



# Extreme Fabric Automation Release Notes

Version 2.4.0

9036930-00 Rev AA  
February 2021



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

## Legal Notice

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

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

## Trademarks

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

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

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

## Open Source Declarations

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



# Table of Contents

---

<b>Preface.....</b>	<b>4</b>
Text Conventions.....	4
Documentation and Training.....	5
Getting Help.....	6
Subscribe to Service Notifications.....	6
Providing Feedback.....	6
<b>Release Notes.....</b>	<b>8</b>
New Features.....	8
CLI Commands.....	10
New commands.....	10
Modified commands.....	11
Supported Platforms and Deployment Models.....	12
Prerequisites for SLX Devices.....	13
Prerequisites for Integration Applications.....	14
EFA Requirements.....	15
General requirements.....	15
High-availability requirements.....	16
EFA Port Requirements.....	17
EFA Installation Modes.....	18
EFA Installation and Upgrade.....	18
Defects Closed with Code Changes .....	19
Defects Closed without Code Changes.....	28
Open Defects.....	31



# Preface

---

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






## Text Conventions

---

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

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

**Table 1: Notes and warnings**

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

**Table 2: Text**

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

**Table 3: Command syntax**

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

## Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

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

[Extreme Optics Compatibility](#)

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

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

---

## Getting Help

---

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

### Extreme Portal

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

### The Hub

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

### Call GTAC

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

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

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

## Subscribe to Service Notifications

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

1. Go to [www.extremenetworks.com/support/service-notification-form](http://www.extremenetworks.com/support/service-notification-form).
2. Complete the form (all fields are required).
3. Select the products for which you would like to receive notifications.



### Note

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

4. Select **Submit**.

---

## Providing Feedback

---

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help

you work better, so we want to hear from you. We welcome all feedback, but we especially want to know about:

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

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

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

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



# Release Notes

---

- [New Features on page 8](#)
- [CLI Commands on page 10](#)
- [Supported Platforms and Deployment Models on page 12](#)
- [Prerequisites for SLX Devices on page 13](#)
- [Prerequisites for Integration Applications on page 14](#)
- [EFA Requirements on page 15](#)
- [EFA Port Requirements on page 17](#)
- [EFA Installation Modes on page 18](#)
- [EFA Installation and Upgrade on page 18](#)
- [Defects Closed with Code Changes on page 19](#)
- [Defects Closed without Code Changes on page 28](#)
- [Open Defects on page 31](#)

## New Features

---

Extreme Fabric Automation 2.4.0 provides several new features and improvements.

**Table 4: Features and improvements**

Feature	Description
EFA as an SNMP trap receiver	A new service, gosnmp-service, enables EFA to persist trap subscribers, receive SNMP traps, and forward the traps to subscribers. For more information, see the <a href="#">Extreme Fabric Automation Administration Guide, 2.4.0</a> .
OpenStack integration enhancements	This release offers numerous improvements, including the following: <ul style="list-style-type: none"><li>• Support for the Victoria release of OpenStack</li><li>• Layer 3 service plugin enhancements</li><li>• Synchronize mechanism in the OSS service</li><li>• Journaling support for the Layer 2 and Layer 3 plugins</li></ul> For more information, see the <a href="#">Extreme Fabric Automation OpenStack Integration Guide, 2.4.0</a> .



**Table 4: Features and improvements (continued)**

Feature	Description
Centralized routing	With centralized routing, routing configuration is performed only on border leaf pairs. For more information, see the <a href="#">Extreme Fabric Automation Administration Guide, 2.4.0</a> .
Multiple management IP networks	You can create multiple IP subnets for support of EFA on all of the subnets. For more information, see the <a href="#">Extreme Fabric Automation Administration Guide, 2.4.0</a> and the <a href="#">Extreme Fabric Automation Deployment Guide, 2.4.0</a> .
Certificate expiration notifications	If a device certificate is within 30 days of expiration, an event is generated. For more information, see the <a href="#">Extreme Fabric Automation Administration Guide, 2.4.0</a> .
New IPv6 Neighbor Discovery options for the <b>efa tenant epg create</b> and <b>efa tenant epg update</b> commands	<p>New options for configuring MTU, router advertisement, prefix lifetimes, and configuration type:</p> <ul style="list-style-type: none"> <li>• --ipv6-nd-mtu</li> <li>• --ipv6-nd-managed-config</li> <li>• --ipv6-nd-other-config</li> <li>• --ipv6-nd-prefix</li> <li>• --ipv6-nd-prefix-valid-lifetime</li> <li>• --ipv6-nd-prefix-preferred-lifetime</li> <li>• --ipv6-prefix-no-advertise</li> <li>• --ipv6-nd-prefix-config-type</li> </ul> <p>For more information, see the <a href="#">Extreme Fabric Automation Command Reference, 2.4.0</a>.</p>
New ctag option for the <b>efa tenant epg create</b> and <b>efa tenant epg update</b> commands	Use the <b>--ctag-description</b> option to assign a unique description for a ctag. For more information, see the <a href="#">Extreme Fabric Automation Command Reference, 2.4.0</a> .
New minimum links option for the <b>efa tenant po create</b> and <b>efa tenant po update</b> commands	Use the <b>--min-link-count</b> option to specify the least number of links that need to be operationally up to declare the port channel as operationally up. For more information, see the <a href="#">Extreme Fabric Automation Command Reference, 2.4.0</a> .
New anycast operations for the <b>efa tenant epg create</b> and <b>efa tenant epg update</b> commands	<p>With two new operations (<b>anycast-ip-add</b> and <b>anycast-ip-delete</b>) you can add or delete anycast addresses using the following options:</p> <ul style="list-style-type: none"> <li>• --anycast-ip</li> <li>• --anycast-ipv6</li> </ul> <p>For more information, see the <a href="#">Extreme Fabric Automation Command Reference, 2.4.0</a>.</p>
New graceful restart toggle for the <b>efa tenant vrf create</b> and <b>efa tenant vrf update</b> commands	You can set the <b>--graceful-restart</b> option to true or false to enable or deactivate the functionality. For more information, see the <a href="#">Extreme Fabric Automation Command Reference, 2.4.0</a> .

**Table 4: Features and improvements (continued)**

Feature	Description
Workarounds for event synchronization and event handling issues	For more information about these workarounds, see the <a href="#">Extreme Fabric Automation Administration Guide, 2.4.0</a> .
Support for upgrade from EFA 2.3.x to any later version of EFA	For more information, see <a href="#">EFA Installation and Upgrade</a> on page 18.
RBAC service in active-standby mode	In previous releases, the RBAC service was in active-active mode. For more information, see the <a href="#">Extreme Fabric Automation Deployment Guide, 2.4.0</a> .
REST API changes for backup, restore, and supportsave	<ul style="list-style-type: none"> <li>Backup, restore, and supportsave APIs are moved from GoSystem to the Monitoring service. You can access these APIs on port 8078.</li> <li>APIs for backup (<code>/v1/system/backup</code>) and restore (<code>/v1/system/restore</code>) are changed from GET to POST.</li> <li>The restore API is now asynchronous and returns an ID when it is triggered. You can use this ID in the status API (<code>/v1/system/restore/&lt;ID&gt;</code>) to return the progress of restore operations.</li> </ul> <p>For more information, see the set of API guides on the Extreme Fabric Automation <a href="#">documentation</a> web page.</p>
The output of the <b>efa system restore</b> command now shows you the backup tar files you can choose from	For more information, see the "Backup and Restore the EFA Database" topic in the <a href="#">Extreme Fabric Automation Administration Guide, 2.4.0</a> .
New REST API for deleting individual backup and supportsave files	<p>Usage: <b>efa system cleanup &lt;flags&gt;</b></p> <p>Flags:</p> <ul style="list-style-type: none"> <li><b>--archive-type:</b> The type of archive to delete. Supported entries are backup or supportsave. The default is backup.</li> <li><b>--name:</b> The name of the archive to delete.</li> </ul> <p>For more information, see the set of API guides on the Extreme Fabric Automation <a href="#">documentation</a> web page.</p>

## CLI Commands

There are new and modified commands in this release.

For more information about each command, see the [Extreme Fabric Automation Command Reference, 2.4.0](#).

