



Extreme Fabric Automation Release Notes

Version 2.4.2

9036930-02 Rev AA
April 2021



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

Legal Notice

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

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

Trademarks

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

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

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

Open Source Declarations

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




Table of Contents

Preface.....	4
Text Conventions.....	4
Documentation and Training.....	5
Getting Help.....	6
Subscribe to Product Announcements.....	6
Providing Feedback.....	6
Release Notes.....	8
New Features.....	8
CLI Commands.....	10
New commands.....	10
Modified commands.....	10
Supported Platforms and Deployment Models.....	11
Prerequisites for SLX Devices.....	12
Prerequisites for Integration Applications.....	13
EFA Requirements.....	14
General requirements.....	14
High-availability requirements.....	15
EFA Port Requirements.....	16
EFA Installation Modes.....	17
EFA Installation and Upgrade.....	17
Defects Closed with Code Changes	18
Defects Closed without Code Changes.....	30
Open Defects.....	34



Preface

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

Text Conventions

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

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

Table 1: Notes and warnings






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

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
Words in 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.
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</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.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member</i> [<i>member</i> ...].
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

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

[Extreme Optics Compatibility](#)

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

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

Getting Help

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

Extreme Portal

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

The Hub

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

Call GTAC

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

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

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

Subscribe to Product Announcements

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

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

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

Providing Feedback

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

- Content errors, or confusing or conflicting information.

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

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

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

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



Release Notes

- [New Features on page 8](#)
- [CLI Commands on page 10](#)
- [Supported Platforms and Deployment Models on page 11](#)
- [Prerequisites for SLX Devices on page 12](#)
- [Prerequisites for Integration Applications on page 13](#)
- [EFA Requirements on page 14](#)
- [EFA Port Requirements on page 16](#)
- [EFA Installation Modes on page 17](#)
- [EFA Installation and Upgrade on page 17](#)
- [Defects Closed with Code Changes on page 18](#)
- [Defects Closed without Code Changes on page 30](#)
- [Open Defects on page 34](#)

New Features

Extreme Fabric Automation 2.4.2 provides the following enhancements.

Table 4: Features and improvements

Feature	Description
Journal clearing	New options for the efa-journal clear command let you clear journal entries by state: pending, failed, completed, processing, all. For more information, see the command topic in the <i>Extreme Fabric Automation OpenStack Integration Guide</i> .
Journal health	Output for the efa-health show and efa-health show --advanced commands now provides information about journal health: journal size, size threshold, whether the size exceeds the threshold, and whether the status is up or down. For more information, see the command topic in the <i>Extreme Fabric Automation OpenStack Integration Guide</i> .
L3 IP MTU	New support for the <code>ip-mtu</code> option at the network level. For more information, see the "Layer 3 IP MTU" topic in the <i>Extreme Fabric Automation OpenStack Integration Guide</i> .

Table 4: Features and improvements (continued)

