

ExtremeCloud Orchestrator Deployment Guide

3.2.1

9037728-01 Rev AB
January 2024



Copyright © 2024 Extreme Networks, Inc. All rights reserved.

Legal Notice

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

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

Trademarks

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

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

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

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
What's New in this Document.....	10
XCO Deployment Preparation.....	11
General Requirements.....	11
OS Versions and Security Updates.....	12
General Restrictions.....	12
NTP.....	12
DNS.....	13
IP Address Restrictions.....	13
TCP and UDP Ports for Single-Node Deployment.....	14
TCP and UDP Ports used for Multi-Node Deployment.....	15
Supported Platforms and Deployment Models for Fabric Skill.....	16
Supported Platforms and Deployment Models for Visibility Skill.....	19
Pure IPv6 Support.....	20
Limitations.....	20
Supported Platforms.....	21
Pre-configuration for Selecting IP Mode.....	21
Interactive Mode.....	22
Non-Interactive Mode.....	23
Single CLI on Server for Single Node.....	24
OVA Support.....	25
Upgrade Scenarios.....	25
XCO Status.....	26
TPVM Based Deployments.....	28
TPVM Overview.....	28
TPVM Deployment Requirements.....	29
Redundant Management Network Overview.....	29
Linux Bonding.....	30
Supported Ports.....	31
No Redundancy Period.....	31
Standby Port Rate Throughput.....	31
Enable Redundant Management.....	31
Redundant Management Data Path.....	33
TPVM High-Availability Requirements.....	34
OS Version.....	34

NTP.....	35
DNS.....	35
IP Addressing.....	35
SSH.....	35
Install XCO Single-Node Fabric Management on TPVM.....	37
Back up Database.....	43
Restore Database.....	44
Upgrade XCO on TPVM in a Single-Node Deployment.....	44
Upgrade TPVM OS for XCO Single-Node Deployment.....	49
Migrate TPVM Single-Node to Multi-Node.....	49
Install XCO on TPVM in a Multi-Node Deployment	52
Back up XCO Database.....	54
Restore Database.....	55
Upgrade XCO on TPVM in a Multi-Node Deployment.....	56
Upgrading SLX-OS, TPVM, and XCO Together.....	58
TPVM Complete Package Upgrade.....	63
TPVM Incremental Upgrade.....	69
Replace a Node in a Multi-node TPVM Deployment.....	74
Maintain TPVM Versions After a Rollback in a Multi-Node Deployment.....	76
Rollback XCO.....	77
Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment.....	77
Flexible XCO Deployment for TPVM.....	78
SLX CLI.....	78
XCO Deployment.....	79
Listing XCO Packages.....	79
Input Parameters on Single-Node Install or Upgrade.....	79
Input Parameters on Multi-Node Install or Upgrade.....	80
XCO Installer Improvements for TPVM-Based Deployment.....	82
No Graphics Mode.....	83
XCO Deployment with Rollback.....	84
Rollback the XCO Upgrade.....	84
Server Deployments.....	88
Server Requirements.....	88
Server and Cloud VM Requirements.....	88
Server and Cloud VM Deployment Models.....	89
High-Availability Requirements.....	89
Single-Node Deployment.....	92
XCO Installer Improvements for Server-Based Deployment.....	92
Install Fabric using Interactive Mode.....	92
Install Fabric using Single Command.....	96
Install Fabric using OVA.....	98
Install Visibility using Interactive Mode.....	99
Install Visibility using Single Command.....	100
Install Visibility using OVA.....	103
Post OVA Install Procedure for Fabric and Visibility.....	104
Backup Database.....	106
Restore Database.....	107
Upgrade of XCO in Single-Node Deployment.....	108
High-Availability Multi-Node Deployment.....	117

XCO Deployment for High Availability	117
Install XCO in a Multi-Node Deployment.....	121
Upgrade XCO in a Multi-Node Deployment.....	123
Upgrade XCO on an OVA.....	124
Recover from an Upgrade Failure.....	124
Upgrade Ubuntu on the XCO Host - Single Node or Multi Node.....	125
Upgrade XCO from a Single-Node to a Multi-Node Deployment.....	127
Replace a Node in a Multi-node Deployment.....	129
Rollback.....	130
Rollback SLX.....	130
Uninstall XCO in a Single-Node or Multi-Node Deployment.....	130



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.



What's New in this Document

The following table describes the information added to this guide for the ExtremeCloud Orchestrator (XCO) 3.2.1 software release.

Table 4: Summary of changes

Description	Topic
Updated the topics for setting the TPVM configuration interface management IP to DHCP.	<ul style="list-style-type: none">• TPVM Upgrade Workflow on page 66• TPVM Upgrade Workflow Dependencies on page 64



XCO Deployment Preparation

[General Requirements](#) on page 11

[TCP and UDP Ports for Single-Node Deployment](#) on page 14

[TCP and UDP Ports used for Multi-Node Deployment](#) on page 15

[Supported Platforms and Deployment Models for Fabric Skill](#) on page 16

[Supported Platforms and Deployment Models for Visibility Skill](#) on page 19

[Pure IPv6 Support](#) on page 20

Learn about the prerequisites of an XCO deployment.

General Requirements

XCO supports Server, OVA, and TPVM deployment models.



Note

- From release 3.2.0 onwards, Extreme Fabric Automation (EFA) is referred to as ExtremeCloud Orchestrator (XCO). The terms EFA and XCO refer to the same product and are used interchangeably.
- Upgrade from Extreme Visibility Manager (XVM) to XCO is not supported.

Be aware of the following prior to deployment:

- [General restriction](#)
- [TCP and UDP Ports for Single-Node Deployment](#) on page 14
- [TCP and UDP Ports used for Multi-Node Deployment](#) on page 15
- [Supported Platforms and Deployment Models for Fabric Skill](#) on page 16
- [Supported Platforms and Deployment Models for Visibility Skill](#) on page 19

OS Versions and Security Updates

Ensure that your operating systems for VM, OVA, and Server deployments are up to date.

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



Note

For best results,

- Perform this task for all non-TPVM deployment models on a regular basis.
- Disable unattended upgrade services on external servers.

General Restrictions

This topic describes common restrictions on user privileges and host names.

User Privileges

To ensure correct installation of components, you must be a root user or have sudo privileges.

Host Names

- Ensure that the host names are unique and consist of lowercase alphanumeric characters and hyphens.

Do not use uppercase alphabetic characters.

- Ensure that the hyphens are the only special characters used.

No other special characters are allowed by the Kubernetes or K3s services for cluster formation.

NTP

Network Time Protocol (NTP) server is used to obtain the correct time from an external time source and adjust the local time in each connected device. When the NTP server functionality is enabled, the NTP server starts listening on the NTP port for client requests and responds with the reference time. The platform on which XCO is installed must use the same NTP or must be synchronized to the correct time and timezone.

Correct time and timezone ensure the following:

- Self-signed certificates have valid start and expiration times.
- XCO logs have the correct time stamp.
- K3s service starts without errors.

You can edit the `/etc/systemd/timesyncd.conf` file 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. Ensure that you select at least three NTP servers from the publicly available servers or the internal NTP servers.

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

## DNS

Ensure that the DNS configuration on the nodes is valid or the `/etc/resolv.conf` file is empty for correct functionality of the DNS resolution of Kubernetes.

Ensure that the `nslookup` returns the correct host name based on the IP address.

## IP Address Restrictions

Ensure that IPv4 and IPv6 addresses are configured on XCO management Interface. Adding an IPv4 address is mandatory during sub-interface creation; IPv6 addresses are optional.

Do NOT use the following IPv4 or IPv6 address subnets. These are either not supported or reserved for K3s:

- 10.42.0.0/16 subnet
- 10.43.0.0/16 subnet
- 169.254.0.0/16 subnet
- fd42::/48 subnet
- fd43::/112 subnet
- 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:10.10.10.10or ::ffff:0a0a:0a0a
```

- Do NOT use IPv6 Link Local addresses.

## TCP and UDP Ports for Single-Node Deployment

The following TCP and UDP ports are used in a single-node deployment:

Port	Protocol	Service	FW state	Communication Path
22	TCP	SSH	Open	SSH to XCO server, and Device Connection
49	TCP	TACACS+ Authentication requests	Open if using TACACS	XCO to TACACS server
80	TCP	HTTP requests	Open	API/Web
162	TCP	HTTPs requests	Open	XCO SNMP Notifications
443	TCP	HTTPs requests	Open	API/Web, Device Connection
514	UDP, TCP	Syslog service	Open	Device to XCO
389	TCP	LDAP Authentication requests	Open if using LDAP	XCO to LDAP server
3306	TCP	MariaDB port		Connection to Database
5672	TCP	RabbitMQ		RabbitMQ listening port
6443	TCP	K3s	Open	Kubernetes API server listening port
6514	TCP	Secure Syslog service	Open	Device to XCO
8078	TCP	XCO Monitoring service	Open	API/Web
8079	TCP	Host Authentication		
8080	TCP	HTTP requests	Open	API/Web
10010	TCP	Containerd service		
15672	UDP, TCP	RabbitMQ Management		API for monitoring RabbitMQ nodes and clusters
30085	UDP, TCP	OpenStack service		API for OpenStack integration
54322	TCP	SLXOS Telemetry	Open for Visibility deployment	SLX stats to XCO

## TCP and UDP Ports used for Multi-Node Deployment

The following TCP and UDP ports are used in a high availability multi-node deployment:

Port	Protocol	Service	Comment	Communication Path
53	TCP	Node local DNS for Kubernetes	Open	Node to Node
4567	UDP, TCP	Galera cluster replication	Open	Node to Node
4568	UDP, TCP	Galera incremental state transfer	Open	Node to Node
24007	UDP, TCP	GlusterFS daemon	Open	Node to Node
24008	UDP, TCP	GlusterFS management	Open	Node to Node
49152 through 49251	TCP	GlusterFS bricks	Open	Node to Node

## 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"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.1.x	More than 24	Yes	16.04, 18.04, and 20.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.2.x	More than 24	Yes	18.04 and 20.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 6: OVA Deployment Models**

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"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>
3.2.x	More than 24	Yes	18.04	<ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 64 GB</li> <li>• RAM: 8 GB</li> </ul>

**Table 7: TPVM Deployment Models**

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"> <li>• SLX 9150</li> <li>• SLX 9250</li> <li>• SLX 9740</li> </ul>	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"> <li>Extreme 8520</li> <li>Extreme 8720</li> </ul>				
3.0.x	<ul style="list-style-type: none"> <li>SLX 9150</li> <li>SLX 9250</li> <li>SLX 9740</li> <li>Extreme 8520</li> <li>Extreme 8720</li> </ul>	Up to 24	Yes	18.04	20.4.2
3.1.x	<ul style="list-style-type: none"> <li>SLX 9150</li> <li>SLX 9250</li> <li>SLX 9740</li> <li>Extreme 8520</li> <li>Extreme 8720</li> <li>Extreme 8820 (20.4.3 onwards only)</li> </ul>	Up to 24	Yes	18.04	20.4.2
3.2.x	<ul style="list-style-type: none"> <li>SLX 9150</li> <li>SLX 9250</li> <li>SLX 9740</li> <li>Extreme 8520</li> <li>Extreme 8720</li> <li>Extreme 8820 (20.4.3 onwards only)</li> </ul>	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
3.1.1	4.5.8	20.4.3

**Table 8: TPVM Software Support (continued)**

XCO Version	TPVM Version	SLX-OS Version
3.2.0	4.5.10	20.4.3a
3.2.1	4.5.12	20.5.1

**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"> <li>• CPU: 4 cores</li> <li>• Storage: 128 GB</li> <li>• RAM: 8 GB</li> </ul> Recommended: <ul style="list-style-type: none"> <li>• CPU: 16 cores</li> <li>• Storage: 200 GB</li> <li>• RAM: 32 GB</li> </ul>
3.2.0	18.04 and 20.04	Minimum: <ul style="list-style-type: none"> <li>• CPU: 4 cores</li> <li>• Storage: 128 GB</li> <li>• RAM: 8 GB</li> </ul> Recommended: <ul style="list-style-type: none"> <li>• CPU: 16 cores</li> <li>• Storage: 200 GB</li> <li>• RAM: 32 GB</li> </ul>

**Table 12: OVA Deployment Models**

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

**Table 12: OVA Deployment Models (continued)**

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

**Table 13: Supported Devices and Software**

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

## Pure IPv6 Support

XCO supports pure IPv6 mode where nodes are configured only with IPv6 addresses on their interfaces.

XCO installer supports IPv4, IPv6, and Dual mode IP deployments. The following table describes the installer changes during installation on both the suites:



**Note**

Fabric suite does not support pure IPv6 mode.

Suite Selected	Modes	Installer changes
Fabric	IPv4	No longer asks for optional IPv6 address. For example, Virtual IPv6. Optional sub-interface IPv6 subnet
	IPv6	<b>Installation exits or asks for restart with an appropriate error shown</b>
	Dual	No changes from previous releases
Packet	IPv4	Skips asking optional IPv6 address
	IPv6	All inputs on IPv6 address
	Dual	No changes from previous releases

## Limitations

1. For Pure IPv6 mode, a static default IPv6 route is required on all nodes.
2. Use of IPv4 mapped IPv6 address and IPv6 link local addresses is not supported.
3. Use of the link local addresses for installation input is not supported.
4. Upgrade from IPv4 mode or Dual mode to IPv6 mode is not supported and vice versa.

## Supported Platforms

1. TPVM (IPv4 and Dual mode)
2. OVA
3. Server

For more information about supported platforms, see [Supported Platforms and Deployment Models for Fabric Skill](#) on page 16 and [Supported Platforms and Deployment Models for Visibility Skill](#) on page 19.

## Pre-configuration for Selecting IP Mode

During a fresh installation, an appropriate IP mode (IPv4 single stack, IPv6 single stack, or Dual IP-stack) is selected as default in the IP configuration wizard. The IP mode selection is based on the IP addresses assigned to the interface. You can accept the default or change the deployment mode.

Ensure that the selected IP mode has an appropriate IP address configured on the VM Interface.

The following table describes the IP mode configuration on interface and the corresponding operational IP mode:

IP Mode Configuration on VM Interface	Operational IP Mode
If IPv4 address is configured and no IPv6 address is configured on gateway interface	IPv4 single-stack mode
If IPv4 address and IPv6 address with global scope are configured on gateway interface	Dual ip-stack mode
If no IPv4 address is configured on gateway interface and only IPv6 address with global scope is configured	IPv6 single-stack mode

During installation, system launches the following wizard to select a deployment mode. You can either accept or change the operational mode.



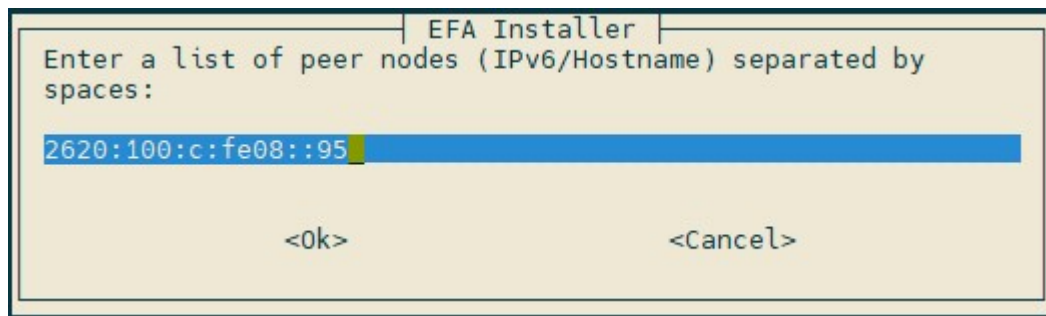
#### Note

- If you select **IPv4 single-stack**, wizards for an optional VIP6 and sub-interface optional IPv6 address will be skipped.
- If you select **IPv6 single-stack**, all inputs will be IPv6 address.
- If you select **Dual IP-stack**, wizard will be the same as the previous release. If you skip adding an optional VIP6, IP mode will change back to IPv4.

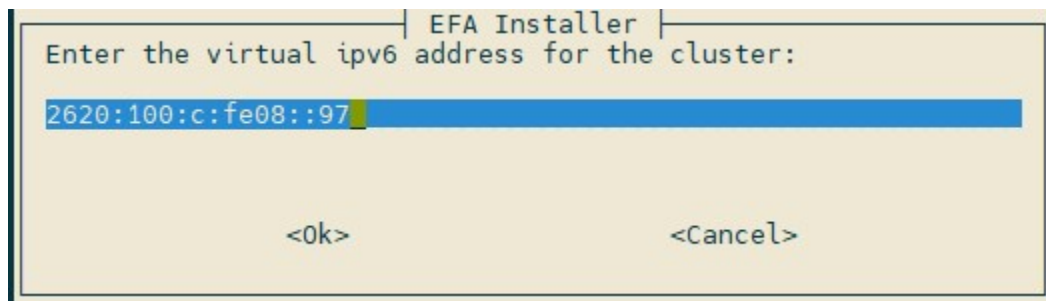
## Interactive Mode

The following snapshots show the changes in the deployment wizard if you select **IPv6 single-stack** as the operational IP mode for a multi-node deployment:

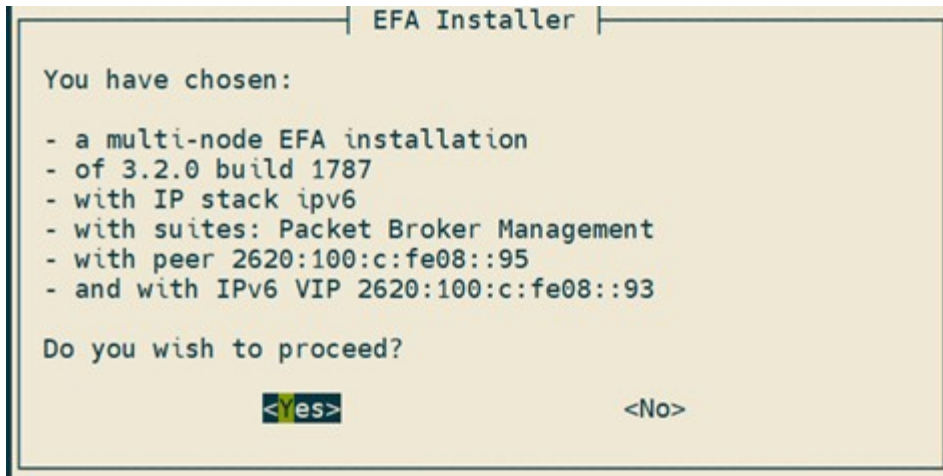
#### Peer IP Wizard



#### Virtual IP Wizard



### Installation Summary



## Non-Interactive Mode

An IPv6 parameter in single-node and multi-node deployment takes only IPv6 inputs for IPv6 single-stack deployment.

Legacy IPv4 or Dual mode will still use the existing mechanism. The installer checks for an optional VIP6 address, based on which operational IP mode is classified as “ipv4” or “dual”.



### Note

Fabric suite does not support pure IPv6 mode. An error on not supporting IPv6 mode appears when you select IPv6 single-stack mode when deploying fabric suite.