## New commands

- `efa inventory device firmware-download commit`
- `efa inventory device firmware-download restore`
- `efa inventory device interface list`
- `efa mgmt subinterface create`
- `efa mgmt subinterface delete`
- `efa mgmt subinterface show`
- `efa snmp subscriber`
- `efa status`

## Modified commands

- `efa fabric device add-bulk`
- `efa inventory device firmware-download execute`
- `efa system settings show`
- `efa system settings update`
- `efa tenant epg create`
- `efa tenant epg show`
- `efa tenant epg update`
- `efa tenant po create`
- `efa tenant po show`
- `efa tenant po update`
- `efa tenant service bgp peer show`
- `efa tenant service bgp peer-group show`
- `efa tenant show`
- `efa tenant vrf create`
- `efa tenant vrf show`
- `efa tenant vrf update`
- `efa-health`
- `efa-journal`
- `efa-sync`
- `openstack network efa-topology-link-map create`
- `openstack network efa-topology-link-map delete`
- `openstack network efa-topology-link-map list`

## Supported Platforms and Deployment Models

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

**Table 5: Bare Metal Deployment Models**

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
EFA 2.2.0, 2.3.0, and 2.4.0	External server (bare metal)	More than 24	Yes	16.04 and 18.04	<ul style="list-style-type: none"> <li>CPU: 4 cores</li> <li>Storage: 50 GB</li> <li>RAM: 8 GB</li> </ul>

**Table 6: OVA Deployment Models**

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
EFA 2.2.0, 2.3.0, and 2.4.0 (Secure mode)	External server (OVA)	More than 24	Yes	18.04	<ul style="list-style-type: none"> <li>CPU: 4 cores</li> <li>Storage: 50 GB</li> <li>RAM: 8 GB</li> </ul>

**Table 7: TPVM Deployment Models**

EFA Version	TPVM Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS Version
2.2.0, 2.3.x	<ul style="list-style-type: none"> <li>SLX 9150</li> <li>SLX 9250</li> </ul>	Up to 24	Yes	18.04	20.1.2x
2.3.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	20.2.2a
2.4.0	<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	20.2.2b

**Table 8: TPVM Software Support**

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

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

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

**Table 9: IP Fabric Topology Matrix**

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

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

EFA Version	Neutron Version	SLX-OS Version
2.3.0	2.3.0_19	20.1.2d
2.3.1	2.3.1_02	20.1.2e
2.3.2	2.3.2_02	20.1.2e, 20.2.2a
2.4.0	3.0.0-23	20.2.2b, 20.2.3, 20.2.3a

## Prerequisites for SLX Devices

Review this list of prerequisites for EFA support of SLX devices.

**Table 11: Prerequisites**

Prerequisite	Description
Supported devices	See <a href="#">Supported Platforms and Deployment Models</a> on page 12.
Operating system	See <a href="#">Supported Platforms and Deployment Models</a> on page 12.

**Table 11: Prerequisites (continued)**

Prerequisite	Description
IP addresses	Management IP addresses configured on all devices.
Security	Proper certificate installation and audit log requires NTP configuration on the SLX device. For example: <b>efa inventory device execute-cli --ip &lt;switch IPs&gt; --command "ntp server &lt;ntp server ip&gt; - config"</b> For more information about NTP requirements, see <a href="#">EFA Requirements</a> on page 15.

## Prerequisites for Integration Applications

**Table 12: vCenter**

Feature	Minimum Requirement
vCenter controller	Windows 2016 with vSphere 6.5.x or 6.7.x
Compute nodes	ESXi 6.x or later

**Table 13: SCVMM (Hyper-V)**

Feature	Minimum Requirement
SCVMM controller	Windows SCVMM 2016
Compute nodes	Windows Hyper-V 2016

**Table 14: OpenStack**

Feature	Minimum Supported Version
Operating system	OpenSuse 15.2 or SL 15 SP2
OpenStack release	Victoria
SLX platforms	For a list of supported SLX-OS versions, see <a href="#">Supported Platforms and Deployment Models</a> on page 12.

## EFA Requirements

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

### General requirements

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

You can edit `/etc/systemd/timesyncd.conf` to select NTP servers in the `[Time]` section of the configuration file. The `NTP=` option takes a space-separated list of host names or IP addresses. NTP suggests selecting as many servers as is feasible, but at least 3. Select from the pool of publicly available servers or your company's internal NTP servers. For example:

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

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

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

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

## High-availability requirements

- **OS:** All nodes in the high-availability cluster must have the same version of the operating system. For more information about supported operating systems, see [Supported Platforms and Deployment Models](#) on page 12.
- **Host names:** High-availability host names must be unique.
- **IP addresses:**
  - High-availability deployments require an extra IP address: virtual IP, cluster IP, or host IP. Ensure that this extra address is an unallocated IP address in the same subnet as the nodes that will form the cluster.
  - All nodes in the cluster must have an IP address in the same subnet as the virtual IP address.
- **SSH:** (For SLX-OS 20.2.3 and later with TPVM version 4.2.2 and later) Before installing EFA, configure SSH passwordless access between TPVM users. You can use the SLX command line and the following commands.
  - To configure a trusted peer: **device# tpvm config trusted-peer add <peer-tpvm-ip> sudo-user <tpvm-sudo-user> password <sudo-userpassword>.**
  - To display trusted peer information: **device# show tpvm config trusted-peer.**
  - To remove a trusted peer: **device# tpvm config trusted-peer remove <peer-tpvm-ip> sudo-user <tpvm-sudo-user> password <sudo-userpassword>.**



### Note

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

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

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



### Note

Modify this script to suit your requirements.