Feature	Description
Mechanism driver and L3 service plug-in	The <code>extreme_efa_v2</code> and <code>extreme_13_efa_v2</code> plug-ins now use the journaling mechanism. You can use the previous plug-ins to enable non-journaling. For more information, see the following topics in the <i>Extreme Fabric Automation OpenStack Integration Guide</i> . <ul style="list-style-type: none"> Mechanism Driver and Plug-ins Enable the Non-Journal Plug-in
L3 flavor configuration	Configuration command examples are updated. For more information, see the "L3 Flavors" topic in the <i>Extreme Fabric Automation OpenStack Integration Guide</i> .
L3 routers	Support for creating routers in distributed mode is disabled. For more information, see the "L2 Routers" topic in the <i>Extreme Fabric Automation OpenStack Integration Guide</i> .
Multiple management IP networks	You can add management routes in multi-node deployments during installation and upgrade. You can also use the CLI to create, delete, and show these management routes. For more information, see the following guides: <ul style="list-style-type: none"> <i>Extreme Fabric Automation Command Line Reference</i> for syntax and examples of the <code>efa mgmt route create</code>, <code>efa mgmt route delete</code>, and <code>efa mgmt route show</code> commands <i>Extreme Fabric Automation Administration Guide</i> for the new "Add and Delete Management Routes" topic <i>Extreme Fabric Automation Deployment Guide</i> for revised multi-node installation and upgrade topics
MTU for endpoint groups	You can configure the value for the maximum transmission unit when you create or update an endpoint group. Use the <code>--ip-mtu</code> parameter (in the format <code>ctag:value</code>) to configure the MTU for the tenant network. This value is then configured on the interface VE on the SLX device. To configure MTU when you create an endpoint group, run the <code>efa tenant epg create</code> command. To configure MTU for an existing endpoint group, run the <code>efa tenant epg update</code> command during <code>vxrf-add</code> or <code>ctag-range-add</code> operations. For more information, see the following guides: <ul style="list-style-type: none"> <i>Extreme Fabric Automation Command Line Reference</i> for syntax and examples of the commands <i>Extreme Fabric Automation Administration Guide</i> for the new "Configure MTU for an Endpoint Group" topic <p>Note: During upgrade from EFA 2.4.1 to 2.4.2, any MTU value configured explicitly on the interface VE remains intact on the SLX device after upgrade.</p>
Neutron 3.0.1-07	For more information, see Supported Platforms and Deployment Models on page 11.

CLI Commands

There are new and modified commands in this release.

New commands

- `efa mgmt route create`
- `efa mgmt route delete`
- `efa mgmt route show`

Modified commands

- `efa tenant epg create`
- `efa tenant epg update`
- `efa-health show`
- `efa-journal clear`
- `openstack router create`

Supported Platforms and Deployment Models

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

Table 5: Bare Metal Deployment Models

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

Table 6: OVA Deployment Models

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

Table 7: TPVM Deployment Models

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

Table 8: TPVM Software Support

TPVM Version	SLX-OS 20.1.2x	SLX-OS 20.2.1x	SLX-OS 20.2.2a	SLX-OS 20.2.2b	SLX-OS 20.2.3	SLX-OS 20.2.3x	SLX-OS 20.3.1	Ubuntu Version	EFA Version
4.0.0-5	Yes	No	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	No	18.04	2.2.0, 2.3.x

Table 8: TPVM Software Support (continued)

TPVM Version	SLX-OS 20.1.2x	SLX-OS 20.2.1x	SLX-OS 20.2.2a	SLX-OS 20.2.2b	SLX-OS 20.2.3	SLX-OS 20.2.3x	SLX-OS 20.3.1	Ubuntu Version	EFA Version
4.1.2	No	No	No	Yes	No	No	No	18.04	2.4.0
4.2.2	No	No	No	No	Yes	No	No	18.04	2.4.0
4.2.3	No	No	No	No	No	Yes	No	18.04	2.4.x
4.2.4	No	No	No	No	No	No	Yes	18.04	2.4.x

Table 9: IP Fabric Topology Matrix

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

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

EFA Version	Neutron Version	SLX-OS Version
2.3.0	2.3.0_19	20.1.2d
2.3.1	2.3.1_02	20.1.2e
2.3.2	2.3.2_02	20.1.2e, 20.2.2a
2.4.0, 2.4.1	3.0.0-23	20.2.2b, 20.2.3, 20.2.3a, 20.2.3b
2.4.2	3.0.1-07	20.2.2b, 20.2.3, 20.2.3a, 20.2.3b, 20.3.1

Prerequisites for SLX Devices

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

Table 11: Prerequisites

Prerequisite	Description
Supported devices	See Supported Platforms and Deployment Models on page 11.
Operating system	See Supported Platforms and Deployment Models on page 11.

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: efa inventory device execute-cli --ip <switch IPs> --command "ntp server <ntp server ip> - config" For more information about NTP requirements, see EFA Requirements on page 14.

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 Supported Platforms and Deployment Models on page 11.

EFA Requirements

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

General requirements

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

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

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

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

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

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

High-availability requirements

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



Note

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

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

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



Note

Modify this script to suit your requirements.

