



ExtremeCloud Orchestrator Deployment Guide

3.2.0

9037728-00 RevAB
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

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>



Table of Contents

Preface.....	6
Text Conventions.....	6
Documentation and Training.....	7
Help and Support.....	8
Subscribe to Product Announcements.....	8
Send Feedback.....	9
About this Document.....	10
What's New in this Document.....	10
XCO Deployment Preparation.....	11
Supported Platforms and Deployment Models for Fabric Skill.....	12
Supported Platforms and Deployment Models for Visibility Skill.....	15
XCO Requirements.....	16
General requirements.....	16
High-availability requirements.....	17
XCO Port Requirements.....	19
XCO Installation on TPVM.....	20
Overview.....	20
Requirements.....	21
Supported deployments.....	21
XCO Installation for Single-Node Deployments.....	22
XCO Installer Improvements for Server-Based Deployment.....	22
Install XCO in a Single-Node Deployment.....	23
Install XCO on TPVM in a Single-Node Deployment	24
OVA Root Login Changes.....	26
Deploy the OVA.....	27
Configure OVA using Postboot Menu.....	28
XCO Installation for Multi-Node Deployments.....	31
XCO Deployment for High Availability	31
Overview.....	31
Install XCO in a Multi-Node Deployment.....	34
Install XCO on TPVM in a Multi-Node Deployment	36
XCO Upgrade.....	39
Upgrade XCO in a Single-node Deployment.....	39
Upgrade XCO on TPVM in a Single-node Deployment.....	40
Upgrade XCO from a Single-Node to a Multi-Node Deployment.....	42
Upgrade XCO in a Multi-Node Deployment.....	44
Upgrade XCO on TPVM in a Multi-Node Deployment.....	45
Upgrade XCO on an OVA.....	47
Upgrading SLX-OS, TPVM, and XCO Together.....	48

Requirement for SCP connections.....	48
Upgrade XCO, SLX-OS, and TPVM Method 1.....	48
Upgrade XCO, SLX-OS, and TPVM Method 2.....	50
Recover from an Upgrade Failure.....	53
Rollback.....	54
Maintain TPVM Versions After a Rollback in a Multi-Node Deployment.....	54
Rollback SLX.....	55
Rollback XCO.....	55
Node Replacement.....	57
Replace a Node in a Multi-node Deployment.....	57
Replace a Node in a Multi-node TPVM Deployment.....	58
XCO Uninstallation.....	60
Uninstall XCO in a Single-Node or Multi-Node Deployment.....	60
Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment.....	60
TPVM Upgrade from XCO.....	62
TPVM Complete Package Upgrade.....	62
Assumptions and Limitations.....	62
TPVM Upgrade Workflow Dependencies.....	63
TPVM Upgrade Workflow.....	65
TPVM Upgrade Workflow States.....	67
TPVM Incremental Upgrade.....	68
TPVM Incremental Upgrade using Auto-Reboot.....	72
Upgrade Ubuntu.....	74
Upgrade Ubuntu on the XCO Host - Single Node or Multi Node.....	74
Upgrade Ubuntu on the XCO Host - Multi-Node.....	76
Redundant Management Network.....	78
Redundant Management Network Overview.....	78
Linux Bonding.....	78
Supported Ports.....	79
No Redundancy Period.....	79
Standby Port Rate Throughput.....	79
Enable Redundant Management.....	79
Redundant Management Data Path.....	81
Flexible XCO Deployment for TPVM.....	83
Flexible XCO Deployment for TPVM.....	83
SLX CLI.....	83
XCO Deployment.....	84
Listing XCO Packages.....	84
Input Parameters on Single-Node Install or Upgrade.....	84
Minimum Required Commands.....	84
Management IP Networks.....	85
Input Parameters on Multi-Node Install or Upgrade.....	85
Deployment Type.....	85
Virtual IPv6.....	85
Management IP Networks.....	86
Build Upgrade and Replacement.....	86
Single CLI for HA Ping-target Parameter.....	86

XCO Installer Improvements for TPVM-Based Deployment.....	87
Upgrades and Service State.....	88
Enable or Disable Services.....	88
No Graphics Mode.....	88
XCO Deployment with Rollback.....	89
Rollback the XCO Upgrade.....	89



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

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic text</i>	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member [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/.

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

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.

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 3.2.0 software release.

Table 4: Summary of changes

Description	Topic
New topic describes using the auto-reboot options for TPVM incremental upgrade	TPVM Incremental Upgrade using Auto-Reboot on page 72
New topic describes the changes of OVA root user	OVA Root Login Changes on page 26



XCO Deployment Preparation

- [Supported Platforms and Deployment Models for Fabric Skill](#) on page 12
- [Supported Platforms and Deployment Models for Visibility Skill](#) on page 15
- [XCO Requirements](#) on page 16
- [XCO Port Requirements](#) on page 19
- [XCO Installation on TPVM](#) on page 20

Supported Platforms and Deployment Models for Fabric Skill

Support includes Server, Open Virtual Appliance (OVA), and TPVM deployment models, supported TPVM versions, supported SLX-OS software versions, and supported SLX devices.



Note

As a best practice, refer to the following Extreme validated support matrices for support platforms and deployment models information.

Table 5: Server Deployment Models

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Server Version	Virtual Machine
2.7.x, 3.0.0	More than 24	Yes	16.04, 18.04	<ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
3.1.x	More than 24	Yes	16.04, 18.04, and 20.04	<ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
3.2.0	More than 24	Yes	18.04 and 20.04	<ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB • RAM: 8 GB

Table 6: OVA Deployment Models

XCO Version	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Virtual Machine
2.7.x, 3.0.0, and 3.1.x	More than 24	Yes	18.04	<ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
3.2.0	More than 24	Yes	18.04	<ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB • RAM: 8 GB

Table 7: TPVM Deployment Models

XCO Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
2.7.x	<ul style="list-style-type: none"> • SLX 9150 • SLX 9250 • SLX 9740 	Up to 24	Yes	18.04	20.4.1

Table 7: TPVM Deployment Models (continued)

XCO Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
	<ul style="list-style-type: none"> Extreme 8520 Extreme 8720 				
3.0.x	<ul style="list-style-type: none"> SLX 9150 SLX 9250 SLX 9740 Extreme 8520 Extreme 8720 	Up to 24	Yes	18.04	20.4.2
3.1.x	<ul style="list-style-type: none"> SLX 9150 SLX 9250 SLX 9740 Extreme 8520 Extreme 8720 Extreme 8820 (20.4.3 onwards only) 	Up to 24	Yes	18.04	20.4.2
3.2.0	<ul style="list-style-type: none"> SLX 9150 SLX 9250 SLX 9740 Extreme 8520 Extreme 8720 Extreme 8820 (20.4.3 onwards only) 	Up to 24	Yes	18.04	20.4.3

Table 8: TPVM Software Support

XCO Version	TPVM Version	SLX-OS Version
2.5.4	4.3.0	20.3.2d
2.5.5		
2.6.0	4.4.0	20.3.4/4a
2.6.1		
2.7.0	4.5.0	20.4.1
2.7.2	4.5.1	20.4.1b
3.0.0	4.5.3	20.4.2
3.1.0	4.5.6	20.4.2a

Table 8: TPVM Software Support (continued)

XCO Version	TPVM Version	SLX-OS Version
3.1.1	4.5.8	20.4.3
3.2.0	4.5.10	20.4.3a

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	✓			✓	
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	✓			✓	✓
Extreme 8820	20.4.3		✓	✓	✓	✓

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

XCO or EFA Version	Neutron Version	SLX-OS Version
2.5.4, 2.5.5	3.1.1-04	20.3.2d

Supported Platforms and Deployment Models for Visibility Skill

Support includes Server, OVA, and supported devices and software.