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

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

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

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

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

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

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

```



```

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

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

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

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

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

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

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

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

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

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

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

echo "Completed passwordless ssh login between TPVMs."

```

## EFA Port Requirements

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

**Table 15: General port requirements**

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

**Table 15: General port requirements (continued)**

Port	Service
30672	Rabbitmq
31672	Rabbitmq management

**Table 16: Port requirements for high availability**

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

## EFA Installation Modes

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

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

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

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



### Note

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

## EFA Installation and Upgrade

You can install and upgrade EFA in single-node environments, in multi-node high-availability environments, and on a TPVM. You can also upgrade from a single-node to a multi-node deployment.

For complete information on the following topics, see the [Extreme Fabric Automation Deployment Guide, 2.4.0](#).

- Installing or upgrading EFA on a single node, including TPVM deployments
- Installing EFA on two nodes for high availability, including TPVM deployments
- Upgrading EFA from a single-node to a multi-node deployment
- Replacing a node in a multi-node deployment
- Deploying or upgrading the OVA

Upgrade support varies by software version. Upgrading from version 2.3.x to any later version is supported. Support for upgrades from versions earlier than version 2.3.x is described in the "Supported Upgrade Paths" topic in the *Extreme Fabric Automation Deployment Guide*.

**Note**

Although the upgrade process allows you do to so, do not upgrade from secure mode to standard mode.

## Defects Closed with Code Changes

The following defects, which were previously disclosed as open, are resolved in EFA 2.4.0.

<b>Parent Defect ID:</b>	EFA-3512	<b>Issue ID:</b>	EFA-3512
<b>Severity:</b>	S4 - Low		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	EPG create might fail with VNI resource not being available in the fabric		
<b>Condition:</b>	Execute EPG create, delete and re-create CLI in quick succession as below: 1. Create EPG/Networks with user-provided VNI parameter. 2. Delete EPG. 3. Create EPG again with the same parameters as in step-1.		

<b>Parent Defect ID:</b>	EFA-3694	<b>Issue ID:</b>	EFA-3694
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Request processing takes longer time and CLI remains hung till the request is processed. EPG update may fail if the wait time is more than 16 minutes.		
<b>Condition:</b>	1. Create EPG with 100 networks. 2. Update EPG with 20 ports/port-channels from one switch using "port-group-add" operation.		

<b>Parent Defect ID:</b>	EFA-3717	<b>Issue ID:</b>	EFA-3717
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.1.0
<b>Symptom:</b>	BD based L3 EPG networks will go into cfg-refresh-err state if out-of-band configured BD-LIFs (without associated VLAN) are present on the switch.		
<b>Condition:</b>	1. Create BD-LIF (without any associated VLAN) configuration manually on the switch followed by EFA inventory update. 2. Create EPG which results in the same BD-LIF on the switch.		

<b>Parent Defect ID:</b>	EFA-4056	<b>Issue ID:</b>	EFA-4056
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0

<b>Symptom:</b>	The second EPG will get created successfully but the anycast-ipv6 provided in the second EPG will not get configured on the switch. The second EPG create command should have failed with an error mentioning a unique anycast-ipv6 needs to be provided for a given BD name.
<b>Condition:</b>	1. Create BD based L3 EPG with some BD label and an anycast-ipv6. 2. Create another BD based L3 EPG with the same BD label (as used in 1) but with a different anycast-ipv6.

<b>Parent Defect ID:</b>	EFA-5064	<b>Issue ID:</b>	EFA-5064
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	EFA configures both "suppress-arp" (needed for ipv4) and "suppress-nd" (needed for ipv6) for the particular network (VLAN/BD). EFA shouldn't have configured "suppress-nd" on the network.		
<b>Condition:</b>	Configure L3 EPG with only anycast-ip and no anycast-ipv6.		

<b>Parent Defect ID:</b>	EFA-5257	<b>Issue ID:</b>	EFA-5257
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	When VRF is added and deleted to/from and Endpoint Group, in quick succession, multiple times, events received from inventory service can get interleaved with the commands. This can cause EFA command execution path to find database entries that are yet to be deleted due to previous command run.		
<b>Condition:</b>	Issue can be observed when vrf-add and vrf-delete operation is executed on Endpoint Group in quick succession multiple times.		

<b>Parent Defect ID:</b>	EFA-5623	<b>Issue ID:</b>	EFA-5623
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Drift and reconcile logs are not notified to the Notification engine.		

<b>Parent Defect ID:</b>	EFA-5675	<b>Issue ID:</b>	EFA-5675
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0

<b>Symptom:</b>	EPG update "vrf-add" operation can fail with the reason as vrf to be added has conflicting vrf on the switch.
<b>Condition:</b>	Execute EPG update "vrf-add", "vrf-delete" and "vrf-add" operation CLI in quick succession as below: 1. Update EPG for operation vrf-add. 2. Update EPG for operation vrf-delete. 3. Update the same EPG again with operation vrf-add for the same vrf which was deleted in step 2.

<b>Parent Defect ID:</b>	EFA-5708	<b>Issue ID:</b>	EFA-5708
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Device status remains in cfg-in-sync when devices have mismatch in selective configs w.r.t to intended configs		
<b>Condition:</b>	Periodic discovery identifies the drift and raise appropriate events; most of the events that are handled do not change the device state from cfg-in-sync to cfg-refreshed		

<b>Parent Defect ID:</b>	EFA-5834	<b>Issue ID:</b>	EFA-5834
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	EPG creation with 4K ctags can result in switch going un-responsive		
<b>Condition:</b>	EPG create operation with 4K ctags.		

<b>Parent Defect ID:</b>	EFA-5936	<b>Issue ID:</b>	EFA-5936
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Overlay-Gateway/Cluster Instance is remaining on the switch.		
<b>Condition:</b>	Configure EPG/Networks on a fabric device and delete device from the fabric.		

<b>Parent Defect ID:</b>	EFA-6318	<b>Issue ID:</b>	EFA-6318
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.1
<b>Symptom:</b>	In Case of Clos Fabric , Cluster drift gets identified for all the devices which are part fabric as soon as fabric gets created even though actual drift is not present on the SLX switch.		
<b>Condition:</b>	This issue observed with below steps, 1)Created CLOS Fabric with multiple MCT-Pairs and a Spine device 2)When checking drift for all devices part of fabric, Cluster Config can be seen as refreshed, but no impact can be seen on fabric show. Also drift status for fabric states false in drift o/p.		

	EFA-7062	<b>Issue ID:</b>	EFA-7062
--	----------	------------------	----------

Parent Defect ID:			
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Device settings for Maintenance mode in EFA is not updated when the device's state changes.		
Condition:	Device may come out of Maintenance mode state based on actions from EFA like Drift-reconcile or a change via CLI in SLX device. The latest state of device will not be automatically updated in the 'efa inventory device settings show' in efa		

Parent Defect ID:	EFA-7080	Issue ID:	EFA-7080
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Tenant commands can take more time (> the timeout =1000 seconds) to complete and hence the failure. The status of timed out fail command can be seen using the relevant "show" command.		
Condition:	The issue can happen in the below scenario: 1. EFA has a scaled tenant configuration. 2. Significant configuration drift introduced between SLX and EFA. 3. Inventory device update is performed resulting in configuration drift derivation of the tenant entities. 4. Tenant commands are executed when (3) is in progress.		

Parent Defect ID:	EFA-7114	Issue ID:	EFA-7114
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Drift is shown for "pseudowire-profile" configuration in command, "efa inventory drift-reconcile execute --ip <device ip>"		