```

SLX# efa deploy non-interactive multi-node

Possible completions:
 ipv6 IPv6 single-stack Deployment
 package Path to EFA bundle
SLX#
SLX# efa deploy non-interactive single-node

Possible completions:
 ipv6 IPv6 single-stack Deployment
 package Path to EFA bundle
SLX#

```

The following example shows the use of IPv6 parameters in a non-interactive multi-node deployment:

```

efa deploy non-interactive multi-node ipv6 package /efaboot/efa-3.2.0-1.tar.gz peer-node
2000::1 vip6 2000:2 ping-target 3000:1 management-ip sub-interface-name sub1 sub-vlan-id
100 external-v6-subnet 2000:1::3/64

```

The following example shows the use of IPv6 parameters in a non-interactive single-node deployment:

```
efa deploy non-interactive single-node ipv6 package /efaboot/efa-3.2.0-1.tar.gz
management-ip sub-interface-name sub1 sub-vlan-id 100 external-v6-subnet 2000::1/64
```



### Note

For syntax and command examples, see [ExtremeCloud Orchestrator Command Reference, 3.2.1](#)

## Single CLI on Server for Single Node

XCO 3.2.0 introduces a new argument `--deploy-ipmode`. If the `--deploy-ipmode` argument is not used, it is assumed as “ipv4” or else you can mention “dual” or “ipv6”.

```
source deployment.sh --help

Usage:
source deployment.sh [options]

Options:
-G | --no-graphics
 Sets the procedure to not use graphical mode. The default is to use graphics.
-g | --graphics use_graphics
 Deprecated: Refer to no-graphics flag
 Sets the procedure to either use graphical mode or not. The supported
 values are 'yes' and 'no'. The default is 'yes'
-I | --non-interactive
 Sets the procedure to use non interactive mode. The default is interactive
 mode. In non-interactive mode, the procedure will not use graphics and the
 -G|-g option is overridden. In interactive mode, additional options other
 than graphics will not be considered and the user will be prompted to
 enter the required information.
-i | --interactive use_interactive
 Deprecated: Refer to non-interactive flag
 Sets the procedure to use interactive or non-interactive mode of
 installation. The supported values are 'yes' and 'no'.
 The default is 'yes'.
-o | --operation deployment_operation
```



```

Sets the value of current mode of operation. The supported values are
'deploy' and 'undeploy'. The operation has to be set to deploy
for a fresh installation or an upgrade to a new version
or to re-deploy the same version. The default is 'deploy'.

-v | --verbose

Send more info to the log during the procedure.

```

## OVA Support

XCO 3.2.0 supports the following two OVAs - each for fabric and packet suites:

- For fabric suites: xco-fabric-3.2.0.ova
- For packet suites: xco-vis-3.2.0.ova

OVAs contain dual IP mode to avoid a separate OVA for each IP mode. In post-boot configuration, at least one IP mode configuration is mandatory. You can either enter 'only IPv4' or 'only IPv6' (only for packet suite) or 'Both' (IPv4 and IPv6) IP addresses.

For post OVA installation procedure for fabric and visibility, see [Post OVA Install Procedure for Fabric and Visibility](#) on page 104.

## Upgrade Scenarios

XCO supports upgrades between similar IP modes.

- Upgrade from IPv4 to dual is incremental and supported if the node interface has the IPv6 address.
- For a multi-node deployment, the installer prompts for an optional Virtual IPv6 address. If you enter Virtual IPv6 address, the IP mode will be updated to Dual or else it remains as IPv4.
- For a single-node deployment, the IP mode will be updated to Dual if node interface has IPv6 address.
- IPv4 and Dual mode use IPv4 address as Peer IP and virtual IP. Similarly, IPv6 mode uses IPv6 addresses. The Inter IP mode upgrades are disruptive and not supported. System will show an appropriate error for invalid IP mode upgrades.

The following table describes the upgrade matrix for IP mode combinations during upgrades:

From	To	Support
Dual	IPv6	No
Dual	IPv4	No
IPv4	IPv6	No
IPv4	Dual	Yes

From	To	Support
IPv6	Dual	No
IPv6	IPv4	No
IPv4	IPv4	Yes
IPv6	IPv6	Yes
Dual	Dual	Yes

Use the following guidelines during an upgrade for best results:

From	To	Comments
IPv4	Dual	Add IPv6 address on node interfaces. Do not remove or change IPv4 addresses.
IPv6	IPv6	Do not remove or change IPv6 addresses on node interfaces.
Dual	Dual	Do not remove or change IPv4 and IPv6 addresses on node interfaces.
IPv4	IPv4	Do not remove or change IPv4 addresses on node interfaces.

## XCO Status

You can use the deployment IP mode logs to view the XCO status.

Ensure that the status of all the nodes are up.

### IPv6 Mode Logs

```
(efa:sbr)sbr@sbrefa1:~/builds/3.2.0/efa$ efa status
+-----+-----+-----+-----+
| Node Name | Role | Status | IP |
+-----+-----+-----+-----+
| sbrefa1 | active | up | 2620:100:c:fe08::96 |
+-----+-----+-----+-----+
| sbrefa2 | standby | up | 2620:100:c:fe08::95 |
+-----+-----+-----+-----+

efa:sbr)sbr@sbrefa1:~/builds/3.2.0/efa$ efa version
Version : 3.2.0
Build: 3100
Time Stamp: 22-12-09:11:06:01
Mode: Secure
Deployment Type: multi-node
Deployment IP Mode: ipv6
Deployment Platform: SERVER
Deployment Suite: Packet Broker Management
Virtual IPv6: 2620:100:c:fe08::97
Node IPv6s: 2620:100:c:fe08::96,2620:100:c:fe08::95
--- Time Elapsed: 4.114014ms ---
```

*Dual Mode Logs*

```
(efa:sbr)sbr@sbrefa1:~/builds/3.2.0/efa$ efa status
+-----+
| Node Name | Role | Status | IP |
+-----+
| sbrefa1 | active | up | 10.32.85.160 |
+-----+
| sbrefa2 | standby | up | 10.32.85.159 |
+-----+
--- Time Elapsed: 3.536148614s ---

efa:sbr)sbr@sbrefa1:~/builds/3.2.0/efa$ efa version
Version : 3.2.0
Build: 3100
Time Stamp: 22-12-09:11:06:01
Mode: Secure
Deployment Type: multi-node
Deployment IP Mode: dual
Deployment Platform: SERVER
Deployment Suite: Packet Broker Management
Virtual IP: 10.32.85.161
Node IPs: 10.32.85.160,10.32.85.159
--- Time Elapsed: 4.114014ms ---
```

*IPv4 Mode Logs*

```
(efa:sbr)sbr@sbrefa1:~/builds/3.2.0/efa$ efa status
+-----+
| Node Name | Role | Status | IP |
+-----+
| sbrefa1 | active | up | 10.32.85.160 |
+-----+
| sbrefa2 | standby | up | 10.32.85.159 |
+-----+
--- Time Elapsed: 3.536148614s ---

efa:sbr)sbr@sbrefa1:~/builds/3.2.0/efa$ efa version
Version : 3.2.0
Build: 3100
Time Stamp: 22-12-09:11:06:01
Mode: Secure
Deployment Type: multi-node
Deployment IP Mode: ipv4
Deployment Platform: SERVER
Deployment Suite: Packet Broker Management
Virtual IP: 10.32.85.161
Node IPs: 10.32.85.160,10.32.85.159
```



# TPVM Based Deployments

---

[TPVM Overview](#) on page 28

[TPVM Deployment Requirements](#) on page 29

[Redundant Management Network Overview](#) on page 29

[TPVM High-Availability Requirements](#) on page 34

[Install XCO Single-Node Fabric Management on TPVM](#) on page 37

[Install XCO on TPVM in a Multi-Node Deployment](#) on page 52

[Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#) on page 77

[Flexible XCO Deployment for TPVM](#) on page 78

Learn about the procedure of an XCO deployment on a Third-party virtual machine (TPVM).

## TPVM Overview

---

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

- 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 an XCO application on a TPVM using one SLX-OS command.
- Post XCO installation on the TPVM, use the command line Interface (CLI) to complete the administrative tasks. For a Web UI, install XCO on a server or VM as described in the [Server Deployments](#) on page 88.
- The XCO, SLX-OS, and TPVM application binary is shipped with the SLX devices. You can decouple XCO from SLX-OS. This allows you to upgrade XCO without a need to upgrade SLX-OS or the TPVM. You can deploy XCO on one of the SLX devices in the fabric.
- The XCO package tar.gz file are stored in the `/efaboot` directory. This is applicable to a fresh install or an upgrade of XCO. For an incremental upgrade of XCO, copy the XCO tar.gz file to the `/efaboot` directory on the SLX device before the deployment.
- With the TPVM 4.0.x and later versions, you can configure DNS, NTP, and LDAP when you deploy the TPVM. For more information, see "Guest OS for TPVM" in the [Extreme SLX-OS Management Configuration Guide](#).

## TPVM Deployment Requirements

---

Complete the following prerequisites before you deploy a TPVM:

- Run the **tpvm deploy** command on the SLX device to ensure that the TPVM is installed and running.
- Ensure that you specify the configuration of TPVM under the config mode.

The following 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
!
```

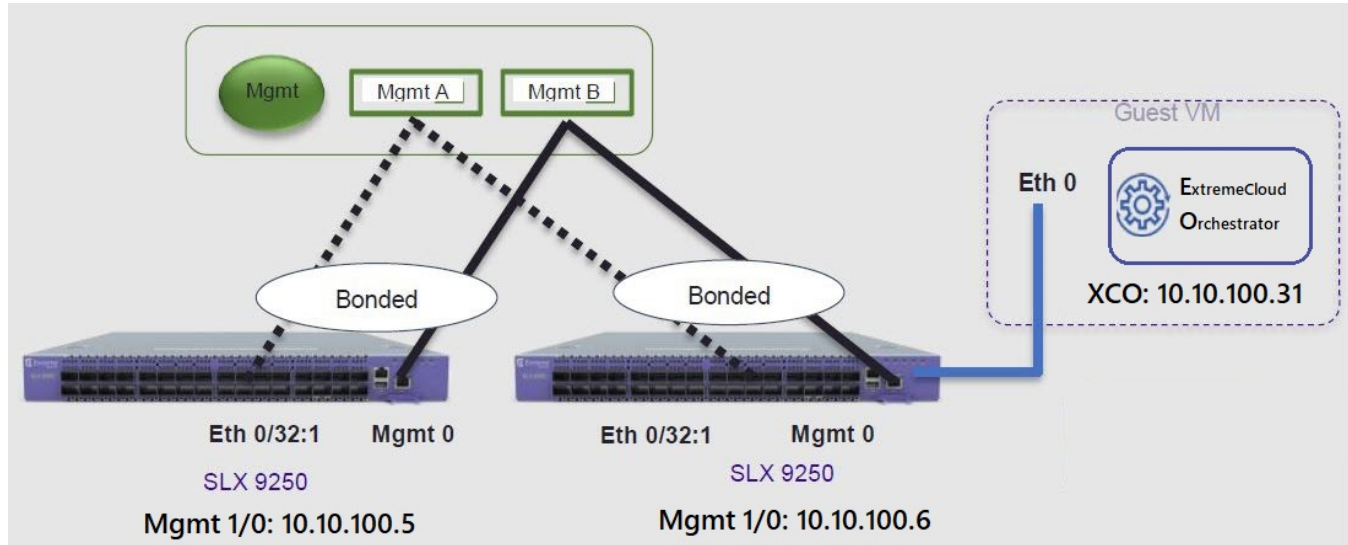
For more information on supported commands on SLX-OS, see the [Extreme SLX-OS Command Reference Guide](#).

For more information on TPVM deployment models, software support, IP fabric topology matrix, and XCO, neutron and SLX-OS compatibility, see [Supported Platforms and Deployment Models for Fabric Skill](#) on page 16

## 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 1: 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 29.

### Procedure

1. Enter global configuration mode.

```
device# configure terminal
```

2. Enter interface configuration mode.

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

3. Enable Redundant Management.

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

## Example

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

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

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

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

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(conf-if-eth-0/32:1)# redundant-management enable
device(conf-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 2: Data path overview**

## TPVM High-Availability Requirements

This topic describes the high availability requirements for XCO deployment on TPVM.

### OS Version

Ensure that all the nodes in a HA cluster have the same operating system version. For more information about supported operating systems, see [Supported Platforms and Deployment Models for Fabric Skill](#) on page 16.

## NTP

XCO installer allows a maximum drift of 10 seconds across the nodes. If the difference is more than 10 seconds, the installer prompts you to synchronize the clocks.

## DNS

Ensure that the DNS servers listed in the `/etc/resolv.conf` file resolve to the IP addresses of all the nodes. For example, ``dig <node_hostname> +short`` must resolve to the correct IP addresses assigned to the host.

## IP Addressing

Ensure that you have configured both IPv4 and IPv6 addresses in the XCO management Interface. If you are creating a sub-interface, the IPv4 address is mandatory; the IPv6 address is optional.

### *Additional IP Addresses*

High availability (HA) deployments require an extra IP address to be used as a virtual IP address for a cluster. Ensure that this extra IP address is an unallocated IP address in the same subnet. Use the virtual IP address to manage the XCO cluster in the HA deployment, and to maintain the HA using VRRP (Virtual Router Redundancy Protocol).

The following is an example of virtual IP address of a cluster:

- Cluster Virtual IP: 10.0.1.10/24
- Node A: 10.0.1.11/24
- Node B: 10.0.1.12/24

## SSH

Before installation, ensure that you have configured the passwordless SSH between the root users of both TPVM nodes.

### Configure Passwordless SSH for SLX-OS 20.2.3 and later with TPVM Version 4.2.2 and Later

You can configure passwordless SSH for SLX-OS and TPVM.

#### About This Task

Follow this procedure to configure passwordless SSH for SLX-OS and TPVM using the SLX command line.



#### Note

- The **tpvm config trusted-peer add** command updates the SSH configuration only for the root user of TPVM. For other users, see [Configure Passwordless SSH for SLX-OS Releases earlier than 20.2.3](#) on page 36 to manually transfer the SSH credentials for SLX-OS releases prior to 20.2.3.
- The **tpvm config trusted-peer add** command generates RSA keys. If your organization needs a different server key type, you must configure it manually using the `ssh-keygen`. Append the `.pub` file contents into the remote TPVM `/root/.ssh/authorized_keys` file.

#### Procedure

1. To configure a trusted peer, run the following command:

```
device# tpvm config trusted-peer add <peer-ip> sudo-user <peer-sudo-user> password
<peer-sudo-user-password>
```

2. To display a trusted peer information, run the following command:

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

3. To remove a trusted peer, run the following command:

```
device# tpvm config trusted-peer remove <peer-ip> sudo-user <peer-sudo-user> password
<peer-sudo-user-password>
```

### Configure Passwordless SSH for SLX-OS Releases earlier than 20.2.3

You can configure passwordless SSH for SLX-OS.

#### About This Task

Follow this procedure to configure passwordless SSH for SLX-OS releases earlier than 20.2.3.

#### Procedure

Run the **ssh-keygen** and **ssh-copy-id** commands on each TPVM to configure the SSH passwordless login. For example, SSH or console into TPVM1 as an admin user and complete the following steps:

- a. On TPVM1 and TPVM2, run the following command:

```
``bash
sudo mkdir -p /root/.ssh
sudo ssh-keygen -b 4096 -t rsa -q -N '' -f /root/.ssh/id_rsa
sudo cat /root/.ssh/id_rsa.pub
```

- b. On TPVM1, paste the contents of the **.pub** file from TPVM2 into the `/root/.ssh/authorized_keys` file.