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

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

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

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

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

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

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

```

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

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

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

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

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

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

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

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

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

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

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

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

EFA Port Requirements

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

Table 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 Administration Guide*.

- 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, 2.4.1, and 2.4.2.

Parent Defect ID:	EFA-3512	Issue ID:	EFA-3512
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG create might fail with VNI resource not being available in the fabric		
Condition:	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.		

Parent Defect ID:	EFA-3694	Issue ID:	EFA-3694
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	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.		
Condition:	1. Create EPG with 100 networks. 2. Update EPG with 20 ports/port-channels from one switch using "port-group-add" operation.		

Parent Defect ID:	EFA-3717	Issue ID:	EFA-3717
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.1.0
Symptom:	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.		
Condition:	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.		

Parent Defect ID:	EFA-4056	Issue ID:	EFA-4056
Severity:	S2 - High		

Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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.		
Condition:	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.		

Parent Defect ID:	EFA-5064	Issue ID:	EFA-5064
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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.		
Condition:	Configure L3 EPG with only anycast-ip and no anycast-ipv6.		

Parent Defect ID:	EFA-5257	Issue ID:	EFA-5257
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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.		
Condition:	Issue can be observed when vrf-add and vrf-delete operation is executed on Endpoint Group in quick succession multiple times.		

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

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

Symptom:	EPG update "vrf-add" operation can fail with the reason as vrf to be added has conflicting vrf on the switch.
Condition:	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.

Parent Defect ID:	EFA-5708	Issue ID:	EFA-5708
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Device status remains in cfg-in-sync when devices have mismatch in selective configs w.r.t to intended configs		
Condition:	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		

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

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

Parent Defect ID:	EFA-6318	Issue ID:	EFA-6318
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.1
Symptom:	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.		
Condition:	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	Issue ID:	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

Symptom:	If any delete nature CLI of "service bgp peer" has failed then bgp service status is not moving to the "delete-pending" state.
Condition:	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.

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

Parent Defect ID:	EFA-7142	Issue ID:	EFA-7142
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	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.		
Condition:	This is not reliably reproducible.		

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

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

Symptom:	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.
Condition:	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.

Parent Defect ID:	EFA-7159	Issue ID:	EFA-7159
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	MCT Port Channel Admin State is not reconciled		
Condition:	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.		

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

Parent Defect ID:	EFA-7200	Issue ID:	EFA-7200
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	"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.		
Condition:	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.		

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

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

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

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

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

Parent Defect ID:	EFA-7322	Issue ID:	EFA-7322
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Firmware-download 'Update State' shows 'In Progress' indefinitely when triggered on Active Node's SLX		