Condition:	This is observed with below steps - Create fabric with back routing enable and configure the fabric. - Create epg and delete the same epg. - Make one of cluster pair device as admin down. - Create same epg which was created and deleted above. - Make the admin down device as admin up. - "pseudowire-profile" still remain as drifted.		

Parent Defect ID:	EFA-7119	Issue ID:	EFA-7119
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0

<b>Symptom:</b>	If any delete nature CLI of "service bgp peer" has failed then bgp service status is not moving to the "delete-pending" state.
<b>Condition:</b>	Any delete nature CLI of "service BGP peer" has failed because any of the below-listed reasons will lead the BGP service to this state. 1. Device connectivity issues 2. Failed to deconfigure the peer on device 3. Targeted topology falls into the APS category.

<b>Parent Defect ID:</b>	EFA-7132	<b>Issue ID:</b>	EFA-7132
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Vlan description is missing after reconciliation		
<b>Condition:</b>	<ul style="list-style-type: none"> <li>- Remove the VLAN configuration from the device</li> <li>- Reload the device with "maintenance mode on reboot enabled"</li> <li>- Description configuration is missing for reconciled VLANs</li> </ul>		

<b>Parent Defect ID:</b>	EFA-7142	<b>Issue ID:</b>	EFA-7142
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Pressing the tab key in the EFA command line when the current command is complete and has no further available options can result in the command being executed silently by the shell.		
<b>Condition:</b>	This is not reliably reproducible.		

<b>Parent Defect ID:</b>	EFA-7144	<b>Issue ID:</b>	EFA-7144
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	'Drift and Reconcile' for a device is skipped when EFA admin of device is performed with 'force' option		
<b>Condition:</b>	When 'efa inventory admin-state up' command is issued with a 'forcce' option		

<b>Parent Defect ID:</b>	EFA-7155	<b>Issue ID:</b>	EFA-7155
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0

<b>Symptom:</b>	Executing the command “efa-sync execute” when one of the SLX switches in cluster is unreachable, could take several minutes to complete. This happens only when multiple neutron entities have to be synced with EFA.
<b>Condition:</b>	The command takes longer time to complete only when: a)Out of sync configuration is large. That is, many neutron entities are out of sync on EFA. b)One of the SLX nodes is not reachable.

<b>Parent Defect ID:</b>	EFA-7159	<b>Issue ID:</b>	EFA-7159
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	MCT Port Channel Admin State is not reconciled		
<b>Condition:</b>	Issue will be observed when 1. MCT PO connecting the leaf pair is administratively shut-down explicitly via out of band management application like CLI and 2. maintenance-mode is set to enable-on-reboot and 3. Device is reloaded. After reload, Drift and reconcile doesn't set the admin state of MCT PO to UP.		

<b>Parent Defect ID:</b>	EFA-7181	<b>Issue ID:</b>	EFA-7181
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Reconciliation of drifted Tenant Configuration fails for scaled configuration		
<b>Condition:</b>	Reconcile fails for a device on which EFA has configured many L3 networks (more than 300+ Networks) by one or more tenants.		

<b>Parent Defect ID:</b>	EFA-7200	<b>Issue ID:</b>	EFA-7200
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	"router bgp" configs will be missing on the devices but Tenant service will not show drift for those configs and hence it will not reconcile those missing configs.		
<b>Condition:</b>	This is observed with the below steps 1. Put SLXs in maintenance mode 2. Remove the router bgp configuration, few VRFs from the device. 3. Reboots both the MCT nodes where TPVMs containing EFA are hosted. 4. Verify after 30 mins, the switches are out of maintenance mode. 5. After a few iterations most of the BGP configs are missing.		

<b>Parent Defect ID:</b>	EFA-7234	<b>Issue ID:</b>	EFA-7234
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0



<b>Symptom:</b>	Reconciliation of drifts in the port-channel and description does not happen		
<b>Condition:</b>	This problem occurs only when drift is introduced specifically on speed and description along with admin-state or cluster-client configuration.		

<b>Parent Defect ID:</b>	EFA-7261	<b>Issue ID:</b>	EFA-7261
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	"efa tenant epg show" and "efa tenant po show" will not reflect correct drift status if aggregator ports are removed.		
<b>Condition:</b>	This display issue is seen only when aggregator port is removed from device.		

<b>Parent Defect ID:</b>	EFA-7304	<b>Issue ID:</b>	EFA-7304
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	While trying to upgrade from secure mode setup, the EFA installer prompts the user to choose between Secure mode and Standard mode.		
<b>Condition:</b>	This will be seen when proceeding through the menu options whenever an upgrade is performed.		

<b>Parent Defect ID:</b>	EFA-7312	<b>Issue ID:</b>	EFA-7312
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Device setting parameters like health check enable, heartbeat miss threshold etc. shows as default values. Drift-reconcile get triggered after default heartbeat miss threshold.		
<b>Condition:</b>	HA double fault failure.		

<b>Parent Defect ID:</b>	EFA-7322	<b>Issue ID:</b>	EFA-7322
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Firmware-download 'Update State' shows 'In Progress' indefinitely when triggered on Active Node's SLX		
<b>Condition:</b>	When SLX firmware upgrade is initiated from the EFA (deployed in HA mode) on the same SLX device where EFA is co-resident and is current active node.		

<b>Parent Defect ID:</b>	EFA-7345	<b>Issue ID:</b>	EFA-7345
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0

<b>Symptom:</b>	A single-node EFA install at version 2.2 failed to upgrade to 2.3.
<b>Condition:</b>	EFA 2.2 to EFA 2.3 upgrade fails only when EFA 2.2 was upgraded (not fresh install) from EFA 2.1.

<b>Parent Defect ID:</b>	EFA-7396	<b>Issue ID:</b>	EFA-7396
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1
<b>Symptom:</b>	If EPG is in delete-pending state with partially cleaned up vlan config on port-channels, deleting the EPG could fail. EPG could remain in delete-pending state unless deleted with force option.		
<b>Condition:</b>	This condition can happen if EPGs in delete-pending state have partially cleaned up vlan configuration on port-channels. That is, the vlan configuration on port-channel is cleaned up on some devices. Deleting the EPG errors out stating that the configuration does not exist.		

<b>Parent Defect ID:</b>	EFA-7404	<b>Issue ID:</b>	EFA-7404
<b>Severity:</b>	S1 - Critical		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1
<b>Symptom:</b>	After firmware download devices from fabric shows a "cfg refresh error" state.		
<b>Condition:</b>	Inventory service discovery happened before lldp links got enabled on devices.		

<b>Parent Defect ID:</b>	EFA-7411	<b>Issue ID:</b>	EFA-7411
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1
<b>Symptom:</b>	EFA firmware download fails intermittently		
<b>Condition:</b>	In cases where the switch config is large, this issue is observed intermittently		

<b>Parent Defect ID:</b>	EFA-7452	<b>Issue ID:</b>	EFA-7452
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1

<b>Symptom:</b>	Maintenance mode triggered DRC (Drift and Reconcile) times out (and fails) with the scaled tenant configuration.
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1) Create Fabric with backup routing enabled and MM enable-on-reboot on all the fabric switches.</li> <li>2) Create tenants.</li> <li>3) Create around 60 POs .</li> <li>4) Create around 8 VRFs with static route bfd.</li> <li>5) Create around 700 L3 EPGs using the POs and VRFs created in (3) and (4) respectively.</li> <li>6) Create around 35 BGP peer group and BGP peers.</li> <li>7) Reboot the fabric switch(es) which will trigger DRC on the rebooted switch(es).</li> </ol>

<b>Parent Defect ID:</b>	EFA-7459	<b>Issue ID:</b>	EFA-7459
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1
<b>Symptom:</b>	EFA firmware download fails		
<b>Condition:</b>	After starting EFA Firmware download the switch does not reboot and firmware download fails intermittently		

<b>Parent Defect ID:</b>	EFA-7538	<b>Issue ID:</b>	EFA-7538
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1
<b>Symptom:</b>	<p>The 'efa login' command does not work and responds with the message                  " CLI Failed.                  EFA services are not operational yet. Please check service state using 'efactl status' command and retry the login once the services are up                  "</p> <p>When executing 'efactl status' command it shows pods in the active node being in "Unkown" state.</p>		
<b>Condition:</b>	Reloading both the SLX nodes multiple times.		

<b>Parent Defect ID:</b>	EFA-7565	<b>Issue ID:</b>	EFA-7565
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Portchannel Deletion failed with error "invalid character 'B' looking for beginning of value"		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1) Tenant service is stopped/not running</li> <li>2) Portchannel delete command run</li> </ol>		

<b>Parent Defect ID:</b>	EFA-7568	<b>Issue ID:</b>	EFA-7568
<b>Severity:</b>	S2 - High		

<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.2