- c. On TPVM2, paste the contents of the **.pub** file from TPVM1 into the `/root/.ssh/authorized_keys` file.
- d. Verify that the root from each TPVM can SSH into the root of the other TPVM with no password prompt.
- e. (Optional) On a Linux server, use the following script and pass the IP address of TPVM1 and TPVM2 as separate arguments:

You will be prompted for a password of each TPVM as it bootstraps.

```

```bash
#!/bin/bash
# Change this to the reference the appropriate local host public key for non-TPVM
linux.
MY_PUB_KEY=`-i ~/.ssh/id_rsa.pub`
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..."
ssh-copy-id $MY_PUB_KEY $SSH_OPTION $TPVM_USER@$TPVM1_IP
ssh-copy-id $MY_PUB_KEY $SSH_OPTION $TPVM_USER@$TPVM2_IP
echo "Generating ssh keypairs for root on TPVMs..."
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM1_IP "sudo mkdir -p /root/.ssh"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM2_IP "sudo mkdir -p /root/.ssh"
# Please note that you can change the key type, bits, and filename here, but the -N ''
should be left alone.
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM1_IP "sudo ssh-keygen -b 4096 -t rsa
-q -N '' -f /root/.ssh/id_rsa"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM2_IP "sudo ssh-keygen -b 4096 -t rsa
-q -N '' -f /root/.ssh/id_rsa"
echo "Setting up passwordless ssh login between TPVMs..."
TPVM1_ROOT_PUB_KEY=`ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM1_IP "sudo cat /
root/.ssh/id_rsa.pub"`
TPVM2_ROOT_PUB_KEY=`ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM2_IP "sudo cat /
root/.ssh/id_rsa.pub"`
echo "Exchanging ssh public keys for root between TPVMs..."
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM1_IP "echo $TPVM2_ROOT_PUB_KEY |
sudo tee -a /root/.ssh/authorized_keys"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM1_IP "echo $TPVM1_ROOT_PUB_KEY |
sudo tee -a /root/.ssh/authorized_keys"
echo "Adding TPVM IPs for root between TPVMs as known hosts to skip first time login
prompts..."
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM1_IP "sudo ssh-keyscan -H $TPVM2_IP
>> /root/.ssh/known_hosts"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $TPVM_USER@$TPVM2_IP "sudo ssh-keyscan -H $TPVM1_IP
>> /root/.ssh/known_hosts"
echo "Completed passwordless ssh login between TPVMs."
```

```

## Install XCO Single-Node Fabric Management on TPVM

You can install single-node XCO on TPVM.

### About This Task

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

## Procedure

1. Download the XCO tar file and the related digests file to a Linux server. Use these files to transfer the tar file to the SLX as a remote host.

```

```bash
remote-server$ ls
efa-3.2.0-digests.tar.gz  efa-3.2.0.tar.gz

```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Verify that the SLX-OS version, SSH keys, and passwordless access configuration are correct for the TPVM via the SLX console or SSH.

For the supported version information, see [Supported Platforms and Deployment Models for Fabric Skill](#) on page 16.

```

SLX# show tpvm status
SSH and Sudo passwordless      :Enabled
AutoStart                      :Enabled
Tpvm status                    :Running
Tpvm version                   :4.5.6
Tpvm additional status         :normal
device# show version
SLX-OS Operating System Version: 20.4.2
Copyright (c) 1995-2022 Extreme Networks, Inc.
Firmware name:                 20.4.2b
Build Time:                    04:24:00 Dec  6, 2022
Install Time:                  03:20:53 Dec 29, 2022
Kernel:                        4.14.67
Control Processor:             Intel(R) Xeon(R) CPU D-1527 @ 2.20GHz,  4 cores
Microcode Version:             0x7000017
Memory Size:                   System Total: 31643 MB
System Uptime:                 0days 0hrs 21mins 58secs

Name      Primary/Secondary Versions
-----
SLX-OS    20.4.2b
          20.4.2b

```

4. Verify the TPVM versions via the TPVM console or SSH.

```

tpvm# lsb_release -a
LSB Version:    core-5.0-noarch:core-5.0-x86_64
Distributor ID: slx
Description:    SLXOS 0.4
Release:        0.4
Codename:       n/a

```

5. Verify that NTP is synchronized by checking the Status line of **show tpvm config ntp** command.

```

SLX# show tpvm config ntp
NTP Servers: 10.20.61.191

Local time: Wed 2022-12-28 19:00:00 PST
Universal time: Thu 2022-12-29 03:00:00 UTC
RTC time: Thu 2022-12-29 03:00:01
Time zone: America/Los_Angeles (PST, -0800)
System clock synchronized: yes
systemd-timesyncd.service active: yes
RTC in local TZ: no

systemd-timesyncd.service - Network Time Synchronization
Loaded: loaded (/lib/systemd/system/systemd-timesyncd.service; enabled; vendor
preset: enabled)
Active: active (running) since Thu 2022-12-22 12:57:24 PST; 6 days ago

```

```
Docs: man:systemd-timesyncd.service(8)
Main PID: 680 (systemd-timesyn)
Status: "Synchronized to time server 10.20.61.191:123 (10.20.61.191)."
```

```
Tasks: 2 (limit: 4630)
CGroup: /system.slice/systemd-timesyncd.service
└─680 /lib/systemd/systemd-timesyncd
```

```
Dec 22 12:57:24 L5tpvm systemd[1]: Starting Network Time Synchronization...
Dec 22 12:57:24 L5tpvm systemd[1]: Started Network Time Synchronization.
Dec 22 12:57:57 L5tpvm systemd-timesyncd[680]: Synchronized to time server
10.20.61.191:123 (10.20.61.191).
```

6. Log in to TPVM and configure the NTP time zone. Press tab after timezone to see the list of available timezones.

```
SLX# tpvm config timezone America/Los_Angeles
```

7. Enter SLX Linux mode and transfer the tarball to the SLX using SCP.

```
SLX# start-shell
device$ scp user@remote-server:~/builds/efa/efa-3.2.0.tar.gz /efaboot/
```

8. To deploy XCO on TPVM from the SLX shell, run the **efa deploy** command.

If multiple XCO tarballs are present in the /efaboot/ directory, choose a target version of XCO to install from the version selection menu.



Note

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

```
SLX# efa deploy
Step 1: Get IP Address assigned to TPVM to deploy EFA 10.x.x.x.
Step 2: Checking for EFA packages in /efaboot directory
1. /efaboot/efa-3.2.0.tar.gz
2. /efaboot/efa-3.0.0.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                                         *
*****
Copying EFA package efa-3.2.0.tar.gz to TPVM 10.x.x.x... done.
Extracting EFA package efa-3.2.0.tar.gz on TPVM 10.x.x.x... done.
Starting EFA installer.
```


```
The EFA Installer will continue with a series of dialogs.
```


```

9. When prompted, select **Single-node deployment** then **OK**.



**Tip**

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

10. When prompted, select the IP mode (IPv4 single stack, or Dual ip-stack) and **OK**.



**Tip**

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

11. When prompted to configure additional management IP networks, take one of the following steps.
  - Select **No** to ignore this optional step or when you have finished entering sub-interface information.
  - 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 for IPv4 mode or IPv6 subnet address for IPv6 mode in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
    - In Dual mode, an IPv6 address is optional, but an IPv4 address is mandatory.

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

12. Verify the installation.
  - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
  - b. From the EFA command line on the TPVM console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.

Ensure that all the containers are in the `Running` state.

- c. Run the **efa status** command for the status information.

Ensure that the status of all the nodes are up.

```
SLX# show efa status
=====
 EFA version details
=====
Version : 3.2.0
Build: 45
Time Stamp: 23-03-16:23:17:04
Mode: Secure
Deployment Type: multi-node
Deployment Platform: TPVM
Deployment Suite: Fabric Automation
Deployment IP Mode: ipv4
Virtual IP: 10.20.55.4
Node IPs: 10.20.55.22,10.20.55.33
--- Time Elapsed: 13.291491ms ---
=====
 EFA Status
=====
+-----+-----+-----+-----+
| Node Name | Role | Status | IP |
+-----+-----+-----+-----+
| Node-22 | active | up | 10.20.55.22 |
+-----+-----+-----+-----+
| Node-33 | standby | up | 10.20.55.33 |
+-----+-----+-----+-----+
--- Time Elapsed: 18.632270244s ---

extreme@tvpm$ sudo efactl status
Node: efa
NAME READY STATUS RESTARTS AGE
```



| IP           | NODE | NOMINATED NODE | READINESS | GATES   |   |     |
|--------------|------|----------------|-----------|---------|---|-----|
| 10.42.194.71 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.74 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.68 | efa  | <none>         | 1/1       | Running | 0 | 13m |
| 10.20.229.54 | efa  | <none>         | 1/1       | Running | 0 | 11m |
| 10.42.194.69 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.70 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.72 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.76 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.75 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.20.229.54 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.42.194.77 | efa  | <none>         | 1/1       | Running | 0 | 12m |
| 10.20.229.54 | efa  | <none>         | 1/1       | Running | 0 | 11m |
| 10.42.194.73 | efa  | <none>         | 1/1       | Running | 0 | 12m |

| NAME                           | TYPE                       | CLUSTER-IP    | EXTERNAL-IP |
|--------------------------------|----------------------------|---------------|-------------|
| service/rabbitmq               | ClusterIP                  | 10.43.219.62  | <none>      |
| 15672/TCP,5672/TCP             | app=rabbitmq               |               |             |
| service/rabbitmq-local         | NodePort                   | 10.43.200.66  | <none>      |
| 15672:30673/TCP,5672:30672/TCP | app=rabbitmq               |               |             |
| service/goinventory-service    | ClusterIP                  | 10.43.96.100  | <none>      |
| 8082/TCP                       | app=goinventory-service    |               |             |
| service/efa-api-docs           | ClusterIP                  | 10.43.127.190 | <none>      |
| 80/TCP                         | app=efa-api-docs           |               |             |
| service/goraslog-service       | ClusterIP                  | 10.43.19.118  | <none>      |
| 8091/TCP                       | app=goraslog-service       |               |             |
| service/gonotification-service | ClusterIP                  | 10.43.11.74   | <none>      |
| 8088/TCP                       | app=gonotification-service |               |             |
| service/goauth-service         | ClusterIP                  | 10.43.133.78  | <none>      |
| 8080/TCP                       | app=goauth-service         |               |             |
| service/gorbac-service         | ClusterIP                  | 10.43.228.98  | <none>      |
| 8089/TCP                       | app=gorbac-service         |               |             |
| service/gofaultmanager-service | ClusterIP                  | 10.43.58.187  | <none>      |
| 8094/TCP                       | app=gofaultmanager-service |               |             |
| service/gosystem-service       | ClusterIP                  | 10.43.141.200 | <none>      |
| 8090/TCP                       | app=gosystem-service       |               |             |
| service/gofabric-service       | ClusterIP                  | 10.43.245.50  | <none>      |
| 8081/TCP                       | app=gofabric-service       |               |             |
| service/gopolicy-service       | ClusterIP                  | 10.43.36.27   | <none>      |
| 8093/TCP                       | app=gopolicy-service       |               |             |
| service/gotenant-service       | ClusterIP                  | 10.43.52.211  | <none>      |
| 8083/TCP                       | app=gotenant-service       |               |             |
| service/goopenstack-service    | ClusterIP                  | 10.43.26.123  | <none>      |
| 8085/TCP                       | app=goopenstack-service    |               |             |
| service/govcenter-service      | ClusterIP                  | 10.43.21.5    | <none>      |
| 8086/TCP                       | app=govcenter-service      |               |             |
| service/gohyperv-service       | ClusterIP                  | 10.43.43.180  | <none>      |
| 8087/TCP                       | app=gohyperv-service       |               |             |
| service/gosnmp-service         | ClusterIP                  | 10.43.78.109  | <none>      |
| 8092/TCP                       | app=gosnmp-service         |               |             |

```

NAME
TO-DATE AVAILABLE NODE SELECTOR DESIRED CURRENT READY UP-
AGE CONTAINERS
IMAGES
 daemonset.apps/goopenstack-service 0 0 0
0 0 non-existing=true 12m openstack
goopenstack:3.2.0 app=goopenstack-service
 daemonset.apps/efa-api-docs 1 1 1
1 1 <none> 12m efa-api-docs efa-api-
docs:3.2.0 app=efa-api-docs
 daemonset.apps/govcenter-service 0 0 0
0 0 non-existing=true 12m vcenter
govcenter:3.2.0 app=govcenter-service
 daemonset.apps/gohyperv-service 0 0 0
0 0 non-existing=true 12m hyperv
gohyperv:3.2.0 app=gohyperv-service
 daemonset.apps/gosystem-service 1 1 1
1 1 <none> 12m gosystem
gosystem:3.2.0 app=gosystem-service
 daemonset.apps/rabbitmq 1 1 1
1 1 <none> 13m rabbitmq-node
rabbitmq:3.2.0 app=rabbitmq
 daemonset.apps/gonotification-service 1 1 1
1 1 <none> 12m gonotification-service
gonotification:3.2.0 app=gonotification-service
 daemonset.apps/goinventory-service 1 1 1
1 1 <none> 12m goinventory-service
goinventory:3.2.0 app=goinventory-service
 daemonset.apps/goauth-service 1 1 1
1 1 <none> 12m go-auth
goauth:3.2.0 app=goauth-service
 daemonset.apps/gofaultmanager-service 1 1 1
1 1 <none> 12m gofaultmanager-service
gofaultmanager:3.2.0 app=gofaultmanager-service
 daemonset.apps/gofabric-service 1 1 1
1 1 <none> 12m gofabric-service
gofabric:3.2.0 app=gofabric-service
 daemonset.apps/gopolicy-service 1 1 1
1 1 <none> 12m gopolicy-service
gopolicy:3.2.0 app=gopolicy-service
 daemonset.apps/goraslog-service 1 1 1
1 1 <none> 12m goraslog-service
goraslog:3.2.0 app=goraslog-service
 daemonset.apps/gotenant-service 1 1 1
1 1 <none> 12m gotenant-service
gotenant:3.2.0 app=gotenant-service
 daemonset.apps/gosnmp-service 1 1 1
1 1 <none> 12m gosnmp-service
gosnmp:3.2.0 app=gosnmp-service
 daemonset.apps/gorbac-service 1 1 1
1 1 <none> 12m go-rbac
gorbac:3.2.0 app=gorbac-service
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:extreme)extreme@tpvm:~$ efa status
+-----+
| Node Name | Role | Status | IP |
+-----+
| tpvm-xco-229-54 | active | up | 10.20.229.54 |
+-----+
--- Time Elapsed: 4.656427295s ---
`--`

```

## Back up Database

You can back up an XCO database after the installation.

### About This Task

Follow this procedure to back up an XCO database after single-node installation of XCO on TPVM.

### Procedure

1. Connect to the TPVM console or SSH session as an Extreme user.
2. Log in to EFA using the **efa login** command.
3. (Optional) Run the **efa system settings update** command to configure XCO to backup to a remote server.

```
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.412206025s ---
(efa:user)@tpvm:~$ efa system settings update --remote-server-ip 10.25.101.74 --
remote-transfer-protocol scp --remote-server-username user --remote-server-directory /
home/user/
Please supply a password for remote server:
Setting Update Successful
--- Time Elapsed: 4.723485224s ---
```

4. If you have configured a remote server, run the **efa system backup --remote** command to generate a local backup, and then copy the resulting tar.gz file to the configured remote device. Otherwise, run the **efa system backup** command.

Backups are stored locally in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` directory. Backup files are date and time stamped in the file name.



### Note

Running the **efa system restore** command without any arguments will run in an interactive mode, and allow you to pick any backup present in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` directory.

For information on additional backup options, see [ExtremeCloud Orchestrator Command Reference, 3.2.1](#) and [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1](#).

### Example

The following example shows **efa system backup** command output with backup location:

```
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.412206025s ---
(efa:extreme)extreme@tpvm:~$ efa system backup --remote
Generating backup of EFA...

Backup Location: user@10.25.101.74:/home/releaseuser/user/
EFA-3.2.0-54-2023-03-28T03-02-13.077.tar
--- Time Elapsed: 20.097407523s ---
```

## Restore Database

You can restore an XCO database from its backup

### About This Task

Follow this procedure to restore an XCO database.

### Procedure

1. Connect to the TPVM console or SSH session as a root user.
2. Log in to EFA using the **efa login** command.
3. (Optional) Transfer the backup from the remote host.
4. Run the **efa system restore** command with the required arguments.



### Note

Running the **efa system restore** command without any arguments will run the command in an interactive mode, and allow you to pick any backup present in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` directory.

For more information, see [ExtremeCloud Orchestrator Command Reference, 3.2.1](#) and [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1](#).

