the /efaboot directory on the SLX device.

As a package name is needed for single-node and multi-node installations, the **show efa packages** command can be used to show the currently available package.

## Input Parameters on Single-Node Install or Upgrade

This topic describes the input parameters for a single-node EFA deployment.



### Note

The deployment parameters, such as -m or -f have been replaced with ?.

## Minimum Required Commands

Following is the minimum required command for a single-node EFA deployment:

```
efa deploy non-interactive single-node package <package-name>
```

For example,

```
efa deploy non-interactive single-node package /efaboot/efa-3.1.0.tar.gz
```

## Management IP Networks

Use the **management-ip** command if you need additional management IP networks.

```
efa deploy non-interactive single-node package /efaboot/efa-3.1.0.tar.gz management-ip  
sub-interface-name sub200 sub-vlan-id 200 external-subnet 10.20.246.99/20
```

Do not use the **management-ip** command if it is not required. EFA supports only adding a single sub-interface.

When deploying the management IP, the command **sub-interface-name** requires the name of the sub-interface. **sub-vlan-id**, which is the VLAN ID, and **external-subnet**, which is the external subnet address in CIDR format.

## Input Parameters on Multi-Node Install or Upgrade

This topic describes the input parameters for a multi-node EFA deployment.

### Deployment Type

Specify the deployment type using the **multi-node** command after the deployment. Use the **peer-node** command followed by the peer node IP address for the installation.

Use the **vip4** command followed by the virtual IP address to provide the virtual IPv4 address for the installation.

```
efa deploy non-interactive multi-node package <package-name> peer-node <peer node ip> vip4 <virtual ip-address>
```

For example,

```
efa deploy non-interactive multi-node package /efaboot/efa-3.1.0.tar.gz peer-node 10.20.246.102 vip4 10.20.246.103
```

### Virtual IPv6

Use the **vip6** command followed by the virtual IP address to provide the virtual IPv6 address for the installation.

As a best practice, do not use the virtual IPv6 address that is converted from IPv4 address.

For example,

```
efa deploy non-interactive multi-node package /efaboot/efa-3.1.0.tar.gz peer-node 10.20.246.102 vip4 10.20.246.103 vip6 fd00::56:45
```

### Management IP Networks

Use the **management-ip** command if you need additional management IP networks. Do not use the command if it is not required. EFA supports adding only a single sub-interface.

```
efa deploy non-interactive multi-node package peer-node <peer ip> vip4 <virtual ip> management-ip sub-interface-name <sub interface name> sub-vlan-id <sub vlan id> external-subnet <virtual ip with subnet> external-v6-subnet <virtual ip with ipv6 subnet>
```

When deploying the management IP, some commands are **sub-interface-name**, which requires the name of the sub-interface. **sub-vlan-id**, which is the VLAN ID, and **external-subnet**, which is the external subnet address in CIDR format.

Example:

```
efa deploy non-interactive single-node package /efaboot/efa-3.1.0-410.tar.gz management-  
ip sub-interface-name sub200 sub-vlan-id 200 external-subnet 10.10.10.1/24 external-v6-  
subnet 2001::1/64
```

## Build Upgrade and Replacement

In a TPVM deployment, following are the available options for a multi-node build upgrade and replacement:

- With node replacement: Use the **replacement-ip** command, followed by the replacement peer node IP address.
- Without node replacement: No action is needed as the default is to deploy with no node replacement.

For example,

```
efa deploy non-interactive multi-node package /efaboot/efa-3.1.0.tar.gz replacement-ip  
10.200.246.155
```

## Single CLI for HA Ping-target Parameter

During installation, EFA connectivity to a gateway is mandatory. It relies on an explicit PING or ICMP request. If the gateway connectivity fails, the installation will fail.

In the scenario where the gateway address is obtained via the standard VRRP, then the PING or ICMP request fails. This results in EFA installation failure.