<b>Symptom:</b>	Unable to re-install EFA in a multi-node environment after TPVM upgrade.		
<b>Condition:</b>	If TPVM upgrade is not performed using the procedure described in the deployment guide, under the section 'Upgrade SLX-OS, TPVM, and EFA Together'		

## Defects Closed without Code Changes

The following defects are closed in EFA 2.4.0.

<b>Parent Defect ID:</b>	EFA-2857	<b>Issue ID:</b>	EFA-2857
<b>Reason Code:</b>	Will Not Fix	<b>Severity:</b>	S3 - Medium
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Addition of fabric ports and MCT ports to the tenant succeeds even though these ports cannot be used to connect to any endpoints.		
<b>Condition:</b>	Addition of fabric ports (ports connecting leaf to spine, spine to super-spine etc) and MCT ports (ports connecting the multi-homed leaf) to the tenant.		
<b>Workaround:</b>	Fabric ports and MCT ports to be not added to the tenant.		

<b>Parent Defect ID:</b>	EFA-4045	<b>Issue ID:</b>	EFA-4045
<b>Reason Code:</b>	Not Reproducible	<b>Severity:</b>	S2 - High
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.1.0
<b>Symptom:</b>	MCT configuration of two leafs on already Configured fabric was modified manually by replacing the existing MCT ports with a different ports and later updating inventory causing the Leaf devices to go into refresh error, expectation was to push the config of old MCT ports to new		
<b>Workaround:</b>	Configure the fabric with force efa fabric configure --name <> --force		
<b>Recovery:</b>	Reconfigure fabric with force		

<b>Parent Defect ID:</b>	EFA-5590	<b>Issue ID:</b>	EFA-5590
<b>Reason Code:</b>	Already Reported	<b>Severity:</b>	S2 - High
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	EFA CLI behavior on pressing <tab> following a complete command will execute the command handler or will get stuck in cli if waiting on user input on background		
<b>Condition:</b>	<pre>(efa:root)root@ubuntu:~/efa/efa_40# efa inventory rma execute --ip 10.24.95.157 --co&lt;tab&gt; - cli gets stuck and needs ^C^C to exit, as it is waiting on user input on backend (efa:extreme)extreme@tpvm2:~\$ efa fabric create --name fab4 - &lt;tab&gt; - this creates a fabric</pre>		

<b>Workaround:</b>	These can be avoided by using '?' before the tab and using the complete command options. In most cases this behavior will not cause an issue, but issue may be seen in cases where there are multiple optional keywords eg in fabric-create, rma execute, fabric add bulk , etc.
<b>Recovery:</b>	Press ^C^C to exit a command thats stuck on cli. Reverse any command executed as above, if it is not intended.

<b>Parent Defect ID:</b>	EFA-5689	<b>Issue ID:</b>	EFA-5689
<b>Reason Code:</b>	Feature/Function Not Supported	<b>Severity:</b>	S2 - High
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	VRF Backup routing configuration on the SLXOS-9140 will be inadequate and hence the backup routing functionality will not work on SLXOS-9140.		
<b>Condition:</b>	L3 EPG create or update operation with the member ports residing on SLXOS-9140.		
<b>Workaround:</b>	Backup routing needs to be disabled at the fabric setting level if the fabric has SLXOS-9140 devices.		
<b>Recovery:</b>	Backup routing needs to be disabled at the fabric setting level if the fabric has SLXOS-9140 devices.		

<b>Parent Defect ID:</b>	EFA-5952	<b>Issue ID:</b>	EFA-5952
<b>Reason Code:</b>	Design Limitation	<b>Severity:</b>	S3 - Medium
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	While executing drift-reconcile port-channels created between MCT pairs and are part of EPG's will be shown as drifted even though the SLX running config is in sync.		
<b>Condition:</b>	This happens when drift is executed while SLX is in maintenance mode.		
<b>Recovery:</b>	Turn off maintenance mode and execute inventory device update which will sync the right status.		

<b>Parent Defect ID:</b>	EFA-6835	<b>Issue ID:</b>	EFA-6835
<b>Reason Code:</b>	Already Implemented	<b>Severity:</b>	S2 - High
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	EFA CLIs did not respond following bulk add of devices		
<b>Condition:</b>	Register more 7 devices into EFA at same time using add-bulk fabric CLI in HA multi-node setup.		
<b>Workaround:</b>	Register less then 5 devices in EFA at same time.		
<b>Recovery:</b>	Restart the EFA sever by reloading the Switch.		

<b>Parent Defect ID:</b>	EFA-7237	<b>Issue ID:</b>	EFA-7237
<b>Reason Code:</b>	Already Reported	<b>Severity:</b>	S3 - Medium

<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	EFA system backup is generated using efa system backup command. Undeply EFA and install fresh EFA Restore EFA backup generated earlier using "efa system restore" command The restore process preserves the config backup interval and count. However the auto config backup process is not taking backups as configured.		
<b>Condition:</b>	Auto Config backup process does not take expected number of backups as configured.		
<b>Workaround:</b>	User can take config backup manually using "efa inventory config-backup execute".		
<b>Recovery:</b>	Fresh install or restart of EFA application would start the auto config backup process		

<b>Parent Defect ID:</b>	EFA-7264	<b>Issue ID:</b>	EFA-7264
<b>Reason Code:</b>	Not Reproducible	<b>Severity:</b>	S2 - High
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	After SLX Switch reload, Configuration drift in port channel members do not get reconciled by EFA		
<b>Condition:</b>	Below are the steps when the configuration drift created does not get reconciled by EFA: - Configuration of fabric/tenant/po/vrf/epg - 'maintenance mode enable-on-reboot ' enabled on the switch - Drift in configuration created by deleting port channel - Reloading of the switch - As part of reconciliation, port channel was created but Eth <port> was not aggregated to channel-group.		
<b>Workaround:</b>	None		
<b>Recovery:</b>	Execute manual Drift and Reconcile using this command: " efa inventory drift-reconcile execute --ip <IP of device> --reconcile "		

<b>Parent Defect ID:</b>	EFA-7326	<b>Issue ID:</b>	EFA-7326
<b>Reason Code:</b>	Cannot Fix	<b>Severity:</b>	S3 - Medium
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	VRF used in EndpointGroup having invalid Target Network cause Drift & Reconcile operation to fail with error- "Target network is invalid"		
<b>Condition:</b>	1) VRF is created with Invalid Target Network in earlier version of EFA (before EFA 2.3.0) 2) EndpointGroup is created with same VRF in earlier version of EFA (before EFA 2.3.0) 3) Drift and Reconcile operation is performed		

<b>Workaround:</b>	NA
<b>Recovery:</b>	1) Delete EndpointGroup using VRF having invalid Target Network 2) Delete VRF 3) Recreate VRF with valid Target Network 4) Recreate EndpointGroup with new VRF

<b>Parent Defect ID:</b>	EFA-7389	<b>Issue ID:</b>	EFA-7389
<b>Reason Code:</b>	Will Not Fix	<b>Severity:</b>	S3 - Medium
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	In a very rare scenario, http daemon may not start		
<b>Workaround:</b>	Reload SLX device		

<b>Parent Defect ID:</b>	EFA-7406	<b>Issue ID:</b>	EFA-7406
<b>Reason Code:</b>	Working as Designed	<b>Severity:</b>	S3 - Medium
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.1
<b>Symptom:</b>	When EPG deleted during admin-down for any device, after admin-up for that device app-state changed to cfg-refreshed		
<b>Condition:</b>	1) Configure non-clos fabric 2) Create tennat, vrf, po and epg's 3) Make any of the device to admin down 4) During admin-down, delete one of the epg's created 5) Make same device admin-up 6) EPG app-state moved to cfg-refreshed instead of cfg-in-sync		
<b>Workaround:</b>	NA		
<b>Recovery:</b>	Delete epg once device admin state is up.		

## Open Defects

The following defects are open in EFA 2.4.0.

<b>Parent Defect ID:</b>	EFA-5592	<b>Issue ID:</b>	EFA-5592
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Certificates need to be manually imported on replaced equipment in-order to perform RMA.		

<b>Condition:</b>	RMA/replaced equipment will not have ssh key and auth certificate, in-order to replay the configuration on new switch user needs to import the certificates manually.
<b>Workaround:</b>	import certificate manually efa certificates device install --ips x,y --certType

<b>Parent Defect ID:</b>	EFA-5732	<b>Issue ID:</b>	EFA-5732
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	When firmware download is in progress, fabric delete command is accepted without an error.		
<b>Condition:</b>	If fabric delete command is submitted when firmware download is in progress, it fails.		
<b>Workaround:</b>	Allow firmware download process to complete. Status of the same can be checked using command efa inventory device firmware-download show --fabric {fabric name}		
<b>Recovery:</b>	Fabric can be deleted once the firmware download is completed		

<b>Parent Defect ID:</b>	EFA-5841	<b>Issue ID:</b>	EFA-5841
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	When firmware download is in progress, tenant create command is accepted without an error.		
<b>Condition:</b>	If tenant commands are submitted when firmware download is in progress, it results in erroneous configuration and some configurations may miss.		
<b>Workaround:</b>	Allow firmware download process to complete. Status of the same can be checked using command efa inventory device firmware-download show --fabric {fabric name}		
<b>Recovery:</b>	Tenant commands can be submitted after the firmware download is completed		

<b>Parent Defect ID:</b>	EFA-5874	<b>Issue ID:</b>	EFA-5874