### Example

The following example shows **efa system restore** command output from a backup:

```
user@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.019739819s ---
(efa:user)user@tpvm:~$ scp user@remotehost:~/EFA-3.2.0-54-2023-03-28T03-02-13.077.tar ./
user@remotehosts password:
EFA-3.2.0-54-2023-03-28T03-02-13.077.tar 100% 876KB 28.8MB/s 00:00
(efa:user)user@tpvm:~$ efa system restore --backup-tar
EFA-3.2.0-54-2023-03-28T03-02-13.077.tar
Performing EFA restore using EFA-3.2.0-54-2023-03-28T03-02-13.077.tar
Restore operation ID: 4d229051-87db-11ed-8921-5254008962d2
Stopping all EFA services...
All pods are terminated
Restoring databases...
Start services after restore
Restore of encryption keys is completed
Extreme Fabric Automation Stack is now Restored and Ready!
Restore operation is successful
--- Time Elapsed: 6m47.278493209s ---
(efa:user)user@tpvm:~$
```

## Upgrade XCO on TPVM in a Single-Node Deployment

You can upgrade single-node XCO on TPVM.

### About This Task

Follow this procedure to upgrade XCO in TPVM in a single-node deployment.

## Procedure

1. Download the XCO tar file and the digests file to a Linux server. Use these files to transfer the tar file to the SLX as a remote host.

```

```bash
remote-server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz

```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. While the upgrade process does make a backup of the XCO state, as a best practice, perform a manual backup before every upgrade to create a remote copy of the backup.
4. Enter SLX Linux mode, and transfer the tarball to the SLX using SCP.



Note

You can copy multiple XCO tarballs to the `/efaboot/` directory.

```

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

```

5. To deploy XCO on TPVM from the SLX command line, run the **efa deploy** command.

If you have multiple XCO tarballs in the `/efaboot/` directory, choose a target version of XCO to install from the version selection menu.



Note

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

```

SLX# efa deploy
Step 1: Get IP Address assigned to TPVM to deploy EFA 10.x.x.x.
Step 2: Checking for EFA packages in /efaboot directory
1. /efaboot/efa-3.1.0.tar.gz
2. /efaboot/efa-3.0.0.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                                        *
*****
Copying EFA package efa-3.1.0.tar.gz to TPVM 10.x.x.x... done.
Extracting EFA package efa-3.1.0.tar.gz on TPVM 10.x.x.x... done.
Starting EFA installer.

The EFA Installer will continue with a series of dialogs.

```

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



Tip

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

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 for IPv4 mode or IPv6 subnet address for IPv6 mode in CIDR format. The subnet must not overlap with any IP subnet that you have already provided.
 - In Dual mode, 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.

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

8. Verify the installation.
 - a. On the SLX device, run the **show efa status** command to see details of the installation and the state of services.
 - b. From the EFA command line on the TPVM console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services. Ensure that all the containers are in the "Running" state.
 - c. Run the **show efa status** command to view the status information. Ensure that the status of all the nodes are up.

```
SLX# show efa status
=====
                        EFA version details
=====
Version : 3.1.0
Build: GA
Time Stamp: 22-11-10:17:11:29
Mode: Secure
Deployment Type: single-node
Deployment Platform: TPVM
Deployment Suite: Fabric Automation
Node IP: 10.20.229.54
--- Time Elapsed: 848.174µs ---

=====
                        EFA Status
=====

+-----+-----+-----+-----+
| Node Name      | Role   | Status | IP           |
+-----+-----+-----+-----+
| tpvm-xco-229-54 | active | up     | 10.20.229.54 |
+-----+-----+-----+-----+
--- Time Elapsed: 6.260854749s ---

extreme@tvpm$ sudo efactl status
Node: efa
NAME                READY   STATUS    RESTARTS   AGE
IP                  NODE    NOMINATED NODE  READINESS GATES
pod/efa-api-docs-sbg9x    1/1     Running   0           12m
```

10.42.194.71	efa	<none>	<none>		
	pod/gosystem-service-m7245		1/1	Running	0 12m
10.42.194.74	efa	<none>	<none>		
	pod/rabbitmq-cf57z		1/1	Running	0 13m
10.42.194.68	efa	<none>	<none>		
	pod/gonotification-service-5nhqb		1/1	Running	0 11m
10.20.229.54	efa	<none>	<none>		
	pod/goinventory-service-m8fgm		1/1	Running	0 12m
10.42.194.69	efa	<none>	<none>		
	pod/goauth-service-qr7zs		1/1	Running	0 12m
10.42.194.70	efa	<none>	<none>		
	pod/gofaultmanager-service-gszrm		1/1	Running	0 12m
10.42.194.72	efa	<none>	<none>		
	pod/gofabric-service-wrl6z		1/1	Running	0 12m
10.42.194.76	efa	<none>	<none>		
	pod/gopolicy-service-4llmn		1/1	Running	0 12m
10.42.194.75	efa	<none>	<none>		
	pod/goraslog-service-2w8hj		1/1	Running	0 12m
10.20.229.54	efa	<none>	<none>		
	pod/gotenant-service-zhbp2		1/1	Running	0 12m
10.42.194.77	efa	<none>	<none>		
	pod/gosnmp-service-5jnvq		1/1	Running	0 11m
10.20.229.54	efa	<none>	<none>		
	pod/gorbac-service-jnttb		1/1	Running	0 12m
10.42.194.73	efa	<none>	<none>		
	NAME		TYPE	CLUSTER-IP	EXTERNAL-IP
	PORT(S)	AGE	SELECTOR		
	service/rabbitmq		ClusterIP	10.43.219.62	<none>
15672/TCP,5672/TCP		13m	app=rabbitmq		
	service/rabbitmq-local		NodePort	10.43.200.66	<none>
15672:30673/TCP,5672:30672/TCP		13m	app=rabbitmq		
	service/goinventory-service		ClusterIP	10.43.96.100	<none>
8082/TCP		12m	app=goinventory-service		
	service/efa-api-docs		ClusterIP	10.43.127.190	<none>
80/TCP		12m	app=efa-api-docs		
	service/goraslog-service		ClusterIP	10.43.19.118	<none>
8091/TCP		12m	app=goraslog-service		
	service/gonotification-service		ClusterIP	10.43.11.74	<none>
8088/TCP		12m	app=gonotification-service		
	service/goauth-service		ClusterIP	10.43.133.78	<none>
8080/TCP		12m	app=goauth-service		
	service/gorbac-service		ClusterIP	10.43.228.98	<none>
8089/TCP		12m	app=gorbac-service		
	service/gofaultmanager-service		ClusterIP	10.43.58.187	<none>
8094/TCP		12m	app=gofaultmanager-service		
	service/gosystem-service		ClusterIP	10.43.141.200	<none>
8090/TCP		12m	app=gosystem-service		
	service/gofabric-service		ClusterIP	10.43.245.50	<none>
8081/TCP		12m	app=gofabric-service		
	service/gopolicy-service		ClusterIP	10.43.36.27	<none>
8093/TCP		12m	app=gopolicy-service		
	service/gotenant-service		ClusterIP	10.43.52.211	<none>
8083/TCP		12m	app=gotenant-service		
	service/goopenstack-service		ClusterIP	10.43.26.123	<none>
8085/TCP		12m	app=goopenstack-service		
	service/govcenter-service		ClusterIP	10.43.21.5	<none>
8086/TCP		12m	app=govcenter-service		
	service/gohyperv-service		ClusterIP	10.43.43.180	<none>
8087/TCP		12m	app=gohyperv-service		
	service/gosnmp-service		ClusterIP	10.43.78.109	<none>
8092/TCP		12m	app=gosnmp-service		
	NAME	DESIRED	CURRENT	READY	UP-

```

TO-DATE    AVAILABLE    NODE SELECTOR    AGE    CONTAINERS
IMAGES
          daemonset.apps/goopenstack-service    0    0    0
0          0          non-existing=true    12m    openstack
goopenstack:3.1.0    app=goopenstack-service
          daemonset.apps/efa-api-docs    1    1    1
1          1          <none>    12m    efa-api-docs    efa-api-
docs:3.1.0    app=efa-api-docs
          daemonset.apps/govcenter-service    0    0    0
0          0          non-existing=true    12m    vcenter
govcenter:3.1.0    app=govcenter-service
          daemonset.apps/gohyperv-service    0    0    0
0          0          non-existing=true    12m    hyperv
gohyperv:3.1.0    app=gohyperv-service
          daemonset.apps/gosystem-service    1    1    1
1          1          <none>    12m    gosystem
gosystem:3.1.0    app=gosystem-service
          daemonset.apps/rabbitmq    1    1    1
1          1          <none>    13m    rabbitmq-node
rabbitmq:3.1.0    app=rabbitmq
          daemonset.apps/gonotification-service    1    1    1
1          1          <none>    12m    gonotification-service
gonotification:3.1.0    app=gonotification-service
          daemonset.apps/goinventory-service    1    1    1
1          1          <none>    12m    goinventory-service
goinventory:3.1.0    app=goinventory-service
          daemonset.apps/goauth-service    1    1    1
1          1          <none>    12m    go-auth
goauth:3.1.0    app=goauth-service
          daemonset.apps/gofaultmanager-service    1    1    1
1          1          <none>    12m    gofaultmanager-service
gofaultmanager:3.1.0    app=gofaultmanager-service
          daemonset.apps/gofabric-service    1    1    1
1          1          <none>    12m    gofabric-service
gofabric:3.1.0    app=gofabric-service
          daemonset.apps/gopolicy-service    1    1    1
1          1          <none>    12m    gopolicy-service
gopolicy:3.1.0    app=gopolicy-service
          daemonset.apps/goraslog-service    1    1    1
1          1          <none>    12m    goraslog-service
goraslog:3.1.0    app=goraslog-service
          daemonset.apps/gotenant-service    1    1    1
1          1          <none>    12m    gotenant-service
gotenant:3.1.0    app=gotenant-service
          daemonset.apps/gosnmp-service    1    1    1
1          1          <none>    12m    gosnmp-service
gosnmp:3.1.0    app=gosnmp-service
          daemonset.apps/gorbac-service    1    1    1
1          1          <none>    12m    go-rbac
gorbac:3.1.0    app=gorbac-service
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:extreme)extreme@tpvm:~$ efa status
+-----+-----+-----+-----+
| Node Name | Role | Status | IP |
+-----+-----+-----+-----+
| tpvm-xco-229-54 | active | up | 10.20.229.54 |
+-----+-----+-----+-----+
--- Time Elapsed: 4.656427295s ---

```


Recover the Upgrade

You can recover from an upgrade failure by rerunning the upgrade, or performing a fresh installation. Then use a backup to restore the system.

Rerunning the Upgrade

You can rerun an XCO upgrade on TPVM.

For more information on rerunning an XCO upgrade, see [Upgrade XCO on TPVM in a Single-Node Deployment](#) on page 44.

Uninstall, Install, and Restore from Backup

You can perform uninstall, install, and restore process.

About This Task

Complete the following tasks to uninstall, install, and restore from XCO backup.

Procedure

1. [Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#) on page 77.
2. [Install XCO Single-Node Fabric Management on TPVM](#) on page 37.
3. [Restore Database](#) on page 44.

Upgrade TPVM OS for XCO Single-Node Deployment

You can upgrade the TPVM version of an XCO single-node deployment.

- If the TPVM version is 4.4.0 or above, see [TPVM Incremental Upgrade](#) on page 69.
- If the TPVM version is lower than 4.4.0, see [Upgrading SLX-OS, TPVM, and XCO Together](#) on page 58.

However, you can migrate the TPVM from a single-node to multi-node deployment. For more information, see [Migrate TPVM Single-Node to Multi-Node](#) on page 49.

Migrate TPVM Single-Node to Multi-Node

You can migrate the SLX TPVM from a single-node to multi-node deployment.

Before You Begin

Ensure that both TPVM instances are correctly configured and meet the high availability multi-node requirements.

About This Task

Follow this procedure to migrate TPVM from a single-node to multi-node deployment.

Procedure

1. Download the XCO tar file and digest file to the Linux server. Use these files to transfer the tar file to the SLX as a remote host.

```

```bash
remote-server$ ls
efa-3.1.0-digests.tar.gz efa-3.1.0.tar.gz

```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. (Optional) Backup the XCO state. While the upgrade process does make a backup of the XCO state, as a best practice, perform a manual backup before every upgrade to get a remote copy of the backup.
4. Enter SLX Linux mode, and transfer the tarball to the SLX using SCP.



### Note

You can copy multiple XCO tarballs to the `/efaboot/` directory.

```

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

```

5. To deploy XCO on TPVM from the SLX command line, run the **efa deploy** command.

If you have multiple XCO tarballs present in the `/efaboot/` directory, choose a target version of XCO to install from the version selection menu.



### Note

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

```

SLX# efa deploy
Step 1: Get IP Address assigned to TPVM to deploy EFA 10.x.x.x.
Step 2: Checking for EFA packages in /efaboot directory
1. /efaboot/efa-3.1.0.tar.gz
2. /efaboot/efa-3.0.0.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 *

Copying EFA package efa-3.1.0.tar.gz to TPVM 10.x.x.x... done.
Extracting EFA package efa-3.1.0.tar.gz on TPVM 10.x.x.x... done.
Starting EFA installer.
```

```

The EFA Installer will continue with 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.
 - 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 **efact1 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.

Install XCO on TPVM in a Multi-Node Deployment

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 16.

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 16.
 - 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.

Back up XCO Database

You can back up an XCO database after the installation.

About This Task

Follow this procedure to back up an XCO database.

Procedure

1. Connect to the TPVM console or SSH session of the redundant IP as an Extreme user.
2. Log in to EFA using the **efa login** command.
3. If needed, run the **efa system settings update** command to configure XCO to back up to a remote server.

```
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.412206025s ---
(efa:user)@tpvm:~$ efa system settings update --remote-server-ip 10.25.101.74 --
remote-transfer-protocol scp --remote-server-username user --remote-server-directory /
home/user/
Please supply a password for remote server:
Setting Update Successful
--- Time Elapsed: 4.723485224s ---
```

- If you have configured a remote server, run the **efa system backup --remote** command to generate a local backup and copy the resulting tar.gz file to the configured remote device.

Otherwise, run the **efa system backup** command. Local backups are stored in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` directory. Backup files are date and time stamped in the file name.



Note

For additional backup options, see the [ExtremeCloud Orchestrator Command Reference, 3.2.1](#) and [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1](#)

Example

The following example shows **efa system backup** command output with backup location:

```
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.412206025s ---
(efa:extreme)extreme@tpvm:~$ efa system backup --remote
Generating backup of EFA...

Backup Location: user@10.25.101.74:/home/user/EFA-3.1.0-2-2022-12-30T01-44-39.193.tar

--- Time Elapsed: 26.700226408s ---
```

Restore Database

You can restore an XCO database from its backup.

About This Task

Follow this procedure to restore an XCO database.

Procedure

- Connect to the TPVM console or SSH session as a root user.
- Log in to EFA using the **efa login** command.
- (Optional) Transfer the backup from the remote host.
- Run the **efa system restore** command with the required arguments.



Note

Running the **efa system restore** command without any arguments will run in an interactive mode. This will allow you to pick any backup present in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` directory. For more information, see [ExtremeCloud Orchestrator Command Reference, 3.2.1](#) and [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1](#).

Example

```
user@tpvm:~$ efa login
Password:
```

```

Login successful.
--- Time Elapsed: 2.019739819s ---
(efa:user)user@tpvm:~$ scp user@remotehost:~/EFA-3.1.0-2-2022-12-30T01-44-39.193.tar ./
user@remotehosts password:
EFA-3.1.0-2-2022-12-30T01-44-39.193.tar          100% 876KB 28.8MB/s 00:00
(efa:user)user@tpvm:~$ efa system restore --backup-tar
EFA-3.1.0-2-2022-12-30T01-44-39.193.tar
Performing EFA restore using EFA-3.1.0-2-2022-12-30T01-44-39.193.tar
Restore operation ID: 4d229051-87db-11ed-8921-5254008962d2
Stopping all EFA services...
All pods are terminated
Restoring databases...
Start services after restore
Restore of encryption keys is completed
Extreme Fabric Automation Stack is now Restored and Ready!
Restore operation is successful
--- Time Elapsed: 6m47.278493209s ---
(efa:user)user@tpvm:~$

```

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 16. 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.

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 16.

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 [Back up XCO Database](#) on page 54 and [Restore Database](#) on page 55.

- 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 73.

- 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 16.

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, 3.2.1](#).

- 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 the standby node of XCO. For more information, see [TPVM Complete Package Upgrade](#) on page 63.

- 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 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)

- 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 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).

- 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 60.

- 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
```

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, 3.2.1](#)

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.

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 Con fig | 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 file. The encoded password string will still configure the SLX TPVM properly. If no password is set then default "password" is used. |
| Interface management
<ul style="list-style-type: none"> • ip • gw | Pre-start only | String | <ul style="list-style-type: none"> • DHCP or IPv4 address • IPv4 address | 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 the active node is expecting the peer node to be configured with a specific IP address. You cannot change the peer node IP without restarting XCO HA cluster daemons on the active node. |

| TPVM Con fig | SLX Command Execution Stage | Type | Value | Description |
|---|-----------------------------|---------|--|---|
| Interface insight
<ul style="list-style-type: none"> • ipv4 • gw | Pre-start only | String | <ul style="list-style-type: none"> • DHCP or IPv4 address • IPv4 address | |
| Host name | 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 | |
| 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 Con fig 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.
 - c. 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-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 --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
```

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| | Version|
+-----+-----+-----+-----+-----+-----+
| 10.24.80.56| 10.24.80.180| node180 | 20.4.2slxos20.4.2_220614_1000 | 4.5.0 |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58| 10.24.80.181| node181 | 20.4.1 | 4.5.0 |
+-----+-----+-----+-----+-----+-----+
    
```

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		State
Version	State							
Type								
+-----+-----+-----+-----+-----+-----+-----+-----+-----+								
+-----+-----+-----+-----+-----+-----+-----+-----+-----+								
1.1.1.1	SLX	4001	BR	4200000000	Leaf	4.5.5		In
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
Required	Waiting for TPVM		Incremental	2022-12-10	2022-12-10			
								Progress
Instance|Instance to come up|           |Upgrade  |19:36:05 -0800 PST| 19:40:31 -0800 PST|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
```

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 [TPVM High-Availability Requirements](#) on page 34.
- 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

1. 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
```

2. 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.