Condition:	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.		

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

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

Parent Defect ID:	EFA-7396	Issue ID:	EFA-7396
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	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.		
Condition:	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.		

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

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

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

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

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

Parent Defect ID:	EFA-7538	Issue ID:	EFA-7538
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	<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>		
Condition:	Reloading both the SLX nodes multiple times.		

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

Parent Defect ID:	EFA-7568	Issue ID:	EFA-7568
Severity:	S2 - High		

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

Parent Defect ID:	EFA-8438	Issue ID:	EFA-8438
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Ipv6/ipv4 mtu not set under ve interface , set only for ipv6 nd		
Condition:			
Parent Defect ID:	EFA-8453	Issue ID:	EFA-8453
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	SNMP subscriber add command failed		
Condition:	EFA deployed in standard (non secure) mode		

Parent Defect ID:	EFA-8472	Issue ID:	EFA-8472
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EFA firmware download fails with the status "Maintenance Mode Enable Failed".		
Condition:	EFA firmware download execution is done with default options and the SLX firmware is being downgraded.		

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

Parent Defect ID:	EFA-8527	Issue ID:	EFA-8527
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0

Symptom:	EFA Update "port-group-delete" operation fail with error "port-group-delete" not allowed for "epg in the port-group-delete-pending state"
Condition:	Perform EPG Update "port-group-delete" when EPG is in "port-group-delete-pending" state

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

Parent Defect ID:	EFA-8567	Issue ID:	EFA-8567
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG update "vrf-delete" fail with error "Operation 'vrf-event-dev-delete' not allowed for 'vrf-create'"		
Condition:	<ol style="list-style-type: none"> 1) Create VRF with routing type 'centralized' 2) Create EPG using VRF created in above step and ports from leaf devices & border-leaf(s) 3) Create another EPG using same VRF and different ports from leaf devices & border-leaf(s) 4) Remove border-leaf devices from fabric. After some time, it is observed both EPGs goes to "vrf-delete-pending" state 5) Perform EPG update "vrf-delete" on one of the EPG 6) Perform EPG update "vrf-delete" on another EPG 		

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

Parent Defect ID:	EFA-8574	Issue ID:	EFA-8574
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0

Symptom:	Networks will remain in ACTIVE state when EFA is not reachable during bulk operations
Condition:	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.

Parent Defect ID:	EFA-8577	Issue ID:	EFA-8577
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	efa-journal should give provision to delete all state(pending/failed/inprogress) entries		
Condition:			

Parent Defect ID:	EFA-8584	Issue ID:	EFA-8584
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG delete fails with error "Service is not available or internal server error has occurred, please try again later"		
Condition:	<ol style="list-style-type: none"> 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. 		

Parent Defect ID:	EFA-8674	Issue ID:	EFA-8674
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Efa-Journal maintains failed calls for longer duration once after max retry attempted.		
Condition:			

Parent Defect ID:	EFA-8677	Issue ID:	EFA-8677
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Creating large number of networks using heat templates could cause 504 gateway timeout errors to be returned by neutron		
Condition:	This issue is seen only when scaled configuration is executed using heat templates		

Parent Defect ID:	EFA-8841	Issue ID:	EFA-8841
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1

Symptom:	Router port status on openstack side its misleading and always the port status shows 'DOWN'
Condition:	

Defects Closed without Code Changes

The following defects were closed in EFA 2.4.0 and EFA 2.4.2. No defects were closed without code changes in EFA 2.4.1.

Parent Defect ID:	EFA-2857	Issue ID:	EFA-2857
Reason Code:	Will Not Fix	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Addition of fabric ports and MCT ports to the tenant succeeds even though these ports cannot be used to connect to any endpoints.		
Condition:	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.		
Workaround:	Fabric ports and MCT ports to be not added to the tenant.		

Parent Defect ID:	EFA-4045	Issue ID:	EFA-4045
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.1.0
Symptom:	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		
Workaround:	Configure the fabric with force efa fabric configure --name <> --force		
Recovery:	Reconfigure fabric with force		

Parent Defect ID:	EFA-5590	Issue ID:	EFA-5590