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	On device registration, the IP of the EFA system is recorded in the logging entry on the device so logs can be forwarded to the EFA system for notification. When the EFA system is backed up and restored on another system with a different IP, the old IP of the EFA system is still present on the devices and the devices will continue to forward logs to the old EFA IP.		
<b>Workaround:</b>	Users will have to manually login to each devices and remove the logging entry for the old EFA IP.		

	EFA-5927	<b>Issue ID:</b>	EFA-5927
--	----------	------------------	----------



<b>Parent Defect ID:</b>			
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Configuration reconciliation fails with an error "drift and reconcile failed waiting for status from tenant." because of the timeout.		
<b>Condition:</b>	<p>When the switch configurations drift from the intended configurations in EFA due to scenarios as follows:</p> <ol style="list-style-type: none"> <li>1. L3 Epg is created with large ctag-range (e.g. 2-2000)</li> <li>2. EFA configured VLANs and PO configurations are manually removed from the switch.</li> <li>3. Switch is reloaded in maintenance mode</li> </ol>		
<b>Recovery:</b>	After the switch is moved out of maintenance mode after reload, configuration drift can be viewed and reconciled using "efa inventory drift-reconcile execute --reconcile --ip <switch-ip>" CLI.		

<b>Parent Defect ID:</b>	EFA-5928	<b>Issue ID:</b>	EFA-5928
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.2.0
<b>Symptom:</b>	Configuring devices to default startup-config and adding them to a non-clos fabric does not enable all MCT ports resulting into fabric validation failure for missing link		
<b>Condition:</b>	Added devices immediately after setting to default startup config		
<b>Workaround:</b>	<pre>Remove the devices from fabric and re-add efa fabric device remove --name &lt;fabric-name&gt; --ip &lt;device-ips&gt; efa inventory device delete --ip &lt;device-ips&gt; efa fabric device add-bulk --name &lt;fabric-name&gt; --rack &lt;rack-name&gt; --username &lt;username&gt; --password &lt;password&gt; --ip &lt;device-ips&gt;</pre>		

<b>Parent Defect ID:</b>	EFA-6501	<b>Issue ID:</b>	EFA-6501
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Configuration Drift for VRF still shown in "cfg-in-sync" though its child configuration are drifted on SLX switch.		
<b>Condition:</b>	<p>With below steps issue can be observed.</p> <ol style="list-style-type: none"> <li>1) Create VRF/EPG having route target, static route and bgp configuration.</li> <li>2) Introduce drift in VRF route target or static route or bgp configuration on SLX switch.</li> <li>3) Update device from efa command "efa inventory device update --ip &lt;device ip&gt;"</li> <li>4) Check device drift using efa command as "efa inventory drift-reconcile execute --ip &lt;device ip&gt;"</li> <li>5) VRF shows as "cfg-in-sync" though its child configuration was drifted.</li> </ol>		
<b>Workaround:</b>	None		
<b>Recovery:</b>	After drift and reconcile all EFA and device configuration will be in sync.		

	EFA-7269	<b>Issue ID:</b>	EFA-7269
--	----------	------------------	----------

<b>Parent Defect ID:</b>			
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	During drift reconcile triggered because of maintenance mode enable-on-reboot configuration, interface configurations shown as drifted even though actual drift is not present on the SLX switch.		
<b>Condition:</b>	This issue observed with below steps, - Configured fabric/tenant/po/vrf/epg/bgp peer/peer-group - Enabled maintenance mode enable-on-reboot on SLX switch. - Reload SLX switch. - Drift and Reconcile process shows drift for interface used in EPG which was not drifted.		
<b>Workaround:</b>	None		
<b>Recovery:</b>	EFA and SLX will be in sync when the drift and reconcile (triggered because of maintenance mode enable-on-reboot) is completed.		

<b>Parent Defect ID:</b>	EFA-7324	<b>Issue ID:</b>	EFA-7324
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.0
<b>Symptom:</b>	Continuous create/delete of BGP peer-group and peer can finally cause CLI errors		
<b>Condition:</b>	When create and delete BGP peer/peer-group is repeatedly done in a loop. This will cause inventory does not have chance to update its DB with the current design so DB can be out of sync between inventory and tenant. When other events happen such as timer collection from inventory to sweep config to tenant, it can cause issues on tenant DB where CLI can fail.		
<b>Workaround:</b>	Avoid such cycles of operations		
<b>Recovery:</b>	Delete the BGP peer/peer-group in problem and recreate them again.		

<b>Parent Defect ID:</b>	EFA-7592	<b>Issue ID:</b>	EFA-7592
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.3.2
<b>Symptom:</b>	"dev-state/app-state" moved to not-provisioned/cfg-ready		
<b>Condition:</b>	1) Configure non-clos fabric 2) Create tenant, vrf, epg 3) Admin down device 4) Create multiple epg's , delete an existing epg 5) Manually delete vrf from admin down device 6) Admin up device 7) After admin up, for epg which is in delete-pending theapp-state moved to cfg-ready		

<b>Workaround:</b>	wait for few minutes after epg delete, before admin up of the device.
<b>Recovery:</b>	force delete the EPGs in question and recreate them.

<b>Parent Defect ID:</b>	EFA-8090	<b>Issue ID:</b>	EFA-8090
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	When a fabric containing more than 15 newly registered devices is deployed using the CLI 'efa fabric configure', an attempt to add ports of any of these devices to a tenant within 5 minutes may fail. The error will indicate that the ports have not yet been registered in the fabric		
<b>Condition:</b>	Attempt to add device ports of a recently configured fabric to a tenant may fail with an error indication that the ports have not yet been registered in the fabric		
<b>Workaround:</b>	Wait for up to 5 minutes after deploying the fabric before adding ports to a tenant		
<b>Recovery:</b>	This is a transient error. Rerunning the port-add operation after a maximum wait time of 5 minutes will succeed		

<b>Parent Defect ID:</b>	EFA-8152	<b>Issue ID:</b>	EFA-8152
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	While graceful-restart(GR) updating with value TRUE and inflight transition triggered as a part of EFA rollover then update will continue as a part of inflight transition.		
<b>Condition:</b>	Update GR with value TRUE and perform EFA rollover on HA setup.		

<b>Parent Defect ID:</b>	EFA-8155	<b>Issue ID:</b>	EFA-8155
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	"cluster-client auto" is not configured under port channel for first reloaded device.		
<b>Condition:</b>	<p>Execute below steps to hit this condition</p> <ol style="list-style-type: none"> <li>1) Create fabric on MCT paired device</li> <li>2) Create Tenant/PO/VRF/EPG</li> <li>3) Enable MM mode on both device</li> <li>4) Perform EFA backup</li> <li>5) Delete EPG/VRF/PO/Tenant</li> <li>6) Delete fabric</li> <li>7) Restore EFA backup</li> <li>8) Reload device one by one</li> </ol> <p>After these steps check PO on both device, "cluster-client auto" will not configured on first reloaded device.</p>		
<b>Workaround:</b>	Instead of reload device in Step (8), perform manual DRC using inventory CLI as "efa inventory drift-reconcile execute --ip <device ip> --reconcile" for each device.		

	EFA-8257	<b>Issue ID:</b>	EFA-8257
--	----------	------------------	----------

<b>Parent Defect ID:</b>			
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EFA is not able to detect drift for configuration like VRF/VE/VLAN/EVPN		
<b>Condition:</b>	Please follow below steps, 1) Create tenant/VRF/PO/EPG 2) As soon as EPG creation pushed configuration on device, remove them from device. 3) Check drift using inventory CLI as "efa inventory drift-reconcile execute --ip --device-ip <device ip>"		
<b>Workaround:</b>	As this is timing issue so we need to wait for 1 min before remove configurations from device.		
<b>Recovery:</b>	We need to delete EPG and recreate it again.		

<b>Parent Defect ID:</b>	EFA-8269	<b>Issue ID:</b>	EFA-8269
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG app-state moved to cfg-refresh-err after epg delete and admin up		
<b>Condition:</b>	1) Configure clos fabric (Medium scale fabric) 2) Create tenant 3) Admin down the devices 4) Create port-channels, vrfs and epgs 5) Admin up the following devices Wait for the DRC to be success 6) Repeat step 3 Wait for the devices to put into maintenance mode 7) Create bgp peer-group and dynamic peers 8) Delete all epg's 9) Repeat step 5 10) Vrfs are getting deleted from admin up devices 11) EPG app-state move to cfg-refresh-err		
<b>Recovery:</b>	Delete the EPGs in cfg-refresh-err state and recreate them.		

<b>Parent Defect ID:</b>	EFA-8273	<b>Issue ID:</b>	EFA-8273
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG Update "vrf-add" operation gives success when EPG is in "vrf-delete-pending" state		
<b>Condition:</b>	Perform EPG Update "vrf-add" operation on an EPG in "vrf-delete-pending" state		
<b>Workaround:</b>	No workaround		
<b>Recovery:</b>	User needs to remove the VRF from EPG using EPG update "vrf-delete" operation before attempting the "vrf-add" operation.		

	EFA-8297	<b>Issue ID:</b>	EFA-8297
--	----------	------------------	----------

<b>Parent Defect ID:</b>			