Note

- Upgrade from XVM (Extreme Visibility Manager) to XCO is not supported.
- XCO supports only a fixed set of special characters for names. Any additional characters configured in MLX or SLX are reconciled in XCO and can be edited or deleted. Any configuration name must start with an alphanumeric character and can contain " a-z A-Z 0-9 _ -"

Table 11: Ubuntu Server Version

XCO Version	Ubuntu Version	Virtual Machine
3.1.x	18.04 and 20.04	Minimum: <ul style="list-style-type: none"> • CPU: 4 cores • Storage: 128 GB • RAM: 8 GB Recommended: <ul style="list-style-type: none"> • CPU: 16 cores • Storage: 200 GB • RAM: 32 GB
3.2.0	18.04 and 20.04	Minimum: <ul style="list-style-type: none"> • CPU: 4 cores • Storage: 128 GB • RAM: 8 GB Recommended: <ul style="list-style-type: none"> • CPU: 16 cores • Storage: 200 GB • RAM: 32 GB

Table 12: OVA Deployment Models

XCO Version	Ubuntu Version	Virtual Machine
3.1.x	18.04	Minimum: <ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB • RAM: 8 GB
3.2.0	18.04	Minimum: <ul style="list-style-type: none"> • CPU: 4 cores • Storage: 64 GB

Table 12: OVA Deployment Models (continued)

XCO Version	Ubuntu Version	Virtual Machine
		<ul style="list-style-type: none"> RAM: 8 GB

Table 13: Supported Devices and Software

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

XCO 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 XCO 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.
 - XCO 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 XCO 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 XCO manages must use NTP to ensure easy audit trails and logging from XCO.
- **NTP:** The XCO 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 XCO 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 XCO 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 Skill](#) 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 XCO, 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 XCO, 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.

XCO Port Requirements

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

Table 14: General port requirements

Port	Service
80	XCO HTTP requests
162	XCO SNMP notifications
443	XCO HTTPs requests
514	Syslog service
3306	MariaDB port

Table 14: General port requirements (continued)

Port	Service
6443	K3S
6514	Secure syslog service
8078	Monitoring service
8079	Host authentication
10010	Containerd
30085	OpenStack 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

XCO Installation on TPVM

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

Overview

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

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

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

You can find the XCO package tar.gz file under the `/efaboot` directory of the SLX device. This applies to a fresh install or upgrade of XCO. For an incremental XCO image upgrade, you can copy the XCO 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 [XCO 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 XCO on TPVM in a single-node deployment or in a multi-node deployment for high availability. For more information, see [Install XCO on TPVM in a Single-Node Deployment](#) on page 24 and [Install XCO on TPVM in a Multi-Node Deployment](#) on page 36.



XCO Installation for Single-Node Deployments

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

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

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

[OVA Root Login Changes](#) on page 26

[Deploy the OVA](#) on page 27

[Configure OVA using Postboot Menu](#) on page 28

XCO Installer Improvements for Server-Based Deployment

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



Note

- During failure, the installer automatically collects supportsave and unwinds the partial installation.
- For any error, check the installer logs in the <Logs directory>/installer directory. For more details, see "Logging and Log Files" in [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).

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 XCO 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 [XCO Requirements](#) on page 16.

About This Task

To install XCO, 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 XCO 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 XCO 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 XCO is deployed.

9. Verify the installation.
 - a. From the XCO 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 XCO on TPVM in a Single-Node Deployment

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

Before You Begin

The XCO tar must be available on the `/efaboot` partition of the SLX device.

About This Task

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

Procedure

1. Verify that the TPVM is set up for an XCO 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 Skill](#) 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 XCO 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 XCO 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 XCO 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 XCO 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 XCO 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 XCO 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.

OVA Root Login Changes

XCO 3.2.0 has no root login for OVA.

Before XCO 3.2.0, a deployment user was the same as the initial installer, and was carried forward even if there was a different user during an upgrade. In XCO 3.2.0, the current user who is installing the system as a deployment user is added with “SystemAdmin” Role.

The name of the user is Ubuntu. The default password for this user is “password”.It is a single user. After boot configuration, you can change the password of this new user.

```

=====
Extreme Networks, Inc. - Extreme Fabric Automation - Modify Settings
=====
If you need to make a change, enter the appropriate number from the choices listed below.
If settings are not done and exited, the default will be applied.
=====
=

1. Set the ubuntu user password
2. Set network settings
3. Apply and Exit Menu

```

```

Enter selection :
=====
Extreme Networks, Inc. - Extreme Fabric Automation - Ubuntu Password Configuration
=====
The ubuntu password is currently set for this appliance.
=====
=
Would you like to set a ubuntu password (y/n) [y]?

```

Deploy the OVA

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

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 Skill](#) 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.

- Use the following credentials and move to [Configure OVA using Postboot Menu](#) on page 28.

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

- 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:
```

- 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
=====
```

```
Would 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]?
```



XCO Installation for Multi-Node Deployments

[XCO Deployment for High Availability](#) on page 31

[Install XCO in a Multi-Node Deployment](#) on page 34

[Install XCO on TPVM in a Multi-Node Deployment](#) on page 36

[XCO Deployment for High Availability](#)

Overview

A high-availability cluster is a group of servers that provide continuous up time or 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, XCO is deployed on the Server. The two XCO instances are clustered and configured with one IP address ensuring that clients need to reach only one endpoint. All XCO services are installed on each node. The node on which XCO is installed is the active node and processes all requests. The other node is the standby node that processes all the requests when the active node fails.