Reason Code:	Already Reported	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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		
Condition:	<pre>(efa:root)root@ubuntu:~/efa/efa_40# efa inventory rma execute --ip 10.24.95.157 --co<tab> - 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 - <tab> - this creates a fabric</pre>		

Workaround:	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.
Recovery:	Press ^C^C to exit a command thats stuck on cli. Reverse any command executed as above, if it is not intended.

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

Parent Defect ID:	EFA-5952	Issue ID:	EFA-5952
Reason Code:	Design Limitation	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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.		
Condition:	This happens when drift is executed while SLX is in maintenance mode.		
Recovery:	Turn off maintenance mode and execute inventory device update which will sync the right status.		

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

Parent Defect ID:	EFA-7237	Issue ID:	EFA-7237
Reason Code:	Already Reported	Severity:	S3 - Medium

Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	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.		
Condition:	Auto Config backup process does not take expected number of backups as configured.		
Workaround:	User can take config backup manually using "efa inventory config-backup execute".		
Recovery:	Fresh install or restart of EFA application would start the auto config backup process		

Parent Defect ID:	EFA-7264	Issue ID:	EFA-7264
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	After SLX Switch reload, Configuration drift in port channel members do not get reconciled by EFA		
Condition:	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.		
Workaround:	None		
Recovery:	Execute manual Drift and Reconcile using this command: " efa inventory drift-reconcile execute --ip <IP of device> --reconcile "		

Parent Defect ID:	EFA-7326	Issue ID:	EFA-7326
Reason Code:	Cannot Fix	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	VRF used in EndpointGroup having invalid Target Network cause Drift & Reconcile operation to fail with error- "Target network is invalid"		
Condition:	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		

Workaround:	NA
Recovery:	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

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

Parent Defect ID:	EFA-7406	Issue ID:	EFA-7406
Reason Code:	Working as Designed	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.1
Symptom:	When EPG deleted during admin-down for any device, after admin-up for that device app-state changed to cfg-refreshed		
Condition:	1) Configure non-clos fabric 2) Create tenant, 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		
Workaround:	NA		
Recovery:	Delete epg once device admin state is up.		

Parent Defect ID:	EFA-8623	Issue ID:	EFA-8623
Reason Code:	Will Not Fix	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	The grep log command returned with "Binary file ./ts/ts-server.log matches". The log file is treated as binary file. This could happen on any log file. Null characters as paddings are added into the log file. This is not an EFA specific issue. It's a broader Linux OS issue when system restarts. NULL bytes in the log are a common effect of restart that prevents the system from cleanly syncing.		
Condition:	When service restarts , there are chances that some log files have null characters filled in the log file.		
Workaround:	Use "grep -a"		

Open Defects

The following defects are open in EFA 2.4.2.

Parent Defect ID:	EFA-5592	Issue ID:	EFA-5592
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Certificates need to be manually imported on replaced equipment in-order to perform RMA.		
Condition:	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.		
Workaround:	import certificate manually efa certificates device install --ips x,y --certType		

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

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

Parent Defect ID:	EFA-5874	Issue ID:	EFA-5874
Severity:	S3 - Medium		

Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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.		
Workaround:	Users will have to manually login to each devices and remove the logging entry for the old EFA IP.		

Parent Defect ID:	EFA-5927	Issue ID:	EFA-5927
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Configuration reconciliation fails with an error "drift and reconcile failed waiting for status from tenant." because of the timeout.		
Condition:	When the switch configurations drift from the intended configurations in EFA due to scenarios as follows: 1. L3 Epg is created with large ctag-range (e.g. 2-2000) 2. EFA configured VLANs and PO configurations are manually removed from the switch. 3. Switch is reloaded in maintenance mode		
Recovery:	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.		

Parent Defect ID:	EFA-5928	Issue ID:	EFA-5928
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	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		
Condition:	Added devices immediately after setting to default startup config		
Workaround:	Remove the devices from fabric and re-add efa fabric device remove --name <fabric-name> --ip <device-ips> efa inventory device delete --ip <device-ips> efa fabric device add-bulk --name <fabric-name> --rack <rack-name> --username <username> --password <password> --ip <device-ips>		

Parent Defect ID:	EFA-6501	Issue ID:	EFA-6501