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG update anycast-ip-delete operation succeeded for deletion of provisioned anycast-ip for admin-down device. This issue is observed only if an update anycast-ip-add operation is performed after device is put in admin down state and the new config is in non-provisioned state followed by anycast-ip-delete operation for already configured anycast-ip.		
<b>Condition:</b>	Steps to reproduce issue: 1) Configure EPG with anycast-ip (ipv4/ipv6) 2) Make one device admin-down 3) Anycast-ip update-add new anycast-ip (ipv6/ipv4) 4) Update-delete provisioned anycast-ip configured in step-1 (ipv4/ipv6) Step (4) should fail as IP is already configured on the device and trying to delete it should fail as part of APS.		
<b>Workaround:</b>	No workaround for this.		
<b>Recovery:</b>	Recovery can be done by configuring EPG again with the required configuration using efa or cleaning device config for anycast-ip on the switch.		

<b>Parent Defect ID:</b>	EFA-8315	<b>Issue ID:</b>	EFA-8315
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	User adds ports in empty EPG and immediately deletes them. The following adding ports into EPG can have error as duplicate entry		
<b>Condition:</b>	1) Add ports in empty EPG 2) Delete ports from epg right away 3) Add ports into EPG. Which can have error.		
<b>Workaround:</b>	After adding ports into EPG, wait certain time before trying to delete ports from EPG.		
<b>Recovery:</b>	Delete the EPG and recreate again.		

<b>Parent Defect ID:</b>	EFA-8319	<b>Issue ID:</b>	EFA-8319
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	PO creation failed with error "Devices are not MCT Pairs".		

<b>Condition:</b>	Please follow below steps, 1) Create fabric/tenant/PO/EPG 2) Take EFA backup 3) Delete EPG/PO/tenant/fabric 4) Restore EFA backup taken in step (2) 5) Delete tenant from which was created before backup 6) Create same tenant again 7) Create PO under same tenant
<b>Workaround:</b>	As after restore MCT peer details are Nil so we need to perform DRC after restore taken backup. After step (4) above, we need to perform DRC using inventory CLI as efa inventory drift-reconcile execute --ip <device ip 1> --reconcile efa inventory drift-reconcile execute --ip <device ip 2> --reconcile

<b>Parent Defect ID:</b>	EFA-8322	<b>Issue ID:</b>	EFA-8322
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG Update "anycast-ip-delete" operation gives different output/result when one of the EPG device is admin down		
<b>Condition:</b>	1) Create L3 EPG with anycast-ip/anycast-ipv6 2) Take one of EPG device administratively down 3) Bring device admin up which was taken down in previous step 4) While device is coming up administratively, try EPG Update "anycast-ip-delete" operation		
<b>Workaround:</b>	No workaround		
<b>Recovery:</b>	No recovery as such. Wait for device to be completely up before trying EPG Update "anycast-ip-delete" operation		

<b>Parent Defect ID:</b>	EFA-8334	<b>Issue ID:</b>	EFA-8334
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.5.0
<b>Symptom:</b>	system backup and restore causes epg state to be in cfg-refresh-err		
<b>Condition:</b>	Tenant DB and inventory DB needs time to be in sync. In a busy and scaled system this in sync can take much longer time to finish. Backup DB during DB un-sync window can cause system saves the DBs for tenant and inventory which are not synced yet and following restore will have issues.		
<b>Workaround:</b>	If there's a need to make system backup, please execute the backup after system have not made any new config for few minutes. It's needed for the inventory and tenant databases to be in sync before executing system backup. In a busy system the DB sync can take longer to finish.		
<b>Recovery:</b>	Delete the EPGs which report errors and recreate them.		

<b>Parent Defect ID:</b>	EFA-8335	<b>Issue ID:</b>	EFA-8335
<b>Severity:</b>	S2 - High		

<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	system backup and restore causes following manual DRC has errors		
<b>Condition:</b>	Tenant DB and inventory DB needs time to be in sync. In a busy and scaled system this in sync can take much longer time to finish. Backup DB during this window will cause system saves the DB for tenant and inventory which is not synced yet and following restore will have issues.		
<b>Workaround:</b>	If there's a need to make system backup, please execute the backup after system have not made any new config for few minutes. It's needed for the inventory and tenant databases to be in sync before executing system backup. In a busy system the DB sync can take longer to finish.		
<b>Recovery:</b>	Delete the epg or tenant with problem and recreate them.		

<b>Parent Defect ID:</b>	EFA-8391	<b>Issue ID:</b>	EFA-8391
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	1) NetConf error for epg update port-group-add operation. Caused due to conflicting anycast-ip added to EPG using anycast-ip-add operation. 2) Error for epg update port-group-add operation. Caused due to configuring different anycast-ip for same vlan in EPG using anycast-ip-add operation.		
<b>Condition:</b>	1) If the user provides conflicting anycast-ip in empty EPG then it should throw an error. * Create EPG1 with port/po, vrf, vlan, anycast-ip * Create EPG2 without port/po * Add new VRF to EPG2 * Add conflicting anycast-ip already used in EPG1 with different vrf * Add port to EPG 2) Multiple EPG sharing the same VRF, VLAN with different anycast-ip should throw an error. * Create EPG1 with port/po, vrf, vlan, anycast-ip * Create EPG2 without port/po * Add vrf to EPG1 which is used in EPG2 * Add new anycast-ip to EPG * Add port to EPG (This will cause conflict)		
<b>Workaround:</b>	NA		
<b>Recovery:</b>	Delete conflicting anycast-ip from EPG. Add correct anycast-ip. Add port/po to EPG.		

<b>Parent Defect ID:</b>	EFA-8408	<b>Issue ID:</b>	EFA-8408
<b>Severity:</b>	S4 - Low		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG create failure while testing brownfield scenario by creating EPG with the same VRF/SR/SRBfd configuration as present on the device.		

<b>Condition:</b>	Creating EPG with the same VRF configuration fails due to static route key mismatch. Steps: 1) Create Tenant, Vrf (with staticRoute), EPG 2) Check "sh run vrf" and "sh run router bgp" on device 3) delete EPG 4) Create VRF on device directly using config in step 2 5) Update inventory service 6) Try to create EPG with same VRF. (compareVrf fails due to key mismatch)
<b>Workaround:</b>	No workaround.
<b>Recovery:</b>	1) Remove Vrf from SLX. 2) Update inventory using "efa inventory device update ..." 3) Create EPG using "efa tenant epg create ..."

<b>Parent Defect ID:</b>	EFA-8443	<b>Issue ID:</b>	EFA-8443
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	For Tenant created with L3 port having multiple ip-address associated with it, "efa tenant show" will have repeated entries of that L3 port.		
<b>Condition:</b>	Steps to reproduce issue: 1) Assign multiple IPs to the physical port on SLX. 2) Create Tenant using same L3 port. 3) Check Tenant show output. L3 ports having multiple IPs will have repeated entry in the "efa tenant show" output.		
<b>Workaround:</b>	No workaround.		
<b>Recovery:</b>	Recovery can be done by removing all but one IP from the L3 port on SLX followed by an inventory device update.		

<b>Parent Defect ID:</b>	EFA-8448	<b>Issue ID:</b>	EFA-8448
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	When the ports provided by the user in "tenant update port-delete operation" contains all the ports owned by the port-channel, the PO goes into delete pending state. However, the ports are not deleted from the PO. They get deleted from the tenant though.		
<b>Condition:</b>	This issue is seen when the ports provided by the user in "tenant update port-delete operation" contains all the ports owned by the port-channel resulting in an empty PO.		
<b>Workaround:</b>	User needs to provide ports for "tenant update port-delete operation" which do not result in an empty PO i.e. PO needs to have at least 1 member port.		
<b>Recovery:</b>	Add the ports back using "tenant port-add operation" so that the port-channel has at least 1 member port. The use "efa configure tenant port-channel" to bring the po back to stable state.		

<b>Parent Defect ID:</b>	EFA-8453	<b>Issue ID:</b>	EFA-8453
<b>Severity:</b>	S3 - Medium		



<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	SNMP subscriber add command failed		
<b>Condition:</b>	EFA deployed in standard (non secure) mode		
<b>Workaround:</b>	EFA as SNMP trap proxy is supported in secure mode deployment. Deploy in secure mode.		

<b>Parent Defect ID:</b>	EFA-8465	<b>Issue ID:</b>	EFA-8465
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	The "efa inventory device firmware-download prepare add" command fails with "Please specify 'fullinstall' option in firmware download cmdline as GLIBC versions change".		
<b>Condition:</b>	Upgrading the SLX firmware from 20.1.2x to 20.2.x requires a 'fullinstall' firmware download in order to proceed.		
<b>Workaround:</b>	There is no workaround from EFA. The firmware download fullinstall must be carried out individually on each SLX device.		

<b>Parent Defect ID:</b>	EFA-8472	<b>Issue ID:</b>	EFA-8472
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EFA firmware download fails with the status "Maintenance Mode Enable Failed".		
<b>Condition:</b>	EFA firmware download execution is done with default options and the SLX firmware is being downgraded.		