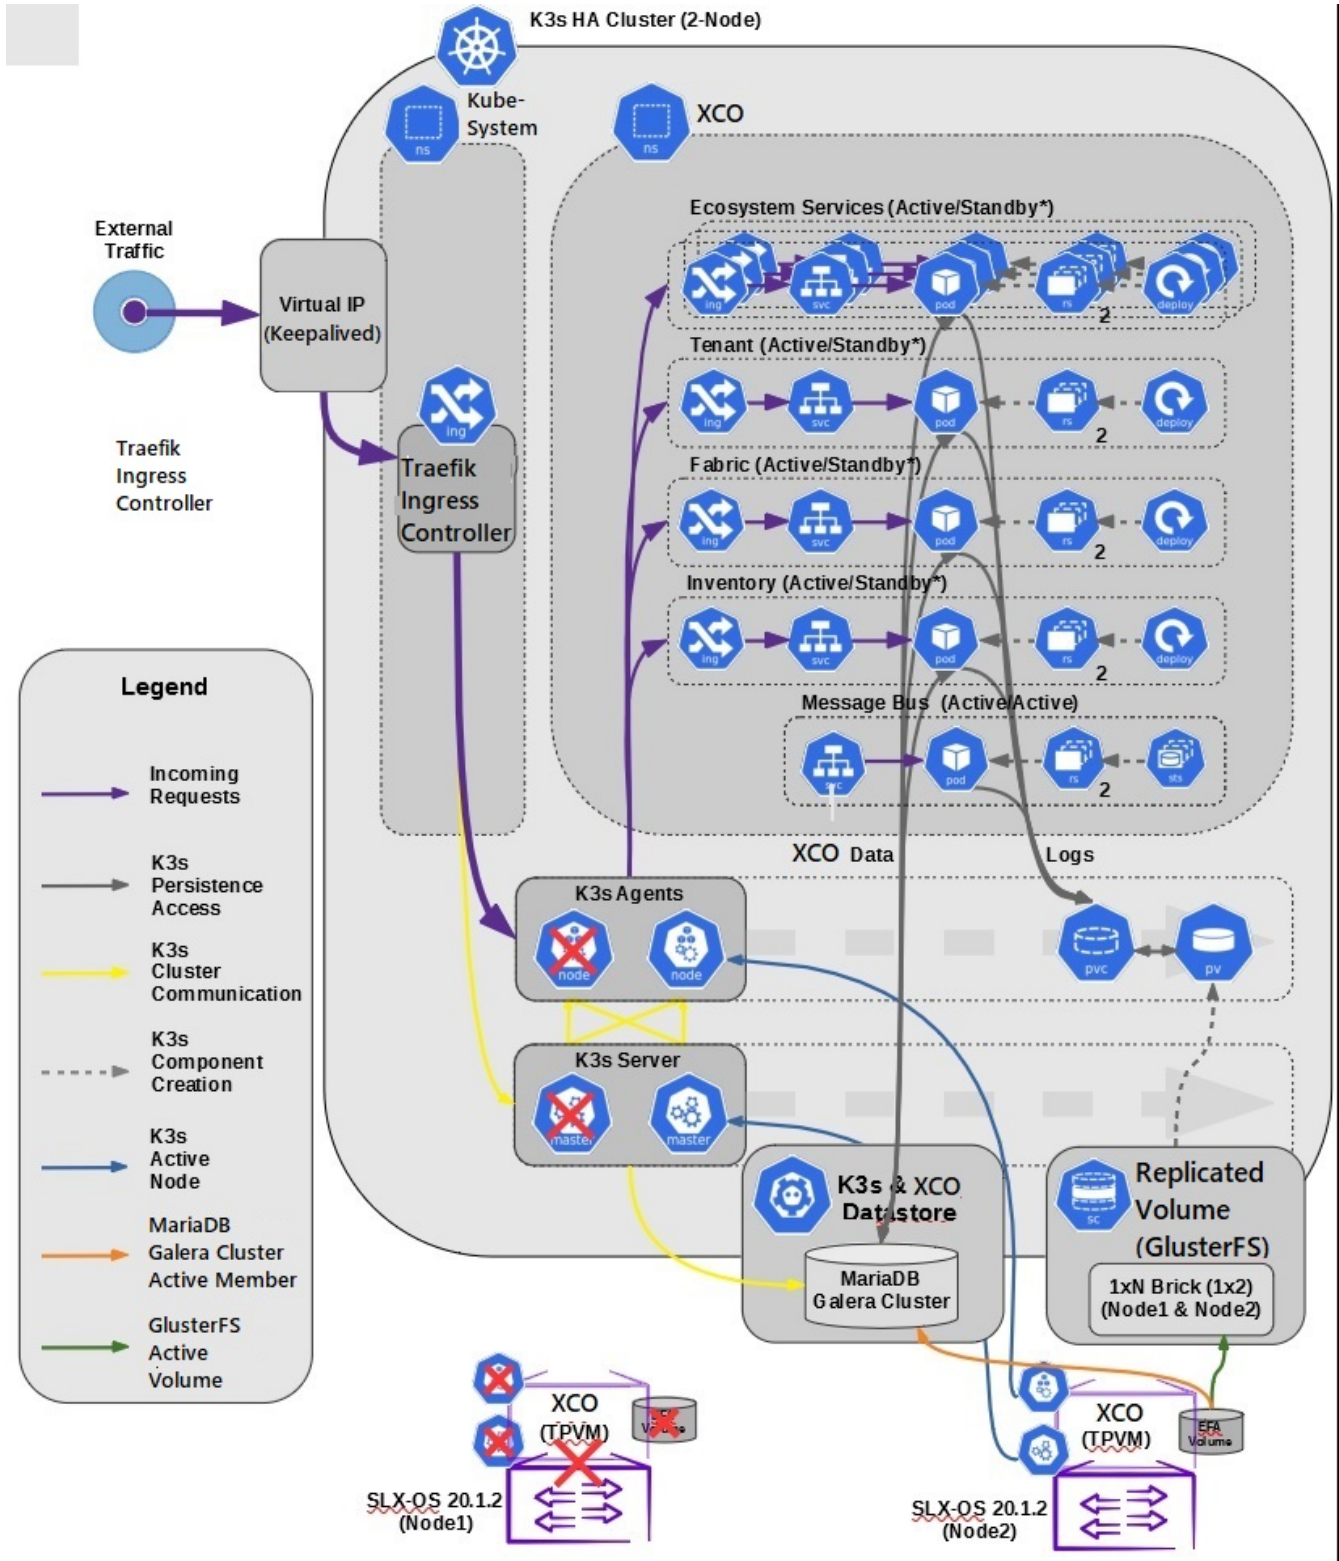


Figure 1: Two-node high-availability deployment

All operations provided by XCO services must be idempotent, meaning they produce the same result for multiple identical requests or operations. For more information, see

the "Idempotency" section in the [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).

XCO 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 performs the active role. Thus, the standby node becomes an active node. Each state change runs a keepalived notify script containing logic to ensure XCO'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 XCO.
- K3s server runs on 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. Therefore, the HA cluster looks like a single-node cluster. However, the HA cluster ties itself to the keepalived-managed virtual IP.
- MariaDB and Galera – XCO business states (device, fabric, and tenant registrations and configuration) are stored in a set of databases managed by MariaDB. Both the nodes run on 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 XCO log files, 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, XCO CLIs still operate correctly when they run from active or standby node. Commands are converted to REST and run 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 XCO services are Ready
 - Kubernetes state is consistent with all the enabled XCO services
 - The host is a member of Galera or MariaDB cluster
- For the standby node:
 - It is reachable via SSH from the active node
 - It is a member of Galera or MariaDB cluster
- For both the nodes:
 - The Galera cluster size is 2 if both the nodes are up. The cluster size is ≥ 1 if the standby node is down.

The following example shows the active and standby TPVM node status in a multi-node TPVM deployment:

```
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 XCO in a Multi-Node Deployment

You can install XCO in a multi-node cluster for high availability.

Before You Begin

- Ensure that the passwordless SSH login is enabled between the two servers. For more information, see [#unique_27](#).
- To install XCO, you must be a root user or have `sudo` privileges.



Note

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

About This Task

Follow this procedure to install XCO in a multi-node deployment.

Procedure

1. Untar the tarball on the primary server.

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

2. Change to the XCO directory.

```
device# cd efa
```

3. Run the installation script.

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

The XCO Installer begins in a series of dialogs.

4. When prompted, select **Multi-node deployment**, and then **Fabric suite**. Then select **OK**.

**Tip**

Use arrow keys to move between options. Press 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. 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 IPv4 or IPv6 addresses.
 - Select **No** to ignore this optional step.
9. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **Yes** and then provide the following information when prompted. Repeat as often as necessary.
 - Sub-interface name, which is a unique name that contains no more than 11 characters, no white space, and no % or / characters.
 - ID of the VLAN that the management network uses to tag traffic. Valid values range from 2 through 4093.
 - IP subnet address in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - 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.
 - 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.

As the installation proceeds, messages display showing the installation progress.

11. Verify the installation.
 - a. From the XCO command line, run the **sudo efact1 status** command to see the status of nodes, pods, and services.
 - b. From the XCO command line, run the **efa status** command to ensure that status of all the nodes are up.

Install XCO on TPVM in a Multi-Node Deployment

You can install XCO 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. XCO on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models for Fabric Skill](#) on page 12.

About This Task

Follow this procedure to install XCO on a TPVM in a multi-node deployment.

Procedure

1. Run the `show tpvm status` command and verify that the TPVM is set up for an XCO deployment.
 - a. Verify the versions of TPVM and SLX-OS.
For the supported version information, see [Supported Platforms and Deployment Models for Fabric Skill](#) 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 the passwordless access is configured.
 - e. (For SLX-OS 20.2.3 and later with TPVM version 4.2.2 and later) Verify that the 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 by running the following command.

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

- g. Verify that NTP is synchronized.
- h. (Optional) Log in to TPVM and configure the NTP time zone from SLX.

```
device# tpvm config timezone
```

- i. (Optional) Configure unique TPVM host names.