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 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, 3.2.1](#).

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 **efact1 status** to verify that all pods on the active node are in Running state.
 - c. Run **efact1 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 52 .

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>
```

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.

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 | <code>#efa deploy non-interactive multi-node package <packagename> peer-node <ipaddress> vip4 <ip-address> ping-target <clear <ip address1>, [ip address2]></code> |
| With sub-interface | <code>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></code>

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 82

[Enable or Disable Services](#) on page 83

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.
Spine1#
```



Server Deployments

- [Server Requirements](#) on page 88
- [Single-Node Deployment](#) on page 92
- [High-Availability Multi-Node Deployment](#) on page 117
- [Upgrade Ubuntu on the XCO Host - Single Node or Multi Node](#) on page 125
- [Upgrade XCO from a Single-Node to a Multi-Node Deployment](#) on page 127
- [Replace a Node in a Multi-node Deployment](#) on page 129
- [Rollback](#) on page 130
- [Uninstall XCO in a Single-Node or Multi-Node Deployment](#) on page 130

Learn about the procedure of an XCO deployment on Server.

Server Requirements

You can install XCO on an Ubuntu Server with appropriate resources for managing IP fabric or network visibility devices. Each management application will have an OVA appliance.

The XCO installation supports CLI and a Web UI.

Server and Cloud VM Requirements

The following table describes minimum and recommended hardware requirements for Server deployments:

Table 14: General Hardware Requirements

| Minimum Requirements | Recommended Requirements |
|---|---|
| OS: Ubuntu 18.04.6 LTS Server (AMD64),
Ubuntu 20.04 LTS Server (AMD64) | OS: Ubuntu 18.04.6 LTS Server (AMD64),
Ubuntu 20.04 LTS Server (AMD64) |
| CPU: 4 cores | CPU: 16 cores |

Table 14: General Hardware Requirements (continued)

| Minimum Requirements | Recommended Requirements |
|----------------------|--------------------------|
| Storage: 128GB | Storage: 200GB |
| RAM: 8GB | RAM: 32GB |

**Note**

Check whether `/usr/local/bin` is added in the following secure path:

Defaults

```
secure_path="/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:/snap/bin"
```

If not, add `/usr/local/bin` and then redeploy.

Server and Cloud VM Deployment Models

The following tables provide software and hardware matrix for IP fabric, SLX-OS, and Visibility Management for Server deployments.

IP Fabric Topology and SLX-OS Release 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 | 20.3.x, 20.4.x | ✓ | | | ✓ | ✓ |

Visibility Management

| Supported Devices | Supported Software Versions |
|-------------------------------|---------------------------------|
| NetIron MLX and XMR platforms | NetIron 6.3.00 and its patches |
| SLX 9140 | SLX-OS 18s.1.03 e and f patches |
| SLX 9240 | TierraOS 21.1.2.x |
| Extreme 9920 | |

High-Availability Requirements

Learn about the high availability requirements for Server deployments.

OS Version

Ensure that all nodes in a high-availability cluster have the same version of the operating system. For more information about supported operating systems, see [Server and Cloud VM Requirements](#) on page 88.

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 the clocks.

DNS

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` must resolve to the exact IP addresses assigned to the host.

IP Addressing

Ensure that you have configured both IPv4 and IPv6 addresses on the XCO management Interfaces. Adding an IPv6 address is optional, but an IPv4 address is mandatory.

Additional IP Addresses

High availability (HA) deployments require an extra IP address to be used as a virtual IP address for a cluster. Ensure that this extra IP address is an unallocated IP address in the same subnet. Use the virtual IP address to manage the XCO cluster in the HA deployment, and to maintain the HA using VRRP (Virtual Router Redundancy Protocol).

The following is an example of virtual IP address of a cluster:

- Cluster Virtual IP: 10.0.1.10/24
- Node A: 10.0.1.11/24
- Node B: 10.0.1.12/24

Configure Passwordless SSH

You can configure passwordless SSH for both the nodes in XCO.

About This Task

Follow this procedure to configure the passwordless SSH between the root users of both nodes prior to installation.

Procedure

1. Run the **ssh-keygen** and **ssh-copy-id** commands on each node to configure the SSH passwordless login. For example, SSH or console into the nodes as an admin user, then run the following commands:

```
sudo sed -i "s/#PermitRootLogin\ prohibit-password/PermitRootLogin\ prohibit-
password/g" /etc/ssh/sshd_config
sudo mkdir -p /root/.ssh
sudo ssh-keygen -b 4096 -t rsa -q -N '' -f /root/.ssh/id_rsa
sudo cat /root/.ssh/id_rsa.pub
```

2. On Node1, paste the contents of the .pub file from Node2 into /root/.ssh/authorized_keys directory. On Node2, paste the contents of the .pub file from Node1 into /root/.ssh/authorized_keys directory.
3. Verify that the root from each node can SSH to the root of the other node with no password prompt.
4. Run the following script and pass the IP address of Node1 and Node2 as separate arguments:

You will be prompted for the password of each node after it bootstraps.



Note

Modify the script to suit your requirements.

```
#!/bin/bash
# Change this to the reference the appropriate local host public key for non-TPVM
Linux.
MY_PUB_KEY=`-i ~/.ssh/id_rsa.pub`
NODE1_IP="$1"
NODE2_IP="$2"
NODE_USER="extreme"
SSH_OPTION="-o StrictHostKeyChecking=no"
echo "Setting up passwordless ssh login from this host to nodes..."
ssh-copy-id $MY_PUB_KEY $SSH_OPTION $NODE_USER@$NODE1_IP
ssh-copy-id $MY_PUB_KEY $SSH_OPTION $NODE_USER@$NODE2_IP
echo "Generating ssh keypairs for root on nodes..."
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE1_IP "sudo mkdir -p /root/.ssh"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE2_IP "sudo mkdir -p /root/.ssh"
# Please note that you can change the key type, bits, and filename here, but the -N ''
should be left alone.
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE1_IP "sudo ssh-keygen -b 4096 -t rsa
-q -N '' -f /root/.ssh/id_rsa"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE2_IP "sudo ssh-keygen -b 4096 -t rsa
-q -N '' -f /root/.ssh/id_rsa"
echo "Setting up passwordless ssh login between nodes..."
NODE1_ROOT_PUB_KEY=`ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE1_IP "sudo cat /
root/.ssh/id_rsa.pub"`
NODE2_ROOT_PUB_KEY=`ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE2_IP "sudo cat /
root/.ssh/id_rsa.pub"`
echo "Exchanging ssh public keys for root between nodes..."
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE1_IP "echo $NODE2_ROOT_PUB_KEY |
sudo tee -a /root/.ssh/authorized_keys"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE1_IP "echo $NODE1_ROOT_PUB_KEY |
sudo tee -a /root/.ssh/authorized_keys"
echo "Adding node IPs for root between nodes as known hosts to skip first time login
prompts..."
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE1_IP "sudo ssh-keyscan -H $NODE2_IP
>> /root/.ssh/known_hosts"
ssh ${MY_PUB_KEY::-4} $SSH_OPTION $NODE_USER@$NODE2_IP "sudo ssh-keyscan -H $NODE1_IP
```

```
>> /root/.ssh/known_hosts"
echo "Completed passwordless ssh login between nodes."
```

Single-Node Deployment

You can install XCO on a single-node server or virtual machine, which is a non-TPVM deployment.

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, 3.2.1](#).

| 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 Fabric using Interactive Mode

You can install XCO using interactive mode.

About This Task

Follow this procedure to install XCO in interactive mode.

Procedure

1. Download the XCO tar file and the digests file to the Linux server. XCO will be installed to this server using an administrative account.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow the administrative user to run commands.

4. Change the current directory to the extracted file location, and run **source deployment.sh**. This starts the XCO installation in a menu driven interactive mode. For command line based interactive mode, run the **source deployment.sh -g no** command.
5. When prompted, select **Single-node deployment** and **Fabric Automation**.

**Tip**

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

```
admin@server:~/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 1 Selected
application suite: Fabric Automation
```

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

**Tip**

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

7. When prompted to configure additional management IP networks, take one of the following steps.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.
 - Select **Yes** and then provide the requested information.
 - 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.

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

8. Verify the installation.
 - From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
 - Run the **efa status** command for concise status information.

Ensure that the status of all the nodes are up.

```
admin@server$ sudo efactl status
Node: efa
NAME                                READY   STATUS    RESTARTS   AGE   IP
NODE  NOMINATED NODE  READINESS GATES
pod/efa-api-docs-sbg9x              1/1     Running   0           12m   10.42.194.71
efa   <none>         <none>
pod/gosystem-service-m7245         1/1     Running   0           12m   10.42.194.74
efa   <none>         <none>
pod/rabbitmq-cf57z                  1/1     Running   0           13m   10.42.194.68
efa   <none>         <none>
```

| | | | | | |
|------------------------------------|-------------------|----------------------------|---------------|--------|-------------------|
| pod/gonotification-service-5nhqb | 1/1 | Running | 0 | 11m | 10.20.229.54 |
| efa <none> | <none> | | | | |
| pod/goinventory-service-m8fgm | 1/1 | Running | 0 | 12m | 10.42.194.69 |
| efa <none> | <none> | | | | |
| pod/goauth-service-qr7zs | 1/1 | Running | 0 | 12m | 10.42.194.70 |
| efa <none> | <none> | | | | |
| pod/gofaultmanager-service-gszrm | 1/1 | Running | 0 | 12m | 10.42.194.72 |
| efa <none> | <none> | | | | |
| pod/gofabric-service-wrl6z | 1/1 | Running | 0 | 12m | 10.42.194.76 |
| efa <none> | <none> | | | | |
| pod/gopolicy-service-4llmn | 1/1 | Running | 0 | 12m | 10.42.194.75 |
| efa <none> | <none> | | | | |
| pod/goraslog-service-2w8hj | 1/1 | Running | 0 | 12m | 10.20.229.54 |
| efa <none> | <none> | | | | |
| pod/gotenant-service-zhbp2 | 1/1 | Running | 0 | 12m | 10.42.194.77 |
| efa <none> | <none> | | | | |
| pod/gosnmp-service-5jnqv | 1/1 | Running | 0 | 11m | 10.20.229.54 |
| efa <none> | <none> | | | | |
| pod/gorbac-service-jnttb | 1/1 | Running | 0 | 12m | 10.42.194.73 |
| efa <none> | <none> | | | | |
| NAME | | TYPE | CLUSTER-IP | | EXTERNAL-IP |
| PORT(S) | | AGE | SELECTOR | | |
| service/rabbitmq | | ClusterIP | 10.43.219.62 | <none> | 15672/ |
| TCP,5672/TCP | 13m | app=rabbitmq | | | |
| service/rabbitmq-local | | NodePort | 10.43.200.66 | <none> | 15672:30673/ |
| TCP,5672:30672/TCP | 13m | app=rabbitmq | | | |
| service/goinventory-service | | ClusterIP | 10.43.96.100 | <none> | |
| 8082/TCP | 12m | app=goinventory-service | | | |
| service/efa-api-docs | | ClusterIP | 10.43.127.190 | <none> | |
| 80/TCP | 12m | app=efa-api-docs | | | |
| service/goraslog-service | | ClusterIP | 10.43.19.118 | <none> | |
| 8091/TCP | 12m | app=goraslog-service | | | |
| service/gonotification-service | | ClusterIP | 10.43.11.74 | <none> | |
| 8088/TCP | 12m | app=gonotification-service | | | |
| service/goauth-service | | ClusterIP | 10.43.133.78 | <none> | |
| 8080/TCP | 12m | app=goauth-service | | | |
| service/gorbac-service | | ClusterIP | 10.43.228.98 | <none> | |
| 8089/TCP | 12m | app=gorbac-service | | | |
| service/gofaultmanager-service | | ClusterIP | 10.43.58.187 | <none> | |
| 8094/TCP | 12m | app=gofaultmanager-service | | | |
| service/gosystem-service | | ClusterIP | 10.43.141.200 | <none> | |
| 8090/TCP | 12m | app=gosystem-service | | | |
| service/gofabric-service | | ClusterIP | 10.43.245.50 | <none> | |
| 8081/TCP | 12m | app=gofabric-service | | | |
| service/gopolicy-service | | ClusterIP | 10.43.36.27 | <none> | |
| 8093/TCP | 12m | app=gopolicy-service | | | |
| service/gotenant-service | | ClusterIP | 10.43.52.211 | <none> | |
| 8083/TCP | 12m | app=gotenant-service | | | |
| service/goopenstack-service | | ClusterIP | 10.43.26.123 | <none> | |
| 8085/TCP | 12m | app=goopenstack-service | | | |
| service/govcenter-service | | ClusterIP | 10.43.21.5 | <none> | |
| 8086/TCP | 12m | app=govcenter-service | | | |
| service/gohyperv-service | | ClusterIP | 10.43.43.180 | <none> | |
| 8087/TCP | 12m | app=gohyperv-service | | | |
| service/gosnmp-service | | ClusterIP | 10.43.78.109 | <none> | |
| 8092/TCP | 12m | app=gosnmp-service | | | |
| NAME | | DESIRED | CURRENT | READY | UP-TO-DATE |
| AVAILABLE | NODE SELECTOR | AGE | CONTAINERS | IMAGES | |
| SELECTOR | | | | | |
| daemonset.apps/goopenstack-service | | 0 | 0 | 0 | 0 |
| 0 | non-existing=true | 12m | openstack | | goopenstack:3.1.0 |
| app=goopenstack-service | | | | | |

```

daemonset.apps/efa-api-docs      1      1      1      1
1      <none>      12m    efa-api-docs      efa-api-docs:3.1.0
app=efa-api-docs
daemonset.apps/govcenter-service  0      0      0      0
0      non-existing=true 12m    vcenter      govcenter:3.1.0
app=govcenter-service
daemonset.apps/gohyperv-service   0      0      0      0
0      non-existing=true 12m    hyperv      gohyperv:3.1.0
app=gohyperv-service
daemonset.apps/gosystem-service   1      1      1      1
1      <none>      12m    gosystem      gosystem:3.1.0
app=gosystem-service
daemonset.apps/rabbitmq          1      1      1      1
1      <none>      13m    rabbitmq-node  rabbitmq:3.1.0
app=rabbitmq
daemonset.apps/gonotification-service 1      1      1      1
1      <none>      12m    gonotification-service  gonotification:3.1.0
app=gonotification-service
daemonset.apps/goinventory-service  1      1      1      1
1      <none>      12m    goinventory-service  goinventory:3.1.0
app=goinventory-service
daemonset.apps/goauth-service      1      1      1      1
1      <none>      12m    go-auth      goauth:3.1.0
app=goauth-service
daemonset.apps/gofaultmanager-service 1      1      1      1
1      <none>      12m    gofaultmanager-service  gofaultmanager:3.1.0
app=gofaultmanager-service
daemonset.apps/gofabric-service     1      1      1      1
1      <none>      12m    gofabric-service  gofabric:3.1.0
app=gofabric-service
daemonset.apps/gopolicy-service     1      1      1      1
1      <none>      12m    gopolicy-service  gopolicy:3.1.0
app=gopolicy-service
daemonset.apps/goraslog-service     1      1      1      1
1      <none>      12m    goraslog-service  goraslog:3.1.0
app=goraslog-service
daemonset.apps/gotenant-service     1      1      1      1
1      <none>      12m    gotenant-service  gotenant:3.1.0
app=gotenant-service
daemonset.apps/gosnmp-service       1      1      1      1
1      <none>      12m    gosnmp-service  gosnmp:3.1.0
app=gosnmp-service
daemonset.apps/gorbac-service       1      1      1      1
1      <none>      12m    go-rbac      gorbac:3.1.0
app=gorbac-service
admin@server:~$ efa login
Password:
Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:admin)admin@server:~$ efa status
+-----+-----+-----+
| Node Name      | Role   | Status | IP          |
+-----+-----+-----+
| server        | active | up     | 10.20.229.54 |
+-----+-----+-----+
--- Time Elapsed: 4.656427295s ---

```

Install Fabric using Single Command

You can install XCO using non-interactive mode.

About This Task

Follow this procedure to install XCO using single command in a non-interactive mode.

Procedure

1. Download the XCO tar file and the digests file to the Linux server. XCO is installed on this server using an administrative account.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.
4. Change the current directory to the extracted file location, and run the **source deployment.sh -i no -deploy-suite fabric** command. This starts the XCO installation in a non-interactive mode with the options of single-node, IP fabric management, and no additional management IP addresses.

To specify an additional management interface for XCO, wait for the installation to complete, un-install XCO, and then re-install using the interactive method.

5. Verify the installation.
 - From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
 - Run the **efa status** command for the status information.