The new parameter `ping-target` is provided in single CLI so that instead of pinging the default gateway, it pings from both the nodes to the `ping-target` IP addresses. If the `ping-target` is not reachable from any of the nodes, then the installation or upgrade fails.

A maximum of two IP addresses can be given as input to `ping-target`. The IP addresses can be IPv4 or IPv6.

If you have installed EFA with `ping-target`, and while upgrading the `ping-target` argument is not given, it will retain old values of `ping-target` and pings to the old `ping-target` IP addresses. If you do not want to ping to the old `ping-target` IP addresses, the “clear” option can be given to `ping-target` so that it clears the old values and ping the default gateway.

If you have installed without `ping-target`, it will ping the default-gateway. If you have given `ping-target` during upgrade, it will ping the new `ping-target` IP addresses, otherwise it will ping the default-gateway.

*Single CLI Commands*

Following table provides the list of single CLI commands supported on SLX:

With or without sub-interface	Commands
Without sub-interface	<code>#efa deploy non-interactive multi-node package &lt;packagename&gt; peer-node &lt;ipaddress&gt; vip4 &lt;ip-address&gt; ping-target &lt;clear &lt;ip address1&gt;,[ip address2]&gt;</code>
With sub-interface	<code>efa deploy non-interactive multi-node package &lt;packagename&gt; peer-node &lt;ipaddress&gt; vip4 &lt;ip-address&gt; ping-target &lt;ip address&gt;,[ip address2] management-ip sub-interface-name &lt;sub-intf-name&gt; sub-vlan-id &lt;vlanid&gt; external-subnet &lt;ip-address&gt; external-v6-subnet &lt;virtual ip with ipv6 subnet&gt;</code>  <b>Note:</b> The <code>external-v6-subnet</code> parameter is not mandatory.

## EFA Installer Improvements for TPVM-Based Deployment

When you install EFA (fresh install or upgrade) on a TPVM, the following services are now disabled by default. REST API calls made to these services return failure.

The updated installer optimizes the TPVM installation to disable the microservices by default from the TPVM.

Using the EFA CLI, you can enable or disable the following 5 microservices:

- OpenStack
- Hyper-V (Microsoft SCVMM)
- vCenter (vMWare)



### Note

By default, these services are enabled on Server-based installations.

All the services will be enabled if you install EFA in server mode.

### Related Topics

[Upgrades and Service State](#) on page 83

[Enable or Disable Services](#) on page 84

## Upgrades and Service State

Services that are disabled prior to upgrade remain disabled after the upgrade. However, the software images for the services are upgraded so that if a disabled service is enabled, it will be consistent with the rest of the EFA installation.

When you disable a service, the corresponding process also gets stopped. When you re-enable a service, the process gets started. The behavior of a microservice post-

enablement is determined by what happens when the process starts. For example, when re-enabled, OpenStack polls Neutron for fresh data if that is its current behavior on startup.

## Enable or Disable Services

EFA provides CLI commands to enable or disable following services for TPVM or Server-based deployments in EFA:

- openstack (OpenStack)
- scvmm (Microsoft HyperV SCVMM)
- vcenter (VMware vCenter)
- notification (NotificationService)
- snmp (SNMP service)

```
efa system service
Microservice-specific commands

Usage:
efa system service [command]

Available Commands:
  enable      Enable and start a service
  disable     Disable and stop a service

Flags:
  -h, --help  help for service

Use "efa system service [command] --help" for more information about a command.
```

Example:

```
efa system service enable -name=openstack
(efa:extreme)extreme@tpvm:~$ efa system service enable --name=foo
Error : Please provide a valid service name: notification, openstack, scvmm, snmp, and
vcenter
efa system service disable --name openstack
```

## No Graphics Mode

Use the **efa deploy nographics** command.

The **efa deploy nographics** command does not display any graphics progress bar. The SLX will only display text about what is getting installed. When you use the "nographics" option, the system prompts you for all the required inputs.

You can also continue using the existing **efa deploy --graphics no** command.