```
device# tpvm config host
```

2. Enter SLX Linux mode.

```
device# start-shell
```

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

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

4. Deploy XCO on TPVM from the SLX command line.

```
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, any new install or upgrade of XCO on TPVM in a multi-node deployment displays the following warning banner:

```
*****
*                               ! ! ! 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 XCO Installer begins in a series of dialogs.

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



Tip

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

- When prompted, select the IP mode and **OK**.



Tip

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

- When prompted, enter the peer IP address or FQDN of the other node in the cluster.
- When prompted, enter the virtual IP address for the cluster.
- 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.
- 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.
- 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.

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

**Note**

XCO Management Interface must have IPv4 address configured. Adding IPv6 address is optional while IPv4 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.

As the installation proceeds, messages display showing the installation progress.

13. 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 XCO command line, run the **efactl status** command to see the status of nodes, pods, and services.
- c. From the XCO command line, run the **efa status** command for concise status information.



XCO Upgrade

- [Upgrade XCO in a Single-node Deployment](#) on page 39
- [Upgrade XCO on TPVM in a Single-node Deployment](#) on page 40
- [Upgrade XCO from a Single-Node to a Multi-Node Deployment](#) on page 42
- [Upgrade XCO in a Multi-Node Deployment](#) on page 44
- [Upgrade XCO on TPVM in a Multi-Node Deployment](#) on page 45
- [Upgrade XCO on an OVA](#) on page 47
- [Upgrading SLX-OS, TPVM, and XCO Together](#) on page 48
- [Recover from an Upgrade Failure](#) on page 53
- [Rollback](#) on page 54

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

Upgrade XCO in a Single-node Deployment

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

About This Task

The upgrade process takes backup of the XCO 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 53.

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 XCO 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 XCO is deployed.
8. Verify the upgrade.
 - a. From the XCO command line, run the `sudo efact1 status` command to see the status of nodes, pods, and services.
 - b. Run the `efa status` command for concise status information.

Upgrade XCO on TPVM in a Single-node Deployment

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

Before You Begin

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

About This Task

XCO does not support *Non-secure mode*. For more information about the modes, see [#unique_32](#).

The upgrade process backs up the XCO system, so that you can easily recover data if the upgrade fails. For more information, see [Recover from an Upgrade Failure](#) on page 53.

Procedure

1. Enter SLX Linux mode.

```
device# start-shell
```

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

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


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

```
device# efa deploy
```



Note

From SLX version 20.4.1 and above, new install or upgrade of XCO 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 XCO 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 XCO 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 XCO command line, run the **sudo efactl status** command to see the status of nodes, pods, and services.
 - c. From the XCO command line, run the **efa status** command for concise status information.

Upgrade XCO from a Single-Node to a Multi-Node Deployment

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

Before You Begin

- Ensure that the single node is running EFA 2.5.5 or later. For more information, see [#unique_34](#) and [#unique_35](#).
- Ensure that you have completed the high-availability prerequisites in [#unique_27](#).

About This Task

The upgrade process takes approximately 15 - 25 minutes. During the upgrade process, XCO services remain down.

The upgrade process backs up the XCO system to recover data if the upgrade fails.



Note

XCO 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 XCO directory.

```
device# cd efa
```

5. Run the deployment script.

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

The XCO Installer begins in a series of dialogs.

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



Tip

Use arrow keys to move between options. Press 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.

As the installation proceeds, messages display showing the installation progress, including when deployment is complete.

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 XCO command line, run the **sudo efact1 status** command to see the status of nodes, pods, and services.
- c. From the XCO command line, run the **efa status** command to ensure that status of all the nodes are up.

Upgrade XCO in a Multi-Node Deployment

You can upgrade XCO in a multi-node, high-availability deployment.

About This Task

The upgrade process takes approximately 10 minutes to complete. During the upgrade process, XCO services are down but users or automated systems can continue to make calls into XCO.

The upgrade process automatically backs up the XCO database to 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 XCO directory.

```
device# cd efa
```

5. Run the deployment script.

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

The XCO 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

XCO 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.

As the installation proceeds, messages display showing the installation progress.

11. Verify the upgrade.

- a. From the XCO command line, run the `sudo efactl status` command to see the status of nodes, pods, and services.
- b. From the XCO command line, run the `efa status` command to ensure that status of all the nodes are up.

Upgrade XCO on TPVM in a Multi-Node Deployment

You can upgrade a multi-node deployment of XCO on TPVM (Third-Party Virtual Machine).

Before You Begin

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

About This Task

XCO on TPVM is supported only on the platforms described in [Supported Platforms and Deployment Models for Fabric Skill](#) on page 12. By default, XCO is installed in secure mode.

Installing XCO 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 XCO tar file to the SLX device.

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



Note

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

```
*****
*                               *
*           !!! 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 XCO on TPVM from the SLX command line.

```
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. Press 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**

XCO 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.

As the installation proceeds, messages display showing the installation progress.

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

Upgrade XCO on an OVA

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

Before you upgrade XCO on an OVA, see [#unique_39](#) and [#unique_40](#) for a list of prerequisites.

To upgrade a VM that is provisioned using an OVA image, follow the procedure explained either [#unique_34](#) in or in [#unique_41](#).

Upgrading SLX-OS, TPVM, and XCO Together

Use the topics to learn how to upgrade an SLX device to the latest supported version of SLX-OS and TPVM.

Requirement for SCP connections

The firmware server must support more than 10 unauthenticated SCP connections. To ensure this requirement, specify an appropriate value of '#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 XCO, 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 XCO, SLX-OS, and TPVM. For more information about supported versions, see [Supported Platforms and Deployment Models for Fabric Skill](#) on page 12.

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

Procedure

1. Upgrade XCO to the latest version.

a. Back up XCO.

```
efa system backup
```

For more information about backup and restore, see [#unique_44](#) and [#unique_45](#).

b. SCP the backup file to a location outside of TPVM, such as the `/efaboot` partition of SLX-OS where the XCO image is kept.

c. Copy the EFA image to the `/efaboot` directory on SLX1.

- d. Deploy XCO 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 a fabric.

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

For more details, refer [TPVM Incremental Upgrade using Auto-Reboot](#) on page 72.

- a. Back up XCO.

```
efa system backup
```

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

```
device# show tpvm status
```

- c. From the active XCO 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>
```

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

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

- e. When the status of the upgrade is complete, perform the following from the XCO 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).
- f. If a “System restart required” message appears when you run the **efa inventory device tpvm-upgrade show** command 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 XCO, 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 XCO, SLX-OS, and TPVM. For more information about supported versions, see [Supported Platforms and Deployment Models for Fabric Skill](#) on page 12.

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

Procedure

1. Upgrade XCO to the latest version.

- a. Back up XCO.

```
efa system backup
```

For more information, see "Back up and Restore the XCO System" in the [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).

- b. SCP the backup file to a location outside of TPVM, such as the `/efaboot` partition of SLX-OS where the XCO image is kept.
- c. Copy the EFA image to the `/efaboot` directory on SLX1.
- d. Deploy XCO 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 XCO 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 XCO 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 XCO 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 XCO. For more information, see [TPVM Complete Package Upgrade](#) on page 62.