Ensure that the status of all the nodes are up.

```
admin@server$ sudo efactl status
Node: efa
NAME                                READY  STATUS   RESTARTS  AGE  IP
NODE  NOMINATED NODE  READINESS GATES
pod/efa-api-docs-sbg9x              1/1    Running  0          12m  10.42.194.71
efa  <none>  <none>
pod/gosystem-service-m7245          1/1    Running  0          12m  10.42.194.74
efa  <none>  <none>
pod/rabbitmq-cf57z                  1/1    Running  0          13m  10.42.194.68
efa  <none>  <none>
pod/gonotification-service-5nhqb    1/1    Running  0          11m  10.20.229.54
efa  <none>  <none>
pod/goinventory-service-m8fgm       1/1    Running  0          12m  10.42.194.69
efa  <none>  <none>
pod/goauth-service-qr7zs            1/1    Running  0          12m  10.42.194.70
efa  <none>  <none>
pod/gofaultmanager-service-gszrm    1/1    Running  0          12m  10.42.194.72
efa  <none>  <none>
pod/gofabric-service-wrl6z          1/1    Running  0          12m  10.42.194.76
efa  <none>  <none>
pod/gopolicy-service-4llmn          1/1    Running  0          12m  10.42.194.75
efa  <none>  <none>
pod/goraslog-service-2w8hj          1/1    Running  0          12m  10.20.229.54
efa  <none>  <none>
pod/gotenant-service-zhbp2          1/1    Running  0          12m  10.42.194.77
efa  <none>  <none>
pod/gosnmp-service-5jmqv            1/1    Running  0          11m  10.20.229.54
```



```

efa      <none>          <none>
pod/gorbac-service-jnttb      1/1      Running    0          12m      10.42.194.73
efa      <none>          <none>

NAME                                TYPE           CLUSTER-IP      EXTERNAL-IP
PORT(S)                             AGE    SELECTOR
service/rabbitmq                    ClusterIP      10.43.219.62    <none>      15672/
TCP, 5672/TCP                        13m    app=rabbitmq
service/rabbitmq-local              NodePort       10.43.200.66    <none>      15672:30673/
TCP, 5672:30672/TCP                 13m    app=rabbitmq
service/goinventory-service          ClusterIP      10.43.96.100    <none>
8082/TCP                             12m    app=goinventory-service
service/efa-api-docs                ClusterIP      10.43.127.190   <none>
80/TCP                                12m    app=efa-api-docs
service/goraslog-service             ClusterIP      10.43.19.118    <none>
8091/TCP                              12m    app=goraslog-service
service/gonotification-service       ClusterIP      10.43.11.74     <none>
8088/TCP                              12m    app=gonotification-service
service/goauth-service               ClusterIP      10.43.133.78    <none>
8080/TCP                              12m    app=goauth-service
service/gorbac-service               ClusterIP      10.43.228.98    <none>
8089/TCP                              12m    app=gorbac-service
service/gofaultmanager-service       ClusterIP      10.43.58.187    <none>
8094/TCP                              12m    app=gofaultmanager-service
service/gosystem-service             ClusterIP      10.43.141.200   <none>
8090/TCP                              12m    app=gosystem-service
service/gofabric-service             ClusterIP      10.43.245.50    <none>
8081/TCP                              12m    app=gofabric-service
service/gopolicy-service             ClusterIP      10.43.36.27     <none>
8093/TCP                              12m    app=gopolicy-service
service/gotenant-service             ClusterIP      10.43.52.211    <none>
8083/TCP                              12m    app=gotenant-service
service/goopenstack-service          ClusterIP      10.43.26.123    <none>
8085/TCP                              12m    app=goopenstack-service
service/govcenter-service            ClusterIP      10.43.21.5      <none>
8086/TCP                              12m    app=govcenter-service
service/gohyperv-service             ClusterIP      10.43.43.180    <none>
8087/TCP                              12m    app=gohyperv-service
service/gosnmp-service               ClusterIP      10.43.78.109    <none>
8092/TCP                              12m    app=gosnmp-service

NAME                                DESIRED    CURRENT    READY    UP-TO-DATE
AVAILABLE  NODE SELECTOR  AGE    CONTAINERS  IMAGES
SELECTOR
daemonset.apps/goopenstack-service  0          0          0        0
0          non-existing=true 12m    openstack   goopenstack:3.1.0
app=goopenstack-service
daemonset.apps/efa-api-docs         1          1          1        1
1          <none>          12m    efa-api-docs efa-api-docs:3.1.0
app=efa-api-docs
daemonset.apps/govcenter-service     0          0          0        0
0          non-existing=true 12m    vcenter     govcenter:3.1.0
app=govcenter-service
daemonset.apps/gohyperv-service      0          0          0        0
0          non-existing=true 12m    hyperv      gohyperv:3.1.0
app=gohyperv-service
daemonset.apps/gosystem-service      1          1          1        1
1          <none>          12m    gosystem    gosystem:3.1.0
app=gosystem-service
daemonset.apps/rabbitmq              1          1          1        1
1          <none>          13m    rabbitmq-node rabbitmq:3.1.0
app=rabbitmq
daemonset.apps/gonotification-service 1          1          1        1
1          <none>          12m    gonotification-service gonotification:3.1.0

```

```

app=gonotification-service
daemonset.apps/goinventory-service      1      1      1      1
1      <none>      12m    goinventory-service      goinventory:3.1.0
app=goinventory-service
daemonset.apps/goauth-service           1      1      1      1
1      <none>      12m    go-auth                  goauth:3.1.0
app=goauth-service
daemonset.apps/gofaultmanager-service   1      1      1      1
1      <none>      12m    gofaultmanager-service   gofaultmanager:3.1.0
app=gofaultmanager-service
daemonset.apps/gofabric-service         1      1      1      1
1      <none>      12m    gofabric-service         gofabric:3.1.0
app=gofabric-service
daemonset.apps/gopolicy-service         1      1      1      1
1      <none>      12m    gopolicy-service         gopolicy:3.1.0
app=gopolicy-service
daemonset.apps/goraslog-service         1      1      1      1
1      <none>      12m    goraslog-service         goraslog:3.1.0
app=goraslog-service
daemonset.apps/gotenant-service         1      1      1      1
1      <none>      12m    gotenant-service         gotenant:3.1.0
app=gotenant-service
daemonset.apps/gosnmp-service           1      1      1      1
1      <none>      12m    gosnmp-service           gosnmp:3.1.0
app=gosnmp-service
daemonset.apps/gorbac-service           1      1      1      1
1      <none>      12m    go-rbac                  gorbac:3.1.0
app=gorbac-service
admin@server:~$ efa login
Password:
Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:admin)admin@server:~$ efa status
+-----+-----+-----+-----+
| Node Name | Role | Status | IP |
+-----+-----+-----+-----+
| server    | active | up | 10.20.229.54 |
+-----+-----+-----+-----+
--- Time Elapsed: 4.656427295s ---

```

Install Fabric using OVA

You can install XCO using OVA.

Before You Begin

- Ensure that you have an understanding of the hypervisor environment, including converting the embedded disk image to a suitable format.
- The XCO OVA image does not support cloud-init, or provisioning the VM with a user and network information via the user interface. Ensure that you are aware of the information on DNS, network address, subnet mask, and gateway requirements to assign to the VM.
- DHCP address of the VM is not supported. A static address is required. Ensure that the selected address complies with the general IP address requirements.

About This Task

Follow this procedure to install XCO using OVA.

Procedure

1. Download the OVA file for fabric management to a Linux device. Use the device to convert or transfer the OVA file to your hypervisor image store.
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. (Optional) Convert the OVA, and import the image or the OVA into your hypervisor.
4. Deploy the OVA or converted image.

When prompted, select virtio or its equivalent on your hypervisor for the drive bus type option for a converted image. If given options are in the form of hardware requirements assigned to the VM, see [Server and Cloud VM Requirements](#) on page 88.

5. When the VM is in a running state, complete the procedure described in [Post OVA Install Procedure for Fabric and Visibility](#) on page 104.

Install Visibility using Interactive Mode

You can install visibility using an interactive mode.

About This Task

Follow this procedure to install visibility in interactive mode.

Procedure

1. Download the XCO tar file and the related digests file to the Linux server. XCO will be installed to this server using an administrative account.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.
4. Change the current directory to the extracted file location, and run the **source deployment.sh**. This starts the XCO installation in a menu driven interactive mode.

For command line based interactive mode, run the **source deployment.sh -g no** command.
5. When prompted, select **Single-node deployment** and **Packet Broker Management**.



Tip

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

```
admin@server:~/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
2
Selected application suite: Packet Broker Management
```

- 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 to configure additional management IP networks, take one of the following steps.
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.
 - Select **Yes** and then provide the requested information.
 - 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.

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

- Verify the installation.
 - From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
 - Run the **efa status** command to ensure that the status of all the nodes are up.

Install Visibility using Single Command

You can install visibility using a single command.

About This Task

Follow this procedure to install visibility using a single command in a non-interactive mode.

Procedure

- Download the XCO tar file and the related digests file to the Linux server. XCO will be installed to this server using an administrative account.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

- Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
- Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.
- Change the current directory to the extracted file location, and run the **source deployment.sh -i no -deploy-suite packet**. This starts the XCO installation in a non-interactive mode with the options of single-node, IP visibility management, and no additional management IP addresses.

To specify an additional management interface for XCO, wait for the installation to complete, un-install XCO, and then re-install using the interactive method.

5. Verify the installation.

- From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
- Run the **efa status** command to ensure that the status of all the nodes are up.

```
admin@server$ sudo efactl status
Node: efa
NAME                READY   STATUS    RESTARTS   AGE   IP
NODE   NOMINATED NODE   READINESS GATES
pod/efa-api-docs-sbg9x          1/1     Running   0           12m   10.42.194.71
efa    <none>          <none>
pod/gosystem-service-m7245     1/1     Running   0           12m   10.42.194.74
efa    <none>          <none>
pod/rabbitmq-cf57z             1/1     Running   0           13m   10.42.194.68
efa    <none>          <none>
pod/gonotification-service-5nhqb 1/1     Running   0           11m   10.20.229.54
efa    <none>          <none>
pod/goinventory-service-m8fgm   1/1     Running   0           12m   10.42.194.69
efa    <none>          <none>
pod/goauth-service-qr7zs       1/1     Running   0           12m   10.42.194.70
efa    <none>          <none>
pod/gofaultmanager-service-gszrm 1/1     Running   0           12m   10.42.194.72
efa    <none>          <none>
pod/gofabric-service-wrl6z     1/1     Running   0           12m   10.42.194.76
efa    <none>          <none>
pod/gopolicy-service-4llmn     1/1     Running   0           12m   10.42.194.75
efa    <none>          <none>
pod/goraslog-service-2w8hj     1/1     Running   0           12m   10.20.229.54
efa    <none>          <none>
pod/gotenant-service-zhbp2     1/1     Running   0           12m   10.42.194.77
efa    <none>          <none>
pod/gosnmp-service-5jnqv       1/1     Running   0           11m   10.20.229.54
efa    <none>          <none>
pod/gorbac-service-jnttb      1/1     Running   0           12m   10.42.194.73
efa    <none>          <none>

NAME                TYPE           CLUSTER-IP      EXTERNAL-IP
PORT(S)            AGE   SELECTOR
service/rabbitmq    ClusterIP      10.43.219.62    <none>      15672/
TCP,5672/TCP        13m   app=rabbitmq
service/rabbitmq-local NodePort       10.43.200.66    <none>      15672:30673/
TCP,5672:30672/TCP 13m   app=rabbitmq
service/goinventory-service ClusterIP      10.43.96.100    <none>
8082/TCP            12m   app=goinventory-service
service/efa-api-docs ClusterIP      10.43.127.190   <none>
80/TCP             12m   app=efa-api-docs
service/goraslog-service ClusterIP      10.43.19.118    <none>
8091/TCP           12m   app=goraslog-service
service/gonotification-service ClusterIP      10.43.11.74     <none>
8088/TCP           12m   app=gonotification-service
service/goauth-service ClusterIP      10.43.133.78    <none>
8080/TCP           12m   app=goauth-service
service/gorbac-service ClusterIP      10.43.228.98    <none>
8089/TCP           12m   app=gorbac-service
service/gofaultmanager-service ClusterIP      10.43.58.187    <none>
8094/TCP           12m   app=gofaultmanager-service
service/gosystem-service ClusterIP      10.43.141.200   <none>
8090/TCP           12m   app=gosystem-service
service/gofabric-service ClusterIP      10.43.245.50    <none>
8081/TCP           12m   app=gofabric-service
service/gopolicy-service ClusterIP      10.43.36.27     <none>
8093/TCP           12m   app=gopolicy-service
```

```

service/gotenant-service      ClusterIP 10.43.52.211 <none>
8083/TCP                      12m app=gotenant-service
service/goopenstack-service   ClusterIP 10.43.26.123 <none>
8085/TCP                      12m app=goopenstack-service
service/govcenter-service     ClusterIP 10.43.21.5 <none>
8086/TCP                      12m app=govcenter-service
service/gohyperv-service      ClusterIP 10.43.43.180 <none>
8087/TCP                      12m app=gohyperv-service
service/gosnmp-service        ClusterIP 10.43.78.109 <none>
8092/TCP                      12m app=gosnmp-service

NAME                               DESIRED  CURRENT  READY  UP-TO-DATE
AVAILABLE  NODE SELECTOR  AGE  CONTAINERS  IMAGES
SELECTOR
daemonset.apps/goopenstack-service 0         0         0         0
0         non-existing=true 12m  openstack  goopenstack:3.1.0
app=goopenstack-service
daemonset.apps/efa-api-docs        1         1         1         1
1         <none>          12m  efa-api-docs  efa-api-docs:3.1.0
app=efa-api-docs
daemonset.apps/govcenter-service    0         0         0         0
0         non-existing=true 12m  vcenter      govcenter:3.1.0
app=govcenter-service
daemonset.apps/gohyperv-service     0         0         0         0
0         non-existing=true 12m  hyperv      gohyperv:3.1.0
app=gohyperv-service
daemonset.apps/gosystem-service     1         1         1         1
1         <none>          12m  gosystem    gosystem:3.1.0
app=gosystem-service
daemonset.apps/rabbitmq             1         1         1         1
1         <none>          13m  rabbitmq-node  rabbitmq:3.1.0
app=rabbitmq
daemonset.apps/gonotification-service 1         1         1         1
1         <none>          12m  gonotification-service  gonotification:3.1.0
app=gonotification-service
daemonset.apps/goinventory-service  1         1         1         1
1         <none>          12m  goinventory-service  goinventory:3.1.0
app=goinventory-service
daemonset.apps/goauth-service       1         1         1         1
1         <none>          12m  go-auth      goauth:3.1.0
app=goauth-service
daemonset.apps/gofaultmanager-service 1         1         1         1
1         <none>          12m  gofaultmanager-service  gofaultmanager:3.1.0
app=gofaultmanager-service
daemonset.apps/gofabric-service     1         1         1         1
1         <none>          12m  gofabric-service  gofabric:3.1.0
app=gofabric-service
daemonset.apps/gopolicy-service     1         1         1         1
1         <none>          12m  gopolicy-service  gopolicy:3.1.0
app=gopolicy-service
daemonset.apps/goraslog-service     1         1         1         1
1         <none>          12m  goraslog-service  goraslog:3.1.0
app=goraslog-service
daemonset.apps/gotenant-service     1         1         1         1
1         <none>          12m  gotenant-service  gotenant:3.1.0
app=gotenant-service
daemonset.apps/gosnmp-service       1         1         1         1
1         <none>          12m  gosnmp-service  gosnmp:3.1.0
app=gosnmp-service
daemonset.apps/gorbac-service       1         1         1         1
1         <none>          12m  go-rbac      gorbac:3.1.0
app=gorbac-service
admin@server:~$ efa login
Password:

```

```

Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:admin)admin@server:~$ efa status
+-----+-----+-----+-----+
| Node Name      | Role   | Status | IP           |
+-----+-----+-----+-----+
| server         | active | up     | 10.20.229.54 |
+-----+-----+-----+-----+
--- Time Elapsed: 4.656427295s ---

```

Install Visibility using OVA

You can install visibility manager using OVA.

Before You Begin

- Ensure that you have an understanding of the hypervisor environment, including converting the embedded disk image to a suitable format.
- The XCO OVA image does not support cloud-init, or provisioning the VM with a user and network information via the user interface. Ensure that you are aware of the information on DNS, network address, subnet mask, and gateway requirements to assign to the VM.
- DHCP address of the VM is not supported. A static address is required. Ensure that the selected address complies with the general IP address requirements.
- Ensure that you use the VNC console. Directly attached serial consoles is not supported for logging in to the VM for the post boot procedure. .

About This Task

Follow this procedure to install visibility manager using OVA.

Procedure

1. Download the OVA file for visibility management to a Linux device. Convert or transfer the OVA file to your hypervisors image store.
2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. (Optional) Convert the OVA, and import the image or the OVA into your hypervisor.
4. Deploy the OVA or converted image.

When prompted, select virtio or its equivalent on your hypervisor for the drive bus type option for a converted image. If given options are in the form of hardware requirements assigned to the VM, see [Server and Cloud VM Requirements](#) on page 88.