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Configuration Drift for VRF still shown in "cfg-in-sync" though its child configuration are drifted on SLX switch.		

Condition:	With below steps issue can be observed. 1) Create VRF/EPG having route target, static route and bgp configuration. 2) Introduce drift in VRF route target or static route or bgp configuration on SLX switch. 3) Update device from efa command "efa inventory device update --ip <device ip>" 4) Check device drift using efa command as "efa inventory drift-reconcile execute --ip <device ip>" 5) VRF shows as "cfg-in-sync" though its child configuration was drifted.
Workaround:	None
Recovery:	After drift and reconcile all EFA and device configuration will be in sync.

Parent Defect ID:	EFA-7269	Issue ID:	EFA-7269
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	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.		
Condition:	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.		
Workaround:	None		
Recovery:	EFA and SLX will be in sync when the drift and reconcile (triggered because of maintenance mode enable-on-reboot) is completed.		

Parent Defect ID:	EFA-7324	Issue ID:	EFA-7324
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Continuous create/delete of BGP peer-group and peer can finally cause CLI errors		
Condition:	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.		
Workaround:	Avoid such cycles of operations		
Recovery:	Delete the BGP peer/peer-group in problem and recreate them again.		

Parent Defect ID:	EFA-8090	Issue ID:	EFA-8090
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0

Symptom:	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
Condition:	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
Workaround:	Wait for up to 5 minutes after deploying the fabric before adding ports to a tenant
Recovery:	This is a transient error. Rerunning the port-add operation after a maximum wait time of 5 minutes will succeed

Parent Defect ID:	EFA-8152	Issue ID:	EFA-8152
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	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.		
Condition:	Update GR with value TRUE and perform EFA rollover on HA setup.		

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

Parent Defect ID:	EFA-8257	Issue ID:	EFA-8257
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EFA is not able to detect drift for configuration like VRF/VE/VLAN/EVPN		

Condition:	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>"
Workaround:	As this is timing issue so we need to wait for 1 min before remove configurations from device.
Recovery:	We need to delete EPG and recreate it again.

Parent Defect ID:	EFA-8269	Issue ID:	EFA-8269
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG app-state moved to cfg-refresh-err after epg delete and admin up		
Condition:	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		
Recovery:	Delete the EPGs in cfg-refresh-err state and recreate them.		

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

Parent Defect ID:	EFA-8297	Issue ID:	EFA-8297
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0

Symptom:	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.
Condition:	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.
Workaround:	No workaround for this.
Recovery:	Recovery can be done by configuring EPG again with the required configuration using efa or cleaning device config for anycast-ip on the switch.

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

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

	EFA-8322	Issue ID:	EFA-8322
--	----------	------------------	----------

Parent Defect ID:			
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EPG Update "anycast-ip-delete" operation gives different output/result when one of the EPG device is admin down		
Condition:	<ol style="list-style-type: none"> 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 		
Workaround:	No workaround		
Recovery:	No recovery as such. Wait for device to be completely up before trying EPG Update "anycast-ip-delete" operation		

Parent Defect ID:	EFA-8334	Issue ID:	EFA-8334
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	system backup and restore causes epg state to be in cfg-refresh-err		
Condition:	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.		
Workaround:	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.		
Recovery:	Delete the EPGs which report errors and recreate them.		

Parent Defect ID:	EFA-8335	Issue ID:	EFA-8335
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	system backup and restore causes following manual DRC has errors		
Condition:	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.		

Workaround:	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.
Recovery:	Delete the epg or tenant with problem and recreate them.

Parent Defect ID:	EFA-8391	Issue ID:	EFA-8391
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	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.		
Condition:	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)		
Workaround:	NA		
Recovery:	Delete conflicting anycast-ip from EPG. Add correct anycast-ip. Add port/po to EPG.		

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

Workaround:	No workaround.
Recovery:	1) Remove Vrf from SLX. 2) Update inventory using "efa inventory device update ..." 3) Create EPG using "efa tenant epg create ..."

Parent Defect ID:	EFA-8443	Issue ID:	EFA-8443