- a. Back up XCO.

```
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, run the following command from the XCO command line on TPVM1 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, run the following command from the XCO command line on TPVM1 to 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. Verify the TPVM is upgraded by running the following command at the XCO command line.

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

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

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 will continue to exist 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 must be manually reconfigured 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 [Extreme SLX-OS 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 XCO synchronizes after the failover, view the output of the following commands to ensure that both nodes are in 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 XCO 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 XCO 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 XCO, SLX-OS, and TPVM Method 2](#) on page 50.
- 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 XCO from a Single-Node to a Multi-Node Deployment](#) on page 42
 - [#unique_48](#)
 - [#unique_34](#) and [#unique_35](#)
 - [Upgrade XCO in a Multi-Node Deployment](#) on page 44
 - [Upgrade XCO on TPVM in a Multi-Node Deployment](#) on page 45
2. To perform a fresh installation and restore the system backup, take the following steps.
 - a. Uninstall XCO to remove any components that might have been installed before the upgrade failed.
 - [Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#) on page 60
 - [Uninstall XCO in a Single-Node or Multi-Node Deployment](#) on page 60
 - b. Follow the steps for the type of installation you need.
 - [Install XCO on TPVM in a Multi-Node Deployment](#) on page 36
 - [Install XCO in a Multi-Node Deployment](#) on page 34
 - [#unique_51](#)
 - [#unique_52](#) and [#unique_53](#)
 - c. Restore the XCO backup.

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

For more information about backup tar files, see the "XCO System Backup and Restoration" section of the [ExtremeCloud Orchestrator Command Reference, 3.2.0](#).

Rollback

You can perform 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 must have the same version of TPVM after an upgrade.

About This Task

Follow this procedure to address a scenario in which TPVM2 (on SLX2) was upgraded, but TPVM1 (on SLX1) was rolled back to a previous version because of an upgrade failure. 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 XCO node. SLX2 and TPVM2 refer to the active XCO node. You can use incremental debian package for the rollback from TPVM version 4.5.10 to 4.5.9.

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 XCO synchronizes after the failover, run the following commands from TPVM to ensure that both nodes are in 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 XCO command line on TPVM1 (the active XCO), 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 XCO 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 XCO on the TPVM.

For more information, see [Install XCO on TPVM in a Multi-Node Deployment](#) on page 36.

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 XCO

Initiate a rollback when there is a deployment failure.

About This Task

Follow this procedure to rollback a deployment failure.

Procedure

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

```
no efa deploy
```

2. Copy the XCO instance.

```
efa deploy
```

3. Use 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. Log in to the XCO system as an Extreme user.
5. Restore the XCO configuration.

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

6. Post XCO rollback, to verify if there are any devices in the `cfg refresh error` state, run the following command:

```
efa fabric show
```

- To remove the devices from the `cfg refresh error` state, run the following command:

```
efa inventory device update --ip <device_ip>
```




Node Replacement

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

[Replace a Node in a Multi-node TPVM Deployment](#) on page 58

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 2.5.5 or later.
- Ensure you have completed the high-availability prerequisites in [#unique_27](#).
- Ensure that XCO 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 XCO is installed.

Procedure

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

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

The XCO 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. Press 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.

- Verify the status of XCO after the node replacement. Ensure that status of all the nodes are up.

```
$ efa status
```

For more information on how to recover SLX configs, refer to the [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).

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 that the cluster with faulty node is running EFA 2.5.5 or later.
- Ensure that you have completed the high-availability prerequisites described in [#unique_40](#).
- Ensure that XCO is not deployed on the replacement node.

About This Task

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

Perform this procedure on the active node where XCO is installed.

Procedure

- Enter SLX Linux mode and copy the XCO tar file to the SLX device.

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

- Deploy XCO 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 XCO Installer begins in a series of dialogs.

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



Tip

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

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

As the node replacement proceeds, messages display showing the replacement progress.

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

```
device# sudo efactl status
```

**Note**

To recover the SLX configuration, see the "Replace a Faulty Device" topic in the [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.0](#).



XCO Uninstallation

[Uninstall XCO in a Single-Node or Multi-Node Deployment](#) on page 60

[Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#) on page 60

Uninstall XCO in a Single-Node or Multi-Node Deployment

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

About This Task

Follow this procedure to uninstall XCO.

Procedure

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

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

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

Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment

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

About This Task

Follow this procedure to uninstall XCO on TPVM in a single-node and multi-node deployment.

Procedure

1. (On a single-node deployment) From the SLX device console, run the following command to uninstall XCO:

```
device# no efa deploy
```

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

As the uninstallation proceeds, messages display showing the uninstallation progress.

2. (On a multi-node deployment) Stop and then start the TPVM to ensure there are no DNS resolution issues.

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

3. (On a multi-node deployment) On the node where XCO is installed, run the following command to uninstall XCO:

```
device# no efa deploy
```

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

As the uninstallation proceeds, messages display showing the uninstallation progress.



TPVM Upgrade from XCO

[TPVM Complete Package Upgrade](#) on page 62

[TPVM Incremental Upgrade](#) on page 68

TPVM Complete Package Upgrade

You can upgrade a device's TPVM image when a TPVM is installed and running an XCO instance which is managing the device. After the TPVM image update, an XCO instance gets reinstalled with the same XCO version as before and rejoines with the active XCO instance.



Note

For more information about commands and supported parameters, see [ExtremeCloud Orchestrator Command Reference, 3.2.0](#)

Assumptions and Limitations

- XCO supports SLX-OS 20.3.2a and later. The TPVM upgrade has SLX-OS dependencies for the new SLX commands: **tpvm upgrade** and **tpvm revert**. The TPVM configurations has already been present in the SLX running-config.
- XCO does not support TPVM upgrade on a single-node TPVM deployment.
- You can deploy a TPVM on multiple high-availability (HA) nodes, but you can upgrade a TPVM only on the standby TPVM node.
- TPVM upgrade is allowed only on the XCO HA nodes that are managing the devices and hosting the XCO HA instances.
- The XCO version is reinstalled and remains the same after the TPVM upgrade. You must not perform an XCO version upgrade during a TPVM upgrade. Ensure that you perform the XCO version upgrade before or after the TPVM upgrade is completed on both the HA nodes.
- You can upgrade only one device's TPVM at a time.

TPVM Upgrade Workflow Dependencies

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

TPVM Configuration Persistence

The TPVM running configuration and operational data (including the TPVM image version and TPVM IP address) from the SLX device continue to exist in the XCO DB. The following table describes the TPVM configurations present in the XCO 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 XCO TPVM.
password	Pre-start only	String	An encoded non-clear text password string. If does not exist then default is "password".	Extreme user password is not a clear-text in the running-config. The 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 XCO 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"> • Host • Port • Secure • basedn • rootdn • Password 	Post-start		<ul style="list-style-type: none"> • FQDN or IPv4 address or IPv6 address • 0-65535 • Exists or not exists • Base domain name • Root domain name • Root domain name password 	
ldap ca-cert <ul style="list-style-type: none"> • protocol • user • password • host • directory • filename 	Post-start	String	<ul style="list-style-type: none"> • scp • Username • Password • IPv4 address • Directory • Filename 	The ca-cert for LDAP must be stored on the firmware-host and for XCO to support the node replacement. The ca-cert can also have IPv6 address.
trusted-peer <ul style="list-style-type: none"> • ip • password • sudo-user 	Post-start	String	<ul style="list-style-type: none"> • IPv4 address • Sudo user password • Sudo username 	Trusted-peer config exists on one of the XCO nodes. Push this config to the correct node after the upgrade.
deploy	Install	Boolean	<ul style="list-style-type: none"> • Exist • Does not exist 	Installs, starts, and applies the configurations to the TPVM instance.

Device Registration Enhancements

The TPVM config information already exist in the XCO DB when a device is registered or during the initial device registration. The **TPVM running-config** information is read and stored during the device discovery so that user visible device registration times are not impacted. The TPVM config is fetched and stored only during initial device registration and not during subsequent device updates.

Timer-based TPVM Config Updates

A timer is set to poll daily data for any TPVM config changes for XCO HA peer managed devices.

TPVM Upgrade Workflow

You can upgrade a TPVM on a single-node and multi-node deployment.

About This Task

Follow this TPVM upgrade procedure to upgrade an existing TPVM.

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 the associated TPVM configuration. It must be one of the XCO HA peers managing the device.
 - b. Ensure that the firmware host is registered prior to TPVM upgrade.
 - c. The TPVM image is validated 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 is pushed to the device in the node replacement case.
 - a. If TPVM is neither configured nor installed, then the TPVM configuration existed in the XCO 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 existing XCO configuration, then the device's configuration has priority, and the XCO DB is updated.

TPVM Configuration Special Handling for RMA Node Replacement Case

- When you set the TPVM configuration interface management IP to DHCP, ensure that the TPVM IP address remains the same. This is due to a dependency on XCO deployment where a peer node is configured with a specific IP address in the active node. You cannot change the peer node IP without restarting XCO HA cluster daemons on an active node.

TPVM Configuration Special Handling for All Cases

- You must re-apply the trusted-peer configuration on the node where it was already applied. It exists on only one of the nodes in the XCO HA cluster. An appropriate node is identified and the trusted peer configuration is pushed to the correct node during TPVM upgrade or node replacement.
3. Run an appropriate SLX command on the device to upgrade or install the TPVM.
 - a. Run the **tpvm upgrade** command on 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, and the existing TPVM configurations are programmed on the running TPVM instance.
 - b. During node replacement, the TPVM configuration is pushed to the device, and the **tpvm deploy** command is run on the device. You do not require a TPVM snapshot because the replacement switch is a new switch without a configured TPVM.
 4. Redeploy XCO on the upgraded or installed TPVM node from the active node. Allow the redeployed peer node to rejoin the XCO HA cluster.

Example

The following example shows an output of **TPVM upgrade execute** command:

```
(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   |2022-11-05   |
|.48.162|    |   -32C |       |       |       |           |Progress|
Started    |         |         |Upgrade    |23:52:29 -0700 PDT|23:52:36 -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   |2022-11-06   |
|.48.162|    |   -32C |       |       |       |           |Workflow
Completed|for TPVM Instance|         |Upgrade    |23:52:29 -0700 PDT|00:01:11 -0700 PDT|
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
```

TPVM Upgrade Workflow States

This topic describes all the upgrade states in a TPVM upgrade workflow.

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	Success: TPVM Config Validation Failure: TPVM Upgrade Workflow Finished	Normal Upgrade Node Replacement	Ensure that the provided device IP has an associated TPVM configurations in the XCO DB, and the device's TPVM IP is one of the XCO peer node IPs.
TPVM Config Validation	-Normal Upgrade: Success: TPVM Upgrade Failure: TPVM Upgrade Workflow Finished -Node Replacement: Success: TPVM Configuration Failure: TPVM Upgrade Workflow Finished	Normal Upgrade Node Replacement	<p>Read TPVM config and operational data from the device and determine if it is a normal TPVM Upgrade or a faulty node replacement.</p> <ol style="list-style-type: none"> 1. If TPVM config and operational data are present on the device and TPVM IP is one of the XCO peers, then it is a normal TPVM upgrade. 2. If there is no TPVM config present on the device, then it is a node replacement. 3. If TPVM config and operational data are present on the device and TPVM IP does not match one of the XCO peers, then validation for a normal TPVM upgrade was unsuccessful. <p>The Detailed Status column in the <code>tpvm-upgrade show</code> command output shows the nature of the issue and possible remedy.</p>
TPVM Configuration	Success: TPVM Installation Failure: TPVM Upgrade Workflow Finished	Node Replacement	Device running-config is programmed using TPVM config data from XCO DB.

TPVM Upgrade State	Next State	Case	Description
TPVM Installation	Success: XCO Deploy Peer and Rejoin Failure: TPVM Upgrade Workflow Finished	Node Replacement	TPVM install and start is invoked on the device.
TPVM Upgrade	Success: XCO Deploy Peer and Rejoin Failure: TPVM Revert	Normal Upgrade	TPVM upgrade is invoked on the device.
TPVM Revert	Success: TPVM Upgrade Workflow Finished Failure: TPVM Upgrade Workflow Finished	Normal Upgrade	On failure of “Upgrading TPVM” or “Deploying XCO for Rejoin”, the TPVM revert state is invoked to rollback the TPVM upgrade failure.
XCO Deploy Peer and Rejoin	Success: TPVM Upgrade Workflow Finished Failure: TPVM Revert	Normal Upgrade Node Replacement	On active XCO node, re-deploying of XCO 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

The TPVM incremental upgrade allows you to upgrade active and standby TPVM nodes. It is applicable for multi-node and single node deployment where XCO is running. It reduces the upgrade time (around 3 minutes) compared to the full upgrade.

Before You Begin

- Ensure that you are using EFA 3.0.0 and later.
- Ensure that the TPVM is running and the SLX version is 20.4.1 or later.

About This Task

XCO 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 the upgrade is 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 the 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 XCO restart or inventory service restart.
- The Detailed Status of TPVM upgrade shows whether a reboot is required for an incremental upgrade. If the output of an incremental upgrade 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 command fetches either the active TPVM IP address or the standby TPVM IP address or both for a TPVM incremental upgrade. XCO 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 command shows the device status for a TPVM upgrade operation.

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

The following are the output examples of a TPVM upgrade show command:

- 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 | 4200000000 | 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 | 4200000000 | 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 | 4200000000 | 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 | 4200000000 | 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 | Model | Chassis
| ASN       | Role | Current TPVM | Target TPVM |
Update | Status | Detailed | Start Time | Last Update Time | Failed |
|           | Name | Name | | |
|           | Version | Version | |
State | Status | State |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX | 4001 | BR-
SLX9640 | 4200000000 | Leaf | 4.5.0 |
| In | TPVM Upgrade | None | 2022-05-26 | 2022-05-26 | |
| | | | | |
| | | | | Progress | Started
| | 13:41:14 -0700 PDT | 13:41:58 -0700 PDT | |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.56 | SLX | 4001 | BR-
SLX9640 | 4200000000 | Leaf | 4.5.0 |
| In | TPVM Upgrade | None | 2022-05-26 | 2022-05-26 | |
| | | | | |
| | | | | Progress | Started
| | 13:41:14 -0700 PDT | 13:41:58 -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 | Model | Chassis |
ASN | Role | Current TPVM | Target TPVM | Update | Status
| Detailed Status | Failed | Upgrade | Start Time | Last Update Time |
| Name | Name | | |
| Version | Version | State |
| State | Type | | |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.56 | SLX | 4001 | BR-SLX9640 | 4200000000 |
Leaf | 4.5.2 | 4.5.2 | Completed | TPVM Upgrade
| Reboot Required | | Incremental | 2022-07-28 | 2022-07-28 |
| | | | | |
| | | | | Workflow Completed |

```

```

for TPVM Instance|          | Upgrade      | 16:19:46 -0700 PDT| 16:25:34 -0700 PDT|
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX | 4001 | BR-SLX9640 |
4200000000| Leaf | 4.5.2 | 4.5.2 | Completed| TPVM Upgrade
| None | | | Incremental| 2022-07-28 | 2022-07-28 |
| | | | | | |
| | | | | Workflow Completed|
| | Upgrade | 16:19:46 -0700 PDT| 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 command shows the information, such as TPVM IP address, SLX IP address, TPVM hostname, TPVM version, and SLX firmware version.