<b>Workaround:</b>	Use the --noMaintMode flag when performing a "efa inventory device firmware-download execute" command.		
<b>Recovery:</b>	Retry the firmware download execution using the --noMaintMode flag.		

<b>Parent Defect ID:</b>	EFA-8497	<b>Issue ID:</b>	EFA-8497
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG delete fails with error "Service is not available or internal server error has occurred, please try again later"		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1) Create Fabric having leaf pair and border-leaf pair.</li> <li>2) Create Tenant with ports from leaf devices.</li> <li>3) Create VRF using routing type as centralized and border-leaf(s) as centralized routers.</li> <li>4) Create L3 EPGs using VRF created in step3 and ports from leaf devices.</li> <li>5) Remove Border-Leaf devices from fabric or inventory.</li> <li>6) Remove Leaf devices from fabric or inventory.</li> <li>7) Delete EPG(s) created in step4.</li> </ol>		

<b>Workaround:</b>	No workaround
<b>Recovery:</b>	Delete Tenant using 'force' option

<b>Parent Defect ID:</b>	EFA-8507	<b>Issue ID:</b>	EFA-8507
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	Certain vlans are missed in configuration when stacks are created in quick succession within a script with no delay.		
<b>Condition:</b>	10 stack creations without much delay leads to missing configuration. Trunk Subport update is not generated from neutron. Issue is seen with only one controller and not seen when more delay is introduced between stack creations. Trunk also remains in DOWN state.		
<b>Workaround:</b>	Workaround is to have delay between stack creation.		
<b>Recovery:</b>	Remove the Trunk Parent port added to the VM and add it back again. e.g. Max-L2-ss3VirtloVM2_Test1==> VM Name Max-L2-ss3VirtloTrunkPort2_Test1 ==> Parent Port of the Sub Port that is down openstack server remove port Max-L2-ss3VirtloVM2_Test1 Max-L2-ss3VirtloTrunkPort2_Test1 openstack server add port Max-L2-ss3VirtloVM2_Test1 Max-L2-ss3VirtloTrunkPort2_Test1		

<b>Parent Defect ID:</b>	EFA-8512	<b>Issue ID:</b>	EFA-8512
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	On SLX there can have partial config of neighbor under router bgp. The "show running command router bgp" from SLX shows invalid command "neighbor pg1" (Assume the bgp-group name is pg1). There's no corresponding command to delete this.		
<b>Condition:</b>	It's found if issue netconf RPC to SLX device with BGP peer group delete operation which the peer-group does not exist, SLX will create the invalid "neighbor pg1".		
<b>Workaround:</b>	Under some admin-down device scenario, avoid delete the same bgp-peer more than once.		
<b>Recovery:</b>	On SLX use the following commands to get rid of the partial bgp-peer. SLX(config)# router bgp SLX(config-bgp-router)# neighbor pg1 peer-group SLX(config-bgp-router)# no neighbor pg1 peer-group		

<b>Parent Defect ID:</b>	EFA-8526	<b>Issue ID:</b>	EFA-8526
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	VRF Update "centralized-router-add" fail with error "[x, y] are MCT pair. Update the VRF with both devices together as centralized routers"		

<b>Condition:</b>	1) In a CLOS fabric setup with MCT pair of border-leaves, create VRF with routing-type as centralized and select MCT pair of border-leaves as centralized routers. 2) Remove one of the MCT pair border-leaf from the fabric 3) Add same/different border-leaf to the fabric and run fabric configure command 4) Wait for sometime and run VRF Update "centralized-router-add" operation to add newly added border-leaf as centralized router
<b>Workaround:</b>	Run VRF Update "centralized-router-add" operation and specify both nodes of MCT pair border-leaves as centralized routers.
<b>Recovery:</b>	No recovery is required.

<b>Parent Defect ID:</b>	EFA-8527	<b>Issue ID:</b>	EFA-8527
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EFA Update "port-group-delete" operation fail with error "port-group-delete" not allowed for "epg in the port-group-delete-pending state"		
<b>Condition:</b>	Perform EPG Update "port-group-delete" when EPG is in "port-group-delete-pending" state		
<b>Workaround:</b>	Run EPG configure as follows: "efa tenant epg configure --name <epg-name> --tenant <tenant-name>"		
<b>Recovery:</b>	No recovery is required.		

<b>Parent Defect ID:</b>	EFA-8535	<b>Issue ID:</b>	EFA-8535
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	On a single-node installation of TPVM, after ip-change, EFA is not operational.		
<b>Condition:</b>	After IP change of the host system, if 'efa-change-ip' script is run by a different user other than the installation user, in that case, EFA is not operational.		
<b>Workaround:</b>	Restart k3s service using the command 'sudo systemctl restart k3s'		

<b>Parent Defect ID:</b>	EFA-8564	<b>Issue ID:</b>	EFA-8564
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	"device delete" and "device firmware-download prepare add" commands failed		
<b>Condition:</b>	This issue surfaces only when firmware download is triggered from the switch directly. EFA restricts certain operations when firmware download is in progress		

<b>Workaround:</b>	Use EFA to do firmware download
<b>Recovery:</b>	<ol style="list-style-type: none"> <li>1) Confirm firmware download is complete on the device, use "show firmwaredownloadstatus" on the device.</li> <li>2) Update device in EFA using command: "efa inventory device update --ip &lt;device ip list&gt;".</li> <li>3) After this other operation can be performed on the device.</li> <li>4) This flag will automatically reset via periodic device update however periodic device update interval defaults to 1 hour.</li> </ol>

<b>Parent Defect ID:</b>	EFA-8567	<b>Issue ID:</b>	EFA-8567
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG update "vrf-delete" fail with error "Operation 'vrf-event-dev-delete' not allowed for 'vrf-create'"		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1) Create VRF with routing type 'centralized'</li> <li>2) Create EPG using VRF created in above step and ports from leaf devices &amp; border-leaf(s)</li> <li>3) Create another EPG using same VRF and different ports from leaf devices &amp; border-leaf(s)</li> <li>4) Remove border-leaf devices from fabric. After some time, it is observed both EPGs goes to "vrf-delete-pending" state</li> <li>5) Perform EPG update "vrf-delete" on one of the EPG</li> <li>6) Perform EPG update "vrf-delete" on another EPG</li> </ol>		
<b>Workaround:</b>	To remove VRF from EPG, delete the EPG followed by re-create EPG as a workaround.		
<b>Recovery:</b>	No recovery is required.		

<b>Parent Defect ID:</b>	EFA-8568	<b>Issue ID:</b>	EFA-8568
<b>Severity:</b>	S2 - High		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	DRC displayed few vlans as drifted with state "cfg-ready" after creating multiple networks/routers from OpenStack. Device has the config in this scenario.		
<b>Condition:</b>	<ol style="list-style-type: none"> <li>1) Create multiple stacks from OpenStack</li> <li>2) Create multiple networks and routers, both centralized and distributed.</li> <li>3) Ensured all epg/vrfs are created in EFA and on switch</li> <li>4) Ran drift for all leaf and border leaf switches. They are showing drift for the some vlans created.</li> </ol>		
<b>Recovery:</b>	Delete the corresponding EPG/VRF and recreate them.		

<b>Parent Defect ID:</b>	EFA-8573	<b>Issue ID:</b>	EFA-8573
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	In few cases, networks in EPG will remain in cfg-in-sync state even if they are created with partial success topology (MCT pair with one admin-up device and one admin-down device).		

<b>Condition:</b>	The issue is seen with the below steps 1) Configure a fabric 2) Create Tenant 3) Create multi-homed portchannel 4) Bring one of the devices of the MCT pair(having the PO created in step 3) admin-down to create a partial success topology 5) Create EPGs on the partial success topology
<b>Recovery:</b>	Bring all the devices in admin-up state. It should push all the configs on devices and everything will be in cfg-in-sync.

<b>Parent Defect ID:</b>	EFA-8574	<b>Issue ID:</b>	EFA-8574
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	Networks will remain in ACTIVE state when EFA is not reachable during bulk operations		
<b>Condition:</b>	When EFA is not reachable and networks are created in bulk, Network Status on Neutron will remain ACTIVE and Network Status moves to DOWN status only when the Network provisioning is attempted on EFA. This can be delayed based on the number of entries in the Journal.		
<b>Workaround:</b>	No Workaround		
<b>Recovery:</b>	The system would recover on it own after the Network Operation is failed and the status would move to DOWN. If there are pending entries, they would recover when EFA becomes reachable		

<b>Parent Defect ID:</b>	EFA-8584	<b>Issue ID:</b>	EFA-8584
<b>Severity:</b>	S3 - Medium		
<b>Product:</b>	Extreme Fabric Automation	<b>Reported in Release:</b>	EFA 2.4.0
<b>Symptom:</b>	EPG delete fails with error "Service is not available or internal server error has occurred, please try again later"		
<b>Condition:</b>	1) Create Tenant with ports. 2) Create EPG(s) on tenant created in step1. 3) Remove device(s) from fabric or inventory which are used in Tenant created in step1. 4) Delete EPG(s) created in step3.		
<b>Workaround:</b>	Add any device to Tenant followed by deleting EPG(s)		
<b>Recovery:</b>	No recovery is required.		