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	For Tenant created with L3 port having multiple ip-address associated with it, "efa tenant show" will have repeated entries of that L3 port.		
Condition:	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.		
Workaround:	No workaround.		
Recovery:	Recovery can be done by removing all but one IP from the L3 port on SLX followed by an inventory device update.		

Parent Defect ID:	EFA-8448	Issue ID:	EFA-8448
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	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.		
Condition:	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.		
Workaround:	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.		
Recovery:	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.		

Parent Defect ID:	EFA-8465	Issue ID:	EFA-8465
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	The "efa inventory device firmware-download prepare add" command fails with "Please specify 'fullinstall' option in firmware download cmdline as GLIBC versions change".		

Condition:	Upgrading the SLX firmware from 20.1.2x to 20.2.x requires a 'fullinstall' firmware download in order to proceed.
Workaround:	There is no workaround from EFA. The firmware download fullinstall must be carried out individually on each SLX device.

Parent Defect ID:	EFA-8507	Issue ID:	EFA-8507
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	Certain vlans are missed in configuration when stacks are created in quick succession within a script with no delay.		
Condition:	10 stack creations without much delay leads to missing configuration. Trunk Support 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.		
Workaround:	Workaround is to have delay between stack creation.		
Recovery:	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		

Parent Defect ID:	EFA-8512	Issue ID:	EFA-8512
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	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.		
Condition:	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".		
Workaround:	Under some admin-down device scenario, avoid delete the same bgp-peer more than once.		
Recovery:	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		

Parent Defect ID:	EFA-8526	Issue ID:	EFA-8526
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0

Symptom:	VRF Update "centralized-router-add" fail with error "[x, y] are MCT pair. Update the VRF with both devices together as centralized routers"
Condition:	1) In a CLOS fabric setup with MCT pair of border-leafs, create VRF with routing-type as centralized and select MCT pair of border-leafs 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
Workaround:	Run VRF Update "centralized-router-add" operation and specify both nodes of MCT pair border-leafs as centralized routers.
Recovery:	No recovery is required.

Parent Defect ID:	EFA-8535	Issue ID:	EFA-8535
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	On a single-node installation of TPVM, after ip-change, EFA is not operational.		
Condition:	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.		
Workaround:	Restart k3s service using the command 'sudo systemctl restart k3s'		

Parent Defect ID:	EFA-8573	Issue ID:	EFA-8573
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	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).		
Condition:	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		
Recovery:	Bring all the devices in admin-up state. It should push all the configs on devices and everything will be in cfg-in-sync.		

Parent Defect ID:	EFA-8950	Issue ID:	EFA-8950
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.2
Symptom:	SRIOV supports status is not coming to ACTIVE intermittently as the transition is not triggered from neutron service.		

Condition:	SRIOV subports status down with ss8/ss9 use cases
Workaround:	In the SS8/SS9 usecases ,when we give little delay(2-3secs) for the network trunk set after sub port creation ,works fine. Sub ports coming to active state.

Parent Defect ID:	EFA-8628	Issue ID:	EFA-8628
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	Tenant does not contain any Ports, state of EPG is "vrf-delete-pending" and EPG Update "vrf-delete" fail with error- "EPG cannot be updated with tenant having no ports"		
Condition:	Step1) Create Fabric and Tenant Step2) Create VRF using Routing type as Centralized and Border-leaf devices as Centralized Routers Step3) Create EPG using VRF created in Step2 with Physical Ports/Portchannel from leaf devices Step4) Delete all the devices from Fabric and/or Inventory Step5) Perform EPG Update "vrf-delete" operation		
Workaround:	No workaround is required.		
Recovery:	Delete and re-create EPG(s).		

Parent Defect ID:	EFA-8665	Issue ID:	EFA-8665
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	VRF configuration present on Border Leaf devices		
Condition:	1) Create Fabric and Tenant 2) Create VRF using Routing-type as "Centralized" and Border-leaf devices as Centralized-Routers 3) Create EPG(s) using VRF created in Step2 and Physical Ports/Portchannel from Leaf devices 4) Remove Leaf devices (having the Physical Ports/Portchannel used in the EPG) from Fabric/Inventory		
Workaround:	No workaround.		
Recovery:	Manually clean up VRF configuration from device		

Parent Defect ID:	EFA-8669	Issue ID:	EFA-8669
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	EFA is not reachable after secondary node down/up of 8720. (kube-system pods are not responding)		