```

efa inventory device tpvm list
+-----+-----+-----+-----+-----+-----+
| Device IP | TPVM IP | TPVM | SLX Firmware Version | TPVM |
| Address   | Address | Hostname| | | | |
+-----+-----+-----+-----+-----+-----+
| 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 | |
+-----+-----+-----+-----+-----+-----+

```

TPVM Incremental Upgrade using Auto-Reboot

TPVM incremental upgrade on a multi-node TPVM in XCO 3.2.0 removes the need of manual reboot of TPVM nodes.

About This Task

TPVM incremental upgrade on a multi-node TPVM in EFA 3.1.0 or earlier requires manual reboot of standby and active TPVM nodes.

XCO 3.2.0 provides an auto-reboot option in the incremental upgrade command. The upgrade process reboots the standby TPVM first (if required) and waits till the standby TPVM comes up. After the standby TPVM is up and running, the upgrade process reboots the active TPVM (if required). Use the `efa status` command to verify the active and standby TPVM node status.

If the auto-reboot feature reboots the active TPVM, ensure that you run the show command on the active TPVM to check the upgrade status. Use the show command in the following two cases:

- Using the IP addresses of the devices.
- Using execution ID: You must store the upgrade status information. If the active TPVM reboots, use this information on the new active TPVM.

Procedure

Run the following TPVM upgrade command:

```
efa inventory device tpvm-upgrade execute --ip <deviceips> --firmware-host <firmwarehost>
--tpvm-image <tpvm_image> --auto-reboot
```

Example

The following examples show the TPVM upgrade configuration with the auto-reboot option:

```
$ efa inventory device tpvm-upgrade execute --ip 1.1.1.1,1.1.1.2 --firmware-host 2.2.2.2
--tpvm-image /buildsjc/sre_fusion/Nightly/tpvm/ci_tpvm/tpvm_inc_upg-4.5.5-5.amd64.deb --
auto-reboot
```

```
TPVM Upgrade Execute [success]
Monitor TPVM upgrade execution progress using:
```

```
efa inventory device tpvm-upgrade show --ip 1.1.1.1,1.1.1.2
efa inventory device tpvm-upgrade show --execution-id
aac8b30e-6911-4bea-8f67-986bd40528a9
```

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

```
$ efa inventory device tpvm-upgrade show --ip 1.1.1.1,1.1.1.2
```

```
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
|IP      |Host|Model|Chassis | ASN      |Role |Current TPVM|Target TPVM|Update  |
Status   |Detailed Status |Failed|Upgrade  | Start Time  | Last Update Time |
|Address|Name|      |Name   |           |      |Version      |           |         |
|Version |State |      |       |           |      |           |           |         |
|Type    |      |      |       |           |      |           |           |         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
|1.1.1.1|SLX |4001 |BR      |4200000000 |Leaf |4.5.5      |           |In      |TPVM
Upgrade |None |     |      |           |      |Incremental|2022-12-10 | 2022-12-10 |
|      |    |     |      |           |      |           |           |         |
|      |Progress|Started |           |           |           |           |           |         |
|Upgrade |19:36:05 -0800 PST| 19:39:26 -0800 PST|           |           |           |           |         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
|1.1.1.2|SLX |4001 |BR      |4200000000 |Leaf |4.5.5      |4.5.5      |In      |Reboot
Required|Waiting for TPVM |           |Incremental|2022-12-10 | 2022-12-10 |           |           |         |
|      |    |     |      |           |      |           |           |Progress|on TPVM
Instance|Instance to come up| Upgrade |19:36:05 -0800 PST| 19:40:31 -0800 PST|           |           |         |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+
+-----+-----+
```



Upgrade Ubuntu

[Upgrade Ubuntu on the XCO Host - Single Node or Multi Node](#) on page 74

[Upgrade Ubuntu on the XCO Host - Multi-Node](#) on page 76

Upgrade Ubuntu on the XCO Host - Single Node or Multi Node

Upgrade Ubuntu in single-node and multi-node deployments.

Before You Begin

- Ensure that XCO is at release 3.2.0 or later.
- Ensure that the nodes are healthy and XCO services are up and running.

About This Task

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

Follow this procedure to Upgrade Ubuntu from 18.04 to 20.04.



Note

- This process is not supported for deployments of XCO on TPVM. For TPVM upgrade, see [Upgrading SLX-OS, TPVM, and XCO Together](#) on page 48, [#unique_69, TPVM Complete Package Upgrade](#) on page 62, and [TPVM Incremental Upgrade](#) on page 68.
- 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. If you already have XCO running on Ubuntu 18.04, take a backup of XCO.
For more information on XCO backup, see [#unique_70](#).

```
$ efa system backup
Generating backup of EFA...
Backup Location: /var/log/efa/backup/EFA-3.2.0-GA-2023-03-23T13-25-55.105.tar
--- Time Elapsed: 37.543063336s ---
```

2. Copy the generated backup archive to a remote server.
3. Uninstall XCO.
For more information, see [Uninstall XCO in a Single-Node or Multi-Node Deployment](#) on page 60.

4. Run the following command to ensure that there are no packages on hold:

```
# apt-mark showhold
```

To un-hold the packages, run the following command.

```
# sudo apt-mark unhold <package-name>
```

5. Update the Ubuntu package, and then update all the Ubuntu packages on all the available nodes.

If needed, perform reboot after the update.

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

6. Upgrade the Ubuntu package on all the available nodes.

Verify that the nodes are at the new version by running the `cat /etc/os-release` command.

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

7. Verify that the nodes are at the new version by running the `cat /etc/os-release` command.

```
# cat /etc/os-release
NAME="Ubuntu"
VERSION="20.04.6 LTS (Focal Fossa)"
ID=ubuntu
ID_LIKE=debian
PRETTY_NAME="Ubuntu 20.04.6 LTS"
VERSION_ID="20.04"
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=focal
UBUNTU_CODENAME=focal
```

8. Install XCO 3.2.0 in a single-node server or multi-node cluster.
9. Copy the backup to the `/var/log/efa/backup/` directory in the local system.
10. Run the following restore command:

For more information on XCO restore, see [#unique_71](#)

```
# efa system restore

1. EFA-3.2.0-GA-2023-03-23T13-25-55.105.tar (Version:3.2.0-GA, Generated by: user)

Choose backup option:1
Selected: EFA-3.2.0-GA-2023-03-23T13-25-55.105.tar
Performing EFA restore using EFA-3.2.0-GA-2023-03-23T13-25-55.105.tar
Restore operation ID: 4a861b2c-c971-11ed-b7fd-000c292bbb0e
Stopping all EFA services...
All pods are terminated
Restoring databases...
Restore of encryption keys is completed
Restore operation is successful
--- Time Elapsed: 4m33.591911082s ---
```

- a. When the restore is complete, run `source /etc/profile`.

You can now log in to XCO.

- b. To enable secure connections, install the certificates on devices.

```
efa certificate device install --ip 10.20.61.92 --cert-type https
```

The command installs the HTTPS or OAuth2 certificate on one or more devices.

- c. To get the current state of the devices, run the **efa inventory device update** command after you run the restore command.
11. Verify that XCO is up and running. Perform the following (from the XCO command line) on both nodes.
 - Run **efa status** to verify that both nodes are up.
 - Run **sudo efact1 status** to verify that all pods on the active node are in Running state.
 - Run **sudo efact1 db-status** to verify that the MariaDB is active (running)

Upgrade Ubuntu on the XCO Host - Multi-Node

Upgrade Ubuntu in multi-node deployments.

Before You Begin

- Ensure that you XCO 3.2.0 is up and running.
- Ensure that the nodes are healthy and the XCO services are up and running.

About This Task

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



Note

- This process is not supported for deployments of XCO on TPVM. For TPVM upgrade, see [Upgrading SLX-OS, TPVM, and XCO Together](#) on page 48, [TPVM Complete Package Upgrade](#) on page 62, and [TPVM Incremental Upgrade](#) on page 68.
- Ensure 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. Run the following command to install the packages that have been kept back:

Ensure that you complete this step one after the other for all the packages that have been kept back.

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

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

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

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

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

- c. Upgrade the second node using the same procedure as used on the first node of the cluster.



Note

If an upgrade fails due to dependencies, run the following command to upgrade all the dependencies at a time:

```
sudo apt-get upgrade <package1> <package2> <package3> .....
```

- d. Verify that the nodes are at the new version by running `uname -a` and `cat /etc/os-release` command.