5. When the VM is in a running state, complete the procedure described in [Post OVA Install Procedure for Fabric and Visibility](#) on page 104.

Post OVA Install Procedure for Fabric and Visibility

The OVA images use a first run script to assign the static IP address and deploy XCO Fabric Manager and XCO Visibility Manager. The script is located at `~/ .startup.sh` for the ubuntu user.

About This Task

Follow this procedure to complete the post OVA installation procedure for fabric and visibility.



Note

As a best practice, change the passwords of the root and ubuntu user of the OVA.

Procedure

1. Log in to the VNC console of the VM deployed from the OVA.
2. Log in to the VM using the Ubuntu user with the following credential:
 - User name/Password: ubuntu/ubuntu

The following welcome screen is displayed. Press **Enter**.

```

=====
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:**

```

A menu opens, with the following choices:

```

=====
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 User Password Configuration
=====
The root user password is currently set for this appliance.
=====

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

4. Press 2 for network settings.

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

Enter selection :
To exit screen, press CTRL-C.
```

5. Press 1 for Static IP configuration.

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

Enter selection :
```

6. To configure static IP, press 1. Enter the IP address in CIDR format.

For IPv4, add AA.BB.CC.DD/EE or for IPv6, add AA:BB::CC/EE. You can add only IPv4 or only IPv6 or both the IP addresses. Based on inputs, optional and required tags will change.

```
=====
Extreme Networks, Inc. - Extreme Fabric Automation Interface Configuration Static
=====
Enter the IPv4 address in cidr format(Optional): 10.10.10.1/24
Enter the IPv4 gateway address (Required): 10.10.10.2
Enter the IPv6 address in cidr format (Optional): 2000:1::1/64
Enter the IPv6 gateway address (Required): 2000:1::2
Enter the IPv4 nameserver address (Optional): 8.8.8.8

These are the current network settings that will be used to configure.
=====
Address type: Static
IPv4 Address: 10.10.10.1/24
IPv6 Address: 2000:1::1/64
IPv4 Gateway Address: 10.10.10.2
IPv6 Gateway Address: 2000:1::2
Nameserver Address: 8.8.8.8
=====

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

7. To confirm the settings, press 3.

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

```

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

Enter selection :

```

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]?

```

Backup Database

You can backup an XCO database after the installation.

About This Task

Follow this procedure to backup an XCO database after single-node installation of XCO on Server.

Procedure

1. Connect to the TPVM console or SSH session of the redundant IP as an Extreme user.
2. Log in to EFA using the **efa login** command
3. (Optional) Run the **efa system settings update** command to configure XCO to backup to a remote server

```
text extreme@tpvm:~$ efa login Password: Login successful. --- Time Elapsed:
2.412206025s --- (efa:user)@tpvm:~$ efa system settings update --remote-server-ip
10.25.101.74 --remote-transfer-protocol scp --remote-server-username user --remote-
server-directory /home/user/ Please supply a password for remote server: Setting
Update Successful --- Time Elapsed: 4.723485224s ---
```

4. If you have configured a remote server, run the **efa system backup --remote** command to generate a local backup and copy the resulting tar.gz file to the configured remote device. Otherwise, run the **efa system backup** command.



Note

- Running the **efa system restore** command without any arguments runs in an interactive mode.
- Backups are stored locally in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` location.
- Backup files are date and time stamped in the filename.

For more information on additional backup options, such as backing up managed device configuration, see [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1](#) and [ExtremeCloud Orchestrator Command Reference, 3.2.1](#).

Example

```
extreme@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.412206025s ---
(efa:extreme)extreme@tpvm:~$ efa system backup --remote
Generating backup of EFA...

Backup Location: user@10.25.101.74:/home/user/EFA-3.1.0-2-2022-12-30T01-44-39.193.tar

--- Time Elapsed: 26.700226408s ---
```

Restore Database

You can restore an XCO database from its backup.

About This Task

Follow this procedure to restore an XCO database from its backup after a single-node installation on Server.

Procedure

1. Connect to the SSH session as a root user.
2. Log in to EFA using the **efa login** command

3. (Optional) Transfer the backup from the remote host.
4. Run the **efa system restore** command with the required arguments.



Note

Running the **efa system restore** command without any arguments runs in an interactive mode. It allows you to pick any backup present in `/var/log/efa/backup/` or `/apps/efa_logs/backup/` location. For more information, see [ExtremeCloud Orchestrator CLI Administration Guide, 3.2.1](#) and [ExtremeCloud Orchestrator Command Reference, 3.2.1](#).

Example

```
user@tpvm:~$ efa login
Password:
Login successful.
--- Time Elapsed: 2.019739819s ---
(efa:user)user@tpvm:~$ scp user@remotehost:~/EFA-3.1.0-2-2022-12-30T01-44-39.193.tar ./
user@remotehosts password:
EFA-3.1.0-2-2022-12-30T01-44-39.193.tar          100% 876KB 28.8MB/s  00:00
(efa:user)user@tpvm:~$ efa system restore --backup-tar
EFA-3.1.0-2-2022-12-30T01-44-39.193.tar
Performing EFA restore using EFA-3.1.0-2-2022-12-30T01-44-39.193.tar
Restore operation ID: 4d229051-87db-11ed-8921-5254008962d2
Stopping all EFA services...
All pods are terminated
Restoring databases...
Start services after restore
Restore of encryption keys is completed
Extreme Fabric Automation Stack is now Restored and Ready!
Restore operation is successful
--- Time Elapsed: 6m47.278493209s ---
(efa:user)user@tpvm:~$
```

Upgrade of XCO in Single-Node Deployment

You can upgrade XCO in a single-node deployment.

To upgrade a VM that is provisioned using an OVA image, see [Upgrade Fabric using Interactive Mode](#) on page 108 or [Upgrade Visibility using Interactive Mode](#) on page 114.

Upgrade Fabric using Interactive Mode

You can upgrade a single-node XCO using interactive mode.

About This Task

Follow this procedure to upgrade a single-node XCO on Server.

Procedure

1. Download the XCO tar file and the digests file to the Linux server. Use an administrative account to install XCO on this server.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.

4. Change the current directory to the extracted file location, and run the **source deployment.sh --deploy-suite fabric** command. This will start the XCO installation in a menu-driven interactive mode.
5. When prompted, select **Upgrade, Single-node deployment, or Fabric Automation.**

**Tip**

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

```
admin@server:~/efa$ source deployment.sh -g no
Step 1: Checking for EFA Stack...
Previous Stack found
Please choose:
1 Remove the current EFA Stack
2 Upgrade/Re-deploy
2 Please choose:
1 Single-node deployment 2 Multi-node deployment 1 Single-node Deployment
```

6. When prompted to configure additional management IP networks, take one of the following steps:
 - Select **No** to ignore this optional step or when you have finished entering sub-interface information.
 - Select **Yes** and then provide the requested information
 - 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.

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

7. Verify the installation.
 - From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
 - Run the **efa status** command to ensure that status of all the nodes are up.

```
admin@server$ sudo efactl status
Node: efa
NAME                                READY   STATUS    RESTARTS   AGE   IP
NODE   NOMINATED NODE   READINESS GATES
pod/efa-api-docs-sbg9x                1/1     Running   0           12m   10.42.194.71
efa   <none>          <none>
pod/gosystem-service-m7245            1/1     Running   0           12m   10.42.194.74
efa   <none>          <none>
pod/rabbitmq-cf57z                    1/1     Running   0           13m   10.42.194.68
efa   <none>          <none>
pod/gonotification-service-5nhqb      1/1     Running   0           11m   10.20.229.54
efa   <none>          <none>
pod/goinventory-service-m8fgm         1/1     Running   0           12m   10.42.194.69
efa   <none>          <none>
pod/goauth-service-qr7zs              1/1     Running   0           12m   10.42.194.70
efa   <none>          <none>
pod/gofaultmanager-service-gszrm      1/1     Running   0           12m   10.42.194.72
```

| | | | | | | | |
|------------------------------------|-------------------|----------|-----|--------------|----------------------------|--------|--------------------|
| efa | <none> | <none> | | | | | |
| pod/gofabric-service-wrl6z | | | 1/1 | Running | 0 | 12m | 10.42.194.76 |
| efa | <none> | <none> | | | | | |
| pod/gopolicy-service-4llmn | | | 1/1 | Running | 0 | 12m | 10.42.194.75 |
| efa | <none> | <none> | | | | | |
| pod/goraslog-service-2w8hj | | | 1/1 | Running | 0 | 12m | 10.20.229.54 |
| efa | <none> | <none> | | | | | |
| pod/gotenant-service-zhbp2 | | | 1/1 | Running | 0 | 12m | 10.42.194.77 |
| efa | <none> | <none> | | | | | |
| pod/gosnmp-service-5jngv | | | 1/1 | Running | 0 | 11m | 10.20.229.54 |
| efa | <none> | <none> | | | | | |
| pod/gorbac-service-jnttb | | | 1/1 | Running | 0 | 12m | 10.42.194.73 |
| efa | <none> | <none> | | | | | |
| | | | | | | | |
| NAME | | | | TYPE | CLUSTER-IP | | EXTERNAL-IP |
| PORT(S) | | | | AGE | SELECTOR | | |
| service/rabbitmq | | | | ClusterIP | 10.43.219.62 | <none> | 15672/ |
| TCP,5672/TCP | 13m | | | app=rabbitmq | | | |
| service/rabbitmq-local | | | | NodePort | 10.43.200.66 | <none> | 15672:30673/ |
| TCP,5672:30672/TCP | 13m | | | app=rabbitmq | | | |
| service/goinventory-service | | | | ClusterIP | 10.43.96.100 | <none> | |
| 8082/TCP | | | | 12m | app=goinventory-service | | |
| service/efa-api-docs | | | | ClusterIP | 10.43.127.190 | <none> | |
| 80/TCP | | | | 12m | app=efa-api-docs | | |
| service/goraslog-service | | | | ClusterIP | 10.43.19.118 | <none> | |
| 8091/TCP | | | | 12m | app=goraslog-service | | |
| service/gonotification-service | | | | ClusterIP | 10.43.11.74 | <none> | |
| 8088/TCP | | | | 12m | app=gonotification-service | | |
| service/goauth-service | | | | ClusterIP | 10.43.133.78 | <none> | |
| 8080/TCP | | | | 12m | app=goauth-service | | |
| service/gorbac-service | | | | ClusterIP | 10.43.228.98 | <none> | |
| 8089/TCP | | | | 12m | app=gorbac-service | | |
| service/gofaultmanager-service | | | | ClusterIP | 10.43.58.187 | <none> | |
| 8094/TCP | | | | 12m | app=gofaultmanager-service | | |
| service/gosystem-service | | | | ClusterIP | 10.43.141.200 | <none> | |
| 8090/TCP | | | | 12m | app=gosystem-service | | |
| service/gofabric-service | | | | ClusterIP | 10.43.245.50 | <none> | |
| 8081/TCP | | | | 12m | app=gofabric-service | | |
| service/gopolicy-service | | | | ClusterIP | 10.43.36.27 | <none> | |
| 8093/TCP | | | | 12m | app=gopolicy-service | | |
| service/gotenant-service | | | | ClusterIP | 10.43.52.211 | <none> | |
| 8083/TCP | | | | 12m | app=gotenant-service | | |
| service/goopenstack-service | | | | ClusterIP | 10.43.26.123 | <none> | |
| 8085/TCP | | | | 12m | app=goopenstack-service | | |
| service/govcenter-service | | | | ClusterIP | 10.43.21.5 | <none> | |
| 8086/TCP | | | | 12m | app=govcenter-service | | |
| service/gohyperv-service | | | | ClusterIP | 10.43.43.180 | <none> | |
| 8087/TCP | | | | 12m | app=gohyperv-service | | |
| service/gosnmp-service | | | | ClusterIP | 10.43.78.109 | <none> | |
| 8092/TCP | | | | 12m | app=gosnmp-service | | |
| | | | | | | | |
| NAME | | | | DESIRED | CURRENT | READY | UP-TO-DATE |
| AVAILABLE | NODE | SELECTOR | AGE | CONTAINERS | | | IMAGES |
| SELECTOR | | | | | | | |
| daemonset.apps/goopenstack-service | | | | 0 | 0 | 0 | 0 |
| 0 | non-existing=true | | 12m | openstack | | | goopenstack:3.1.0 |
| app=goopenstack-service | | | | | | | |
| daemonset.apps/efa-api-docs | | | | 1 | 1 | 1 | 1 |
| 1 | <none> | | 12m | efa-api-docs | | | efa-api-docs:3.1.0 |
| app=efa-api-docs | | | | | | | |
| daemonset.apps/govcenter-service | | | | 0 | 0 | 0 | 0 |
| 0 | non-existing=true | | 12m | vcenter | | | govcenter:3.1.0 |
| app=govcenter-service | | | | | | | |
| daemonset.apps/gohyperv-service | | | | 0 | 0 | 0 | 0 |

```

0          non-existing=true  12m  hyperv          gohyperv:3.1.0
app=gohyperv-service
daemonset.apps/gosystem-service      1      1      1      1
1          <none>          12m  gosystem        gosystem:3.1.0
app=gosystem-service
daemonset.apps/rabbitmq              1      1      1      1
1          <none>          13m  rabbitmq-node   rabbitmq:3.1.0
app=rabbitmq
daemonset.apps/gonotification-service 1      1      1      1
1          <none>          12m  gonotification-service gonotification:3.1.0
app=gonotification-service
daemonset.apps/goinventory-service    1      1      1      1
1          <none>          12m  goinventory-service goinventory:3.1.0
app=goinventory-service
daemonset.apps/goauth-service         1      1      1      1
1          <none>          12m  go-auth         goauth:3.1.0
app=goauth-service
daemonset.apps/gofaultmanager-service 1      1      1      1
1          <none>          12m  gofaultmanager-service gofaultmanager:3.1.0
app=gofaultmanager-service
daemonset.apps/gofabric-service       1      1      1      1
1          <none>          12m  gofabric-service gofabric:3.1.0
app=gofabric-service
daemonset.apps/gopolicy-service       1      1      1      1
1          <none>          12m  gopolicy-service gopolicy:3.1.0
app=gopolicy-service
daemonset.apps/goraslog-service       1      1      1      1
1          <none>          12m  goraslog-service goraslog:3.1.0
app=goraslog-service
daemonset.apps/gotenant-service       1      1      1      1
1          <none>          12m  gotenant-service gotenant:3.1.0
app=gotenant-service
daemonset.apps/gosnmp-service         1      1      1      1
1          <none>          12m  gosnmp-service  gosnmp:3.1.0
app=gosnmp-service
daemonset.apps/gorbac-service         1      1      1      1
1          <none>          12m  go-rbac         gorbac:3.1.0
app=gorbac-service
admin@server:~$ efa login
Password:
Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:admin)admin@server:~$ efa status
+-----+-----+-----+-----+
| Node Name | Role | Status | IP |
+-----+-----+-----+-----+
| server    | active | up | 10.20.229.54 |
+-----+-----+-----+-----+
--- Time Elapsed: 4.656427295s ---

```

Upgrade Fabric using Single Command

You can upgrade a single-node XCO using single command.

About This Task

Follow this procedure to upgrade a single-node XCO in a non-interactive mode on Server.

To add or update additional management interfaces, see [Upgrade Fabric using Interactive Mode](#) on page 108.

Procedure

1. Download the XCO tar file and the digests file to the Linux server. Use an administrative account to install XCO on this server.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.
4. Change the current directory to the extracted file location, and run the **source deployment.sh -i no --deploy-suite fabric** command. This will start the XCO upgrade in a non-interactive mode.

```
admin@server:~/efa$ ~/efa$ source deployment.sh -i no --deploy-suite fabric
Checking for EFA Stack...
Deployment mode is upgrade
You have entered:
```

5. Verify the installation.

- From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
- Run the **efa status** command to ensure that status of all the nodes are up..