Condition:	This issue occurs only when the management cables are unplugged for under 10 seconds and plugged back in. What happens is that the EFA failover keepalived scripts, which does labeling and then switches the pod over to the other node, doesn't finish to completion causing the node labeling to not reflect right status where the VIP is assigned.
Workaround:	Either keep the management ports unplugged for more than 20 seconds or reboot the tpvm.
Recovery:	Reboot the tpvm.

Parent Defect ID:	EFA-8701	Issue ID:	EFA-8701
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	EFA becomes unresponsive and the OS reports that there is no disk space.		
Condition:	Galera is technology used by mysql to keep data in sync between nodes. Occasionally, such a transfer between nodes can fail, resulting in a log file of the failed transaction, and some binary data. These accumulate in /apps/efadata/mysql. EFA will tar.gz these logs hourly and delete; but if there is a rapid accumulation, this process can be slow, and under extreme circumstances, not complete before the next hour starts. Eventually, the system is unable to proceed.		
Workaround:	Delete files matching GRA_*.log in /apps/efadata/mysql. These are not needed by EFA.		
Recovery:	Delete files matching GRA_*.log in /apps/efadata/mysql. These are not needed by EFA. Kill all tar czvf processes working on GRA*.log files.		

Parent Defect ID:	EFA-8754	Issue ID:	EFA-8754
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Shutting both mgmt and RME ports can occasionally cause EFA's kubernetes deployment to become stuck in an initializing state.		
Condition:	Turning off and on both mgmt and RME ports.		
Recovery:	Restart kubernetes on both nodes of the cluster. As root user: \$ systemctl restart k3s		

Parent Defect ID:	EFA-8781	Issue ID:	EFA-8781
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	'efa login' does not work after a double fault with 'rabbitmq' container in 'ContainerCreating' state		

Condition:	When reload is performed on both the devices where EFA is deployed.
Workaround	Restart k3s service on both the nodes using the command 'sudo systemctl restart k3s'

Parent Defect ID:	EFA-8802	Issue ID:	EFA-8802
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	Same epg will be allowed to update with different port properties value, however the first configured value will take effect		
Condition:	<ol style="list-style-type: none"> 1. Create L3 epg e1 with port-group P1 port-group property switch port mode trunk 2. Update epg e1 by port-group-add P2 and port-group property switch port mode as trunk 3. Update same epg e1 by port-group-add P2 and different port-group property switch port mode trunk-no-default-native <p>Above steps are applicable for all port properties like switch port mode , switchport-native-vlan-tagging, switchport-native-vlan, single-homed-bfd-session-type. No error is seen while updating with different port property</p>		
Workaround:	Updating the epg with similar port property values since idempotency doesn't work		

Parent Defect ID:	EFA-8827	Issue ID:	EFA-8827
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1
Symptom:	CEP port will remain in "cfg-refreshed" state even after the DRC is successful.		
Condition:	<ol style="list-style-type: none"> 1. Create EPG with CEP ports. 2. Configure "MM enable on reboot" on the SLX and modify any of the CEP port properties. 3. Reboot the SLX. 4. MM triggered DRC will reconcile all the port properties but the CEP will remain in cfg-refreshed state. 		
Workaround:	CEP in cfg-refreshed state can be removed from the EPG and readded to the EPG.		

Parent Defect ID:	EFA-8848	Issue ID:	EFA-8848
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.1

Symptom:	When a configuration drift is introduced on a physical port on the SLX device, followed by a DRC on the device, the reconcile status displays "Portchannel" and "Vlan" reconciliation status as success even though the same were not drifted.
Condition:	Below are the steps to reproduce the issue: 1) Create Fabric/Tenant using MCT paired SLX devices. 2) Create an EPG using the physical port (P1) from an MCT SLX device (D1). 3) Introduce drift on port (P1) on the SLX device(D1). 4) Perform DRC of the SLX device (D1). 5) DRC executed successfully and all configuration pushed on the SLX device. 6) DRC output displays "Portchannel" and "Vlan" reconciliation status as success even though the same were not drifted.

Parent Defect ID:	EFA-8966	Issue ID:	EFA-8966
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	BGP peer group deletion fails when the deletion is attempted using a tenant not owning the BGP peer group.		
Condition:	Deletion of the BGP peer group which is not owned by an existing tenant.		

Parent Defect ID:	EFA-8967	Issue ID:	EFA-8967
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.4.0
Symptom:	Error during epg update ctag-range-add operation if different anycastIP is provided for new ctag range with same l2-vni already configured in EPG. Ex. Error : EPG: e1(11:11.11.1/24) shares network name with EPG: e1(10:10.10.1/24) should have same Anycast IP		
Condition:	1. Create L3 EPG epg1 with ctag-1, l2vni-1, anycastIP1 2. Update EPG epg1 with ctag-2, l2vni-1, anycastIP2 For the same l2vni, anycastIP must be the same, the same condition will be verified as part of epg validation which provided above-mentioned error.		
Workaround:	1. For same l2vni, anycastIP must be same. 2. Different l2-vni can be used for using different anycastIP.		