```
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="20.04 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
```

- e. Verify that XCO is operational. Perform the following (from the XCO 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 78

[Enable Redundant Management](#) on page 79

[Redundant Management Data Path](#) on page 81

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.

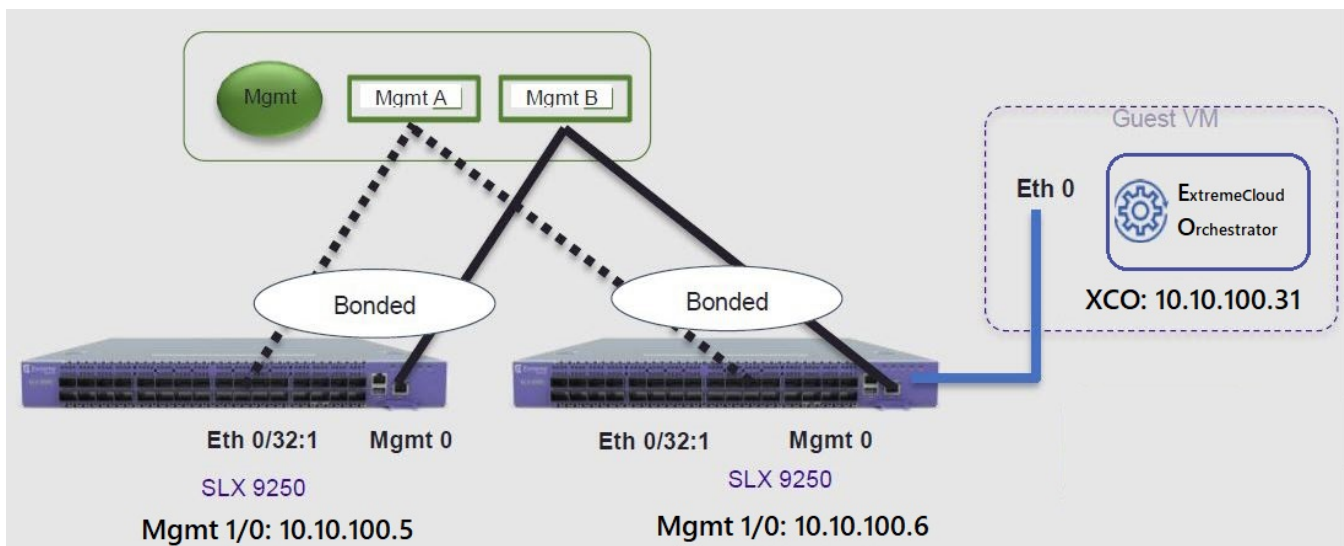


Figure 2: Redundant Management Network Overview

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 stand by 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 - It has a dual management port and does not need RME CLI.
- Extreme 8820 - Redundant Management is supported on the Extreme 8820 devices.
- Extreme 8520 - Preferred ports are 10G/1G port in 1G mode.

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 independent of front panel physical port limit and capability.

Enable Redundant Management

Redundant management provides fault tolerance for the management path.

About This Task

Follow this procedure on a supported SLX device. For more information, see [Redundant Management Network Overview](#) on page 78.

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

Example

The following example configures Ethernet 0/32 at 10G speed.

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

The following example configures Ethernet 0/32 at 1G speed.

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

The following example 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(config-if-eth-0/32:1)# redundant-management enable
device(config-if-eth-0/32:1)# no shut
```

Example

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

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




Figure 3: Data path overview



Flexible XCO Deployment for TPVM

[Flexible XCO Deployment for TPVM](#) on page 83

[Input Parameters on Single-Node Install or Upgrade](#) on page 84

[Input Parameters on Multi-Node Install or Upgrade](#) on page 85

[XCO Installer Improvements for TPVM-Based Deployment](#) on page 87

[No Graphics Mode](#) on page 88

[XCO Deployment with Rollback](#) on page 89

[Rollback the XCO Upgrade](#) on page 89

Flexible XCO Deployment for TPVM

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



Note

- Ensure that you have valid XCO packages and TPVM installations stored 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 you have not update the SLX 20.4.2 or above to reflect the XCO changes, the parameter change does not appear in deployment. Use the non-interactive parameters with the XCO deployment commands.
- When TPVM or XCO deployment is in progress, do not reboot or toggle management ports on the target devices. Avoid using 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

Use the commands directly on SLX command line to specify the parameters for 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 ...
```

XCO Deployment

Use of the graphical interface is the default procedure for deploying XCO. If you do not want to use graphical interface, use the optional parameters to deploy XCO.

The following table describes the minimal required commands to start the XCO deployment:

Deployment type	Commands
Single-node install or upgrade	efa deploy non-interactive single-node package / efaboot/efa-3.1.0.tar.gz
Multi-node install or upgrade	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
Multi-node upgrade with replacement	efa deploy non-interactive multi-node package /efaboot/efa-3.1.0.tar.gz replacement-ip 10.200.246.155

Listing XCO Packages

XCO packages are stored in the /efaboot directory on an SLX device.

Ensure that you have a package name for single-node and multi-node installations. Use the **show efa packages** command to show the available packages.

Input Parameters on Single-Node Install or Upgrade

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



Note

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

Minimum Required Commands

The following is the minimum required command for a single-node XCO deployment:

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

The following example deploys XCO on a single-node in a non-interactive mode:

```
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. XCO supports only adding a single sub-interface.

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

Input Parameters on Multi-Node Install or Upgrade

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

Deployment Type

Specify the deployment type using the **multi-node** command after the deployment. For installation, use the **peer-node** command followed by the peer node IP address.

Use **vip4** command followed by a virtual IP address to provide a 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 a virtual IP address to provide a virtual IPv6 address for the installation.

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

The following example shows the vip6 attributes in an XCO multi-node deployment in a non-interactive mode:

```
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. XCO 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 a management IP, the **sub-interface-name** command requires a name of the sub-interface. Ensure that the VLAN ID **sub-vlan-id** and the external subnet address **external-subnet** must be in the CIDR format.

The following example shows the management IP attributes for a single-node XCO deployment in non-interactive mode:

```
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, the following options are available 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.

The following example shows multi-node build upgrade with node replacement in a non-interactive mode:

```
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

Ensure that XCO is connected to a gateway during installation. If the gateway connectivity fails, the installation will fail.

When you use the standard VRRP to obtain a gateway address, the XCO installation fails.

The **ping-target** parameter in single CLI pings from both the nodes to the **ping-target** IP addresses. If the **ping-target** is not reachable from any of the nodes, the installation or upgrade fails.

You can provide a maximum of two IP addresses as input to **ping-target**. The IP addresses can be IPv4 or IPv6.

If you have installed XCO with `ping-target`, but you have not given the `ping-target` argument during upgrade, 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, provide the “clear” option to `ping-target` to clear the old values and ping the default gateway.

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

Single CLI Commands

The following table provides a list of single CLI command supported on SLX:

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

XCO Installer Improvements for TPVM-Based Deployment

When you fresh install or upgrade XCO on a TPVM, the following services are 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 XCO CLI, you can enable or disable the following 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 XCO in server mode.

Related Topics

[Upgrades and Service State](#) on page 88

[Enable or Disable Services](#) on page 88

Upgrades and Service State

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

Enable or Disable Services

You can enable or disable the following services for TPVM or Server-based deployments in XCO:

- 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.
```

The following examples show how to enable or disable a service:

```
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 if you do not want to view the graphic progress bar during installation process.

The **efa deploy nographics** command does not display any graphic 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.

XCO Deployment with Rollback

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

Choose from the following options to deploy XCO with rollback:

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

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

2. Deploy XCO 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 XCO 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 XCO 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 XCO Upgrade

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

About This Task

Follow this procedure to rollback the XCO 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 available 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 XCO 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 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

The 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#
```

The 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#
```