```
admin@server$ sudo efactl status
Node: efa
NAME                                READY   STATUS    RESTARTS   AGE   IP
NODE  NOMINATED NODE  READINESS GATES
pod/efa-api-docs-sbg9x              1/1     Running   0           12m   10.42.194.71
efa  <none>          <none>
pod/gosystem-service-m7245         1/1     Running   0           12m   10.42.194.74
efa  <none>          <none>
pod/rabbitmq-cf57z                  1/1     Running   0           13m   10.42.194.68
efa  <none>          <none>
pod/gonotification-service-5nhqb   1/1     Running   0           11m   10.20.229.54
efa  <none>          <none>
pod/goinventory-service-m8fgm      1/1     Running   0           12m   10.42.194.69
efa  <none>          <none>
pod/goauth-service-qr7zs           1/1     Running   0           12m   10.42.194.70
efa  <none>          <none>
pod/gofaultmanager-service-gszrm   1/1     Running   0           12m   10.42.194.72
efa  <none>          <none>
pod/gofabric-service-wrl6z         1/1     Running   0           12m   10.42.194.76
efa  <none>          <none>
pod/gopolicy-service-4llmn         1/1     Running   0           12m   10.42.194.75
efa  <none>          <none>
pod/goraslog-service-2w8hj         1/1     Running   0           12m   10.20.229.54
efa  <none>          <none>
pod/gotentant-service-zhbp2        1/1     Running   0           12m   10.42.194.77
efa  <none>          <none>
pod/gosnmp-service-5jnqv           1/1     Running   0           11m   10.20.229.54
efa  <none>          <none>
pod/gorbac-service-jnttb           1/1     Running   0           12m   10.42.194.73
efa  <none>          <none>

NAME                                TYPE           CLUSTER-IP      EXTERNAL-IP
PORT(S)                              AGE  SELECTOR
service/rabbitmq                      ClusterIP      10.43.219.62    <none>          15672/
TCP,5672/TCP                          13m  app=rabbitmq
service/rabbitmq-local                 NodePort       10.43.200.66    <none>          15672:30673/
TCP,5672:30672/TCP 13m  app=rabbitmq
```


| | | | |
|---------------------------------------|---------------|----------------------------|---|
| service/goinventory-service | ClusterIP | 10.43.96.100 | <none> |
| 8082/TCP | 12m | app=goinventory-service | |
| service/efa-api-docs | ClusterIP | 10.43.127.190 | <none> |
| 80/TCP | 12m | app=efa-api-docs | |
| service/goraslog-service | ClusterIP | 10.43.19.118 | <none> |
| 8091/TCP | 12m | app=goraslog-service | |
| service/gonotification-service | ClusterIP | 10.43.11.74 | <none> |
| 8088/TCP | 12m | app=gonotification-service | |
| service/goauth-service | ClusterIP | 10.43.133.78 | <none> |
| 8080/TCP | 12m | app=goauth-service | |
| service/gorbac-service | ClusterIP | 10.43.228.98 | <none> |
| 8089/TCP | 12m | app=gorbac-service | |
| service/gofaultmanager-service | ClusterIP | 10.43.58.187 | <none> |
| 8094/TCP | 12m | app=gofaultmanager-service | |
| service/gosystem-service | ClusterIP | 10.43.141.200 | <none> |
| 8090/TCP | 12m | app=gosystem-service | |
| service/gofabric-service | ClusterIP | 10.43.245.50 | <none> |
| 8081/TCP | 12m | app=gofabric-service | |
| service/gopolicy-service | ClusterIP | 10.43.36.27 | <none> |
| 8093/TCP | 12m | app=gopolicy-service | |
| service/gotenant-service | ClusterIP | 10.43.52.211 | <none> |
| 8083/TCP | 12m | app=gotenant-service | |
| service/goopenstack-service | ClusterIP | 10.43.26.123 | <none> |
| 8085/TCP | 12m | app=goopenstack-service | |
| service/govcenter-service | ClusterIP | 10.43.21.5 | <none> |
| 8086/TCP | 12m | app=govcenter-service | |
| service/gohyperv-service | ClusterIP | 10.43.43.180 | <none> |
| 8087/TCP | 12m | app=gohyperv-service | |
| service/gosnmp-service | ClusterIP | 10.43.78.109 | <none> |
| 8092/TCP | 12m | app=gosnmp-service | |
| NAME | | | |
| AVAILABLE | NODE SELECTOR | AGE | DESIRED CURRENT READY UP-TO-DATE |
| SELECTOR | | | CONTAINERS IMAGES |
| daemonset.apps/goopenstack-service | | 0 | 0 0 0 |
| 0 non-existing=true | | 12m | openstack goopenstack:3.1.0 |
| app=goopenstack-service | | | |
| daemonset.apps/efa-api-docs | | 1 | 1 1 1 |
| 1 <none> | | 12m | efa-api-docs efa-api-docs:3.1.0 |
| app=efa-api-docs | | | |
| daemonset.apps/govcenter-service | | 0 | 0 0 0 |
| 0 non-existing=true | | 12m | vcenter govcenter:3.1.0 |
| app=govcenter-service | | | |
| daemonset.apps/gohyperv-service | | 0 | 0 0 0 |
| 0 non-existing=true | | 12m | hyperv gohyperv:3.1.0 |
| app=gohyperv-service | | | |
| daemonset.apps/gosystem-service | | 1 | 1 1 1 |
| 1 <none> | | 12m | gosystem gosystem:3.1.0 |
| app=gosystem-service | | | |
| daemonset.apps/rabbitmq | | 1 | 1 1 1 |
| 1 <none> | | 13m | rabbitmq-node rabbitmq:3.1.0 |
| app=rabbitmq | | | |
| daemonset.apps/gonotification-service | | 1 | 1 1 1 |
| 1 <none> | | 12m | gonotification-service gonotification:3.1.0 |
| app=gonotification-service | | | |
| daemonset.apps/goinventory-service | | 1 | 1 1 1 |
| 1 <none> | | 12m | goinventory-service goinventory:3.1.0 |
| app=goinventory-service | | | |
| daemonset.apps/goauth-service | | 1 | 1 1 1 |
| 1 <none> | | 12m | go-auth goauth:3.1.0 |
| app=goauth-service | | | |
| daemonset.apps/gofaultmanager-service | | 1 | 1 1 1 |
| 1 <none> | | 12m | gofaultmanager-service gofaultmanager:3.1.0 |
| app=gofaultmanager-service | | | |

```

daemonset.apps/gofabric-service      1      1      1      1
1      <none>      12m      gofabric-service      gofabric:3.1.0
app=gofabric-service
daemonset.apps/gopolicy-service      1      1      1      1
1      <none>      12m      gopolicy-service      gopolicy:3.1.0
app=gopolicy-service
daemonset.apps/goraslog-service      1      1      1      1
1      <none>      12m      goraslog-service      goraslog:3.1.0
app=goraslog-service
daemonset.apps/gotenant-service      1      1      1      1
1      <none>      12m      gotenant-service      gotenant:3.1.0
app=gotenant-service
daemonset.apps/gosnmp-service      1      1      1      1
1      <none>      12m      gosnmp-service      gosnmp:3.1.0
app=gosnmp-service
daemonset.apps/gorbac-service      1      1      1      1
1      <none>      12m      go-rbac      gorbac:3.1.0
app=gorbac-service
admin@server:~$ efa login
Password:
Login successful.
--- Time Elapsed: 4.121103171s ---
(efa:admin)admin@server:~$ efa status
+-----+-----+-----+-----+
| Node Name      | Role   | Status | IP           |
+-----+-----+-----+-----+
| server         | active | up     | 10.20.229.54 |
+-----+-----+-----+-----+
--- Time Elapsed: 4.656427295s ---

```

Upgrade Visibility using Interactive Mode

You can upgrade visibility manager using interactive mode.

About This Task

Follow this procedure to upgrade visibility manager in a single-node XCO using a interactive mode on Server.

Procedure

1. Download the XCO tar file and the digests file to the Linux server. Use an administrative account to install XCO on this server.

```

admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz

```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.
4. Change the current directory to the extracted file location. Run the **source deployment.sh --deploy-suite packet** command. This will start the XCO installation in a menu-driven interactive mode.

For command line based interactive mode, run the **source deployment.sh --deploy-suite packet -g no** command.

5. When prompted, select **Upgrade, Single-node deployment**, or **Packet**.



Tip

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

```
text admin@server:~/efa$ source deployment.sh -g no Step 1:
Checking for EFA Stack...
Previous Stack found
Please choose: 1 Remove the current EFA Stack 2 Upgrade/Re-deploy 2 Please choose: 1
Single-node deployment 2 Multi-node deployment 1 Single-node Deployment`
```

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

- Select **No** to ignore this optional step or when you have finished entering sub-interface information.
- Select **Yes** and then provide the requested information
 - 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.

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

7. Verify the installation.

- From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
- Run the **efa status** command to ensure that status of all the nodes are up.

```
user@dev-server:/opt/efa/efa$ sudo efactl status
Node: efa
NAME                                READY  STATUS   RESTARTS   AGE
IP          NODE    NOMINATED NODE  READINESS GATES
pod/efa-api-docs-jpppb               1/1    Running   0           28m
10.42.194.72 efa    <none>          <none>
pod/ui-service-4dggw                 1/1    Running   0           28m
10.42.194.73 efa    <none>          <none>
pod/rabbitmq-tdq8q                   1/1    Running   1 (27m ago) 28m
10.42.194.77 efa    <none>          <none>
pod/gosystem-service-wkbn9           1/1    Running   0           28m
10.42.194.75 efa    <none>          <none>
pod/gofaultmanager-service-n9c5s     1/1    Running   0           28m
10.42.194.74 efa    <none>          <none>
pod/gonotification-service-ljnb4     1/1    Running   0           28m
10.37.34.129 efa    <none>          <none>
pod/goinventory-service-2hk7t        1/1    Running   0           28m
10.42.194.69 efa    <none>          <none>
pod/gorbac-service-d6bxx             1/1    Running   0           28m
10.42.194.70 efa    <none>          <none>
pod/goauth-service-9n544             1/1    Running   0           28m
10.42.194.76 efa    <none>          <none>
pod/goevm-service-d45jl              1/1    Running   0           28m
10.42.194.71 efa    <none>          <none>
pod/goraslog-service-t6n6n          1/1    Running   0           28m
```

| NAME | TYPE | CLUSTER-IP | EXTERNAL-IP |
|--------------------------------|--------------------------------|---------------|---------------------|
| 10.37.34.129 efa <none> | <none> | | |
| service/rabbitmq | ClusterIP | 10.43.136.36 | <none> 15672/ |
| TCP, 5672/TCP | 28m app=rabbitmq | | |
| service/rabbitmq-local | NodePort | 10.43.115.197 | <none> 15672:30673/ |
| TCP, 5672:30672/TCP | 28m app=rabbitmq | | |
| service/goinventory-service | ClusterIP | 10.43.194.233 | <none> |
| 8082/TCP | 28m app=goinventory-service | | |
| service/efa-api-docs | ClusterIP | 10.43.160.75 | <none> |
| 80/TCP | 28m app=efa-api-docs | | |
| service/goraslog-service | ClusterIP | 10.43.28.179 | <none> 8091/ |
| TCP, 54322/TCP | 28m app=goraslog-service | | |
| service/gonotification-service | ClusterIP | 10.43.10.148 | <none> |
| 8088/TCP | 28m app=gonotification-service | | |
| service/goauth-service | ClusterIP | 10.43.186.66 | <none> |
| 8080/TCP | 28m app=goauth-service | | |
| service/gorbac-service | ClusterIP | 10.43.116.36 | <none> |
| 8089/TCP | 28m app=gorbac-service | | |
| service/gofaultmanager-service | ClusterIP | 10.43.101.181 | <none> |
| 8094/TCP | 28m app=gofaultmanager-service | | |
| service/gosystem-service | ClusterIP | 10.43.51.133 | <none> |
| 8090/TCP | 28m app=gosystem-service | | |
| service/goevm-service | ClusterIP | 10.43.243.102 | <none> 8101/ |
| TCP, 8106/TCP | 28m app=goevm-service | | |
| service/ui-service | ClusterIP | 10.43.158.199 | <none> |
| 3000/TCP | 28m app=ui-service | | |

| NAME | AVAILABLE | NODE SELECTOR | AGE | DESIRED CONTAINERS | CURRENT | READY IMAGES | UP-TO-DATE |
|---------------------------------------|-----------|---------------|-----|--------------------------|---------|------------------------|------------|
| daemonset.apps/efa-api-docs | 1 | <none> | 28m | 1 efa-api-docs | 1 | 1 efa-api-docs:3.2.0 | 1 |
| app=efa-api-docs | | | | | | | |
| daemonset.apps/ui-service | 1 | <none> | 28m | 1 ui | 1 | 1 ui:3.2.0 | 1 |
| app=ui-service | | | | | | | |
| daemonset.apps/rabbitmq | 1 | <none> | 28m | 1 rabbitmq-node | 1 | 1 rabbitmq:3.2.0 | 1 |
| app=rabbitmq | | | | | | | |
| daemonset.apps/gosystem-service | 1 | <none> | 28m | 1 gosystem | 1 | 1 gosystem:3.2.0 | 1 |
| app=gosystem-service | | | | | | | |
| daemonset.apps/gofaultmanager-service | 1 | <none> | 28m | 1 gofaultmanager-service | 1 | 1 gofaultmanager:3.2.0 | 1 |
| app=gofaultmanager-service | | | | | | | |
| daemonset.apps/gonotification-service | 1 | <none> | 28m | 1 gonotification-service | 1 | 1 gonotification:3.2.0 | 1 |
| app=gonotification-service | | | | | | | |
| daemonset.apps/goinventory-service | 1 | <none> | 28m | 1 goinventory-service | 1 | 1 goinventory:3.2.0 | 1 |
| app=goinventory-service | | | | | | | |
| daemonset.apps/gorbac-service | 1 | <none> | 28m | 1 go-rbac | 1 | 1 gorbac:3.2.0 | 1 |
| app=gorbac-service | | | | | | | |
| daemonset.apps/goauth-service | 1 | <none> | 28m | 1 go-auth | 1 | 1 goauth:3.2.0 | 1 |
| app=goauth-service | | | | | | | |
| daemonset.apps/goevm-service | 1 | <none> | 28m | 1 evm | 1 | 1 goevm:3.2.0 | 1 |
| app=goevm-service | | | | | | | |
| daemonset.apps/goraslog-service | 1 | | | 1 | 1 | 1 | 1 |

```
1          <none>          28m  goraslog-service          goraslog:3.2.0
  app=goraslog-service
```

Upgrade Visibility using Single Command

You can upgrade a visibility manager using single command.

Before You Begin

To add or update additional management interfaces, see [Upgrade Visibility using Interactive Mode](#) on page 114.

About This Task

Follow this procedure to upgrade a visibility manager in a non-interactive mode on Server.

Procedure

1. Download the XCO tar file and the digests file to the Linux server. Use an administrative account to install XCO on this server.

```
admin@server$ ls
efa-3.1.0-digests.tar.gz  efa-3.1.0.tar.gz
```

2. Verify the PGP signature as described in article 48172 on the [Extreme Portal](#).
3. Extract the XCO tarball to a desired location. Ensure that this location has the permissions to allow your administrative user to run commands.
4. Change the current directory to the extracted file location, and run the **source deployment.sh -i no --deploy-suite packet** command. This will start the XCO upgrade in a non-interactive mode.

```
text admin@server:~/efa$ ~/efa$ source deployment.sh -i no --deploy-suite packet
Checking for EFA Stack...
Deployment mode is upgrade
You have entered:
```

5. Verify the installation.
 - From the EFA command line on the console or SSH session, run the **sudo efactl status** command to see the status of nodes, pods, and services.
 - Run the **efa status** command to ensure that status of all the nodes are up.

High-Availability Multi-Node Deployment

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

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 3: 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, 3.2.1](#).

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, litemstream. 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 [High-Availability Requirements](#) on page 89.
- 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.

- 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.

- 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.
- 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.

- Verify the installation.
 - From the XCO command line, run the **sudo efact1 status** command to see the status of nodes, pods, and services.
 - 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.

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

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 [TPVM Deployment Requirements](#) on page 29 and [TPVM High-Availability Requirements](#) on page 34 for a list of prerequisites.

To upgrade a VM that is provisioned using an OVA image, follow the procedure explained either [Upgrade Fabric using Interactive Mode](#) on page 108 in or in [Upgrade Visibility using Interactive Mode](#) on page 114.

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.

About This Task

Follow this procedure to recover from an upgrade failure.

Procedure

1. To rerun the upgrade, complete the following steps for the type of upgrade you were attempting:
 - a. [Upgrade XCO from a Single-Node to a Multi-Node Deployment](#) on page 127
 - b. [Upgrade XCO on TPVM in a Single-Node Deployment](#) on page 44
 - c. [Upgrade Fabric using Interactive Mode](#) on page 108 and [Upgrade Fabric using Single Command](#) on page 111
 - d. [Upgrade XCO in a Multi-Node Deployment](#) on page 123
 - e. [Upgrade XCO on TPVM in a Multi-Node Deployment](#) on page 56
2. To perform a fresh installation and restore the system backup, complete the following steps:
 - a. Uninstall XCO to remove any components that might have been installed before the upgrade failed.
 - i. [Uninstall XCO on TPVM in a Single-Node and Multi-Node Deployment](#) on page 77
 - ii. [Uninstall XCO in a Single-Node or Multi-Node Deployment](#) on page 130
 - b. Follow the steps for the type of installation you need.
 - i. [Install XCO on TPVM in a Multi-Node Deployment](#) on page 52
 - ii. [Install XCO in a Multi-Node Deployment](#) on page 121
 - iii. [Install XCO Single-Node Fabric Management on TPVM](#) on page 37
 - iv. [Install Fabric using Interactive Mode](#) on page 92 and [Install Fabric using Single Command](#) on page 96
 - 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, 3.2.1](#).

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 16. 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 58, [Upgrade TPVM OS for XCO Single-Node Deployment](#) on page 49, [TPVM Complete Package Upgrade](#) on page 63, and [TPVM Incremental Upgrade](#) on page 69.
- 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 [Backup Database](#) on page 106.

```
$ 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 130.

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 [Restore Database](#) on page 107

```
# 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 efactl status** to verify that all pods on the active node are in Running state.
- Run **sudo efactl db-status** to verify that the MariaDB is active (running)

Upgrade 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 [Upgrade Fabric using Interactive Mode](#) on page 108 and [Upgrade Fabric using Single Command](#) on page 111.
- Ensure that you have completed the high-availability prerequisites in [High-Availability Requirements](#) on page 89.

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.

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 [High-Availability Requirements](#) on page 89.
- 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.

6. 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, 3.2.1](#).

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.

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>
```

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.