### Note

You cannot use the **efa deploy nographics** and **efa deploy --graphics no** commands together.

## EFA Deployment with Rollback

Ensure that the EFA tar is available on the `/efaboot` partition of the SLX device.

Following are the various options for EFA deployment with rollback:

1. Deploy EFA with rollback on an SLX TPVM in a single-node deployment without `slx-peer` parameters.

```
efa deploy --nographics with-rollback
```

2. Deploy EFA with rollback on an SLX TPVM in a multi-node deployment with `slx-peer` parameters.

```
efa deploy nographics with-rollback slx-peer-ip 10.20.246.2 slx-peer-user admin
slxpeer-
password pass
```

3. Deploy EFA with rollback on an SLX TPVM in a single-node deployment.

```
efa deploy non-interactive with-rollback single-node package /efaboot/
efa-3.1.0.tar.gz ...
```

4. Deploy EFA with rollback on an SLX TPVM in a multi-node deployment.

```
efa deploy non-interactive with-rollback slx-peer-ip 10.20.246.2 slx-peer-user admin
slx-peer-password pass multi-node package /efaboot/efa-3.1.0.tar.gz ...
```

## Rollback the EFA Upgrade

### About This Task

Rollback the EFA upgrade when there is a upgrade failure. After the rollback, EFA operates in the previous state by canceling the upgrade.



#### Note

- Ensure that the minimum EFA version is 3.1.0 and above and the minimum SLX version is 20.4.2 and above.
- Ensure that the minimum disk space is 2 GB on each SLX TPVM partition. Log in to SLX as a root user and run the following command:

```
[root@SLX-1]# df -h
Filesystem      Size  Used Avail Use% Mounted on
/dev/root        16G   4.6G   11G   32% /
```

### Procedure

Deploy EFA with rollback.

```
efa deploy nographics with-rollback slx-peer-ip <ip-address> slx-peer-user <user name>
slx-peer-password <password>
```

When you use the `with-rollback` option, the following parameters are required:

- `slx-peer-ip` - SLX IP, which hosts peer TPVM
- `slx-peer-user` - SLX user, which hosts peer TPVM
- `slx-peer-password` - SLX password, which hosts peer TPVM

## Example

Following is an example of a deployment with rollback option:

```

efa deploy nographics with-rollback slx-peer-ip 10.20.54.62 slx-peer-user admin slx-peer-
password password

Step 1: Get IP Address assigned to TPVM to deploy EFA 10.20.63.128.
Step 2: Checking for EFA packages in /efaboot directory
1. /efaboot/efa-2.7.2-32.tar.gz
2. /efaboot/efa-2.7.2-31.tar.gz
Enter option: 1
*****
*                ! ! ! WARNING ! ! !                *
* Proceeding with Extreme Fabric Automation deployment *
*       1. Do not reboot device(s) or TPVM(s)         *
*       2. Do not toggle management port on device(s) or TPVM(s) *
*       3. Avoid CTRL+C on the installer window        *
*****
Ensuring TPVM 10.20.63.129 is deployed on remote SLX 10.20.54.62... done.
Ensuring EFA supports this rollback procedure... done.
Putting EFA into quiescent state..... done.
Stopping database on standby TPVM.... done.
Stopping database on active TPVM.... done.
Taking snapshot of active TPVM... done.
Taking snapshot of standby TPVM... done.
Starting database on active TPVM..... done.
Starting database on standby TPVM.... done.
Waiting for EFA to start..... done.
Completed EFA install preparation.
Copying EFA package efa-2.7.2-32.tar.gz to TPVM 10.20.63.128... done.
Extracting EFA package efa-2.7.2-32.tar.gz on TPVM 10.20.63.128... done.
Starting EFA installer.
Step 3: Checking for EFA Stack...
Previous Stack found
Are you sure you want to re-deploy EFA? (yes/no)
no
Do you wish to restart the install? (yes/no)
no
Preserving EFA supportsave... done.
Powering off standby TPVM... done.
Powering off active TPVM... done.
Reverting to saved EFA state on active... done.
Waiting for TPVM to boot on active..... done.
Reverting to saved EFA state on standby... done.
Copying EFA supportsave back to TPVM... done.
Waiting for EFA to start..... done.
Completed EFA revert procedure.
EFA revert succeeded.
EFA deployment discontinued or failed.
Spinel#

```

Following is an example of an upgrade failure with rollback option:

```

Spinel# efa deploy non-interactive with-rollback slx-peer-ip 10.20.54.62 slx-peer-user
admin slx-peer-password password multi-node package /efaboot/efa-2.7.2-32.tar.gz
Initializing...
*****
*                ! ! ! WARNING ! ! !                *
* Proceeding with Extreme Fabric Automation deployment *
*       1. Do not reboot device(s) or TPVM(s)         *
*       2. Do not toggle management port on device(s) or TPVM(s) *
*       3. Avoid CTRL+C on the installer window        *
*****
Ensuring TPVM 10.20.63.129 is deployed on remote SLX 10.20.54.62... done.
Ensuring EFA supports this rollback procedure... done.

```

```

Putting EFA into quiescent state..... done.
Stopping database on standby TPVM.... done.
Stopping database on active TPVM.... done.
Taking snapshot of active TPVM... done.
Taking snapshot of standby TPVM... done.
Starting database on active TPVM..... done.
Starting database on standby TPVM.... done.
Waiting for EFA to start..... done.
Completed EFA install preparation.
Copying EFA package efa-2.7.2-32.tar.gz to TPVM 10.20.63.128... done.
Extracting EFA package efa-2.7.2-32.tar.gz on TPVM 10.20.63.128... done.
Starting EFA installer.
Checking for EFA Stack...
Deployment mode is upgrade
Verifying connectivity to 10.20.63.129...
You have entered:
- to redeploy EFA at version 2.7.2 build 32
- with peer 10.20.63.129
- and VIP 10.20.63.127
- with additional HA health ping check IP(s) 10.20.54.63,10.20.54.64
Making backup
Removing legacy EFA installation
Stopping EFA services
Undeploying EFA application...
Undeploying ecosystem services
Undeploying core services
Removed current application deployment successfully.
Removing EFA container images
Removing container images on 10.20.63.128 10.20.63.129...
Removing EFA OS services
Removing k3s container orchestration
Removing database
Removing cluster filesystem
Removing keepalived for cluster virtual IP
Removing database sync tools
Removing EFA services and utilities
Proceeding with new EFA installation
Verifying system requirements
Verifying system requirements on all nodes
Ensuring networking components are ready
Installing software dependencies
Started installing helm
Installing database migrate client
Installing glusterfs filesystem software
Installing glusterfs 7.2...
GlusterFS Installation Success
Creating clustered filesystem
Configuring glusterfs volumes
Mounting efa volumes for replication to start
Mounting gluster units
Done with mounting of glusterfs efa volumes on nodes
Completed configuring glusterfs
Setting up EFA database
Installing and configuring mariadb server for HA...
Installing perl dependency for database use
Installing database client
Installing database server
Installing mariadb 10.4 server...
MariaDB 10.4 Installation Success
Configuring database server
Failed.
Failed.
Please wait while supportsave runs...
Supportsave complete - /apps/efa_logs/efa_2022-04-26T11-26-40.185.logs.zip

```

```
46.016172224s
Preserving EFA supportsave... done.
Powering off standby TPVM... done.
Powering off active TPVM... done.
Reverting to saved EFA state on active... done.
Waiting for TPVM to boot on active..... done.
Reverting to saved EFA state on standby... done.
Copying EFA supportsave back to TPVM... done.
Waiting for EFA to start..... done.
Completed EFA revert procedure.
EFA revert succeeded.
EFA deployment discontinued or failed.
Spinel#
```