Extreme Fabric Automation, 2.3.0

Release Notes

9036773-00 Rev AA
October 2020

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/

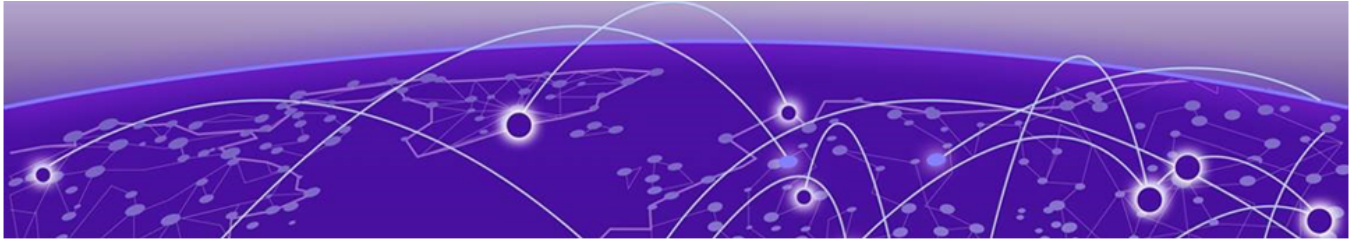

Table of Contents

Preface

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.

Table 1: Notes and warnings (continued)

Icon	Notice type	Alerts you to...
	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 appears on the screen.
The words *enter* and *type*	When you see the word *enter* 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 *type*.
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.
italic 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, *member*[*member*...].
\	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 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.
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.

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.

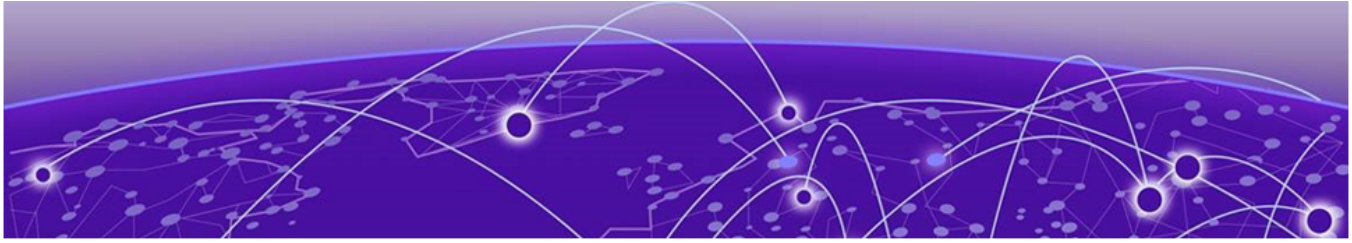

Release Notes

New Features

Extreme Fabric Automation 2.3.0 provides several new features and improvements.

Table 4: 2.3.0 features

Feature	Description
High availability support	You can deploy EFA in a 2-node high-availability cluster. EFA support for high availability also includes handling service restarts after failover. For more information, see the *Extreme Fabric Automation Deployment Guide, 2.3.0*.
Layer 2 and Layer 3 tenant configuration improvements	• VRF static route configuration • EPG configuration • CEP reload delay • CEP cluster tracking • EFA command improvements to support VLAN and bridge domain configuration, EPG configuration • EFA command improvements to support admin state, port speed, MTU, and breakout ports For a list of new and updated commands, see CLI Commands on page 10.

Table 4: 2.3.0 features (continued)

Feature	Description
supportsave enhancement	The **efa show-running-config** command is now part of the **supportsave** script. The supportsave zip file contains output from this command.
Resilient hashing support	Resilient hashing helps to lessen the possibility that a destination path will be remapped when a LAG (link aggregation group) link fails. Use the **--rh-ecmp-enable** and **--rh-max-path** parameters when you create or update a tenant VRF.
Syslog over RELP support for notifications	Any external server that is configured with RELP can be configured as a subscriber to EFA notifications. For more information, see the *Extreme Fabric Automation Administration Guide, 2.3.0*.
SLX 9740 support	EFA now supports ExtremeRouting SLX 9740 devices. For more information, see Supported Platforms and Deployment Models on page 12.
Backup and restore support	New commands let you determine a backup schedule, a purge schedule, a supportsave schedule, and the number of backup files to keep before purging. For more information, see CLI Commands on page 10 and the *Extreme Fabric Automation Administration Guide, 2.3.0*.
EFA deployment	With this release, you can change an IP address or host name after deployment. For more information, see the Administration Guide *Extreme Fabric Automation Administration Guide, 2.3.0*. This release also offers procedures for recovering from upgrade failures. For more information, see the *Extreme Fabric Automation Deployment Guide, 2.3.0*.
Device configuration backup and replay	The **efa inventory config-backup** and **efa inventory config-replay** commands now copy and replay the backup file to the startup-config.
OpenStack integration	This release offers numerous improvements, including TCP keepalive, journal disabling, performance optimization, and link-mapping and host-mapping. For more information, see the *Extreme Fabric Automation OpenStack Integration Guide, 2.3.0*.

CLI Commands

New commands

- **efa inventory admin-state up**
- **efa inventory admin-state down**
- **efa inventory admin-state delete**
- **efa inventory admin-state detail**
- **efa inventory admin-state history**
- **efa inventory admin-state show**
- **efa inventory debug devices-lock**
- **efa notification subscribers add-syslog-relp**
- **efa system backup**
- **efa system backup-list**
- **efa system feature show**
- **efa system feature update**
- **efa system restore**
- **efa system settings show**
- **efa system settings update**
- **efa tenant service bgp peer-group configure**
- **efa tenant service bgp peer-group create**
- **efa tenant service bgp peer-group delete**
- **efa tenant service bgp peer-group show**
- **efa tenant service bgp peer-group update**
- **efa-change-hostname**
- **efa-change-ip**
- **efa-change-vip**
- **efa-sync execute**

Modified commands

- **efa inventory device interface set-admin-state**
- **efa inventory device interface set-breakout**
- **efa inventory device interface set-mtu**
- **efa inventory device interface set-speed**
- **efa inventory device interface unset-breakout**
- **efa inventory device system set-mtu**
- **efa inventory config-replay**
- **efa tenant vrf create**

- **efa tenant vrf update**
- **efa tenant update**
- **efa tenant epg show**
- **efa fabric setting update**
- **efa tenant service bgp configure → efa tenant service bgp peer configure**
- **efa tenant service bgp create → efa tenant service bgp peer create**
- **efa tenant service bgp delete → efa tenant service bgp peer delete**
- **efa tenant service bgp show → efa tenant service bgp peer show**
- **efa tenant service bgp update → efa tenant service bgp peer update**

Deprecated commands

- **efa backup**
- **efa fabric import**
- **efa fabric export**
- **efa restore**
- **efa supportsave**

Supported Platforms and Deployment Models

Table 5: Bare Metal Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirement
EFA 2.1.0	External server (bare metal)	More than 24	Yes	16.04	CPU: 4 cores Storage: 50 GB RAM: 8 GB
EFA 2.2.0 and later	External server (bare metal)	More than 24	Yes	16.04, 18.04	

Table 6: OVA Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirement
EFA 2.1.0	External server (OVA)	More than 24	Yes	16.04	CPU: 4 cores Storage: 50 GB RAM: 8 GB
EFA 2.2.0 and later, in secure mode	External server (OVA)	More than 24	Yes	18.04	

Table 7: TPVM Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS version
EFA 2.1.0	SLX 9150 TPVM SLX 9250 TPVM	Up to 24	Yes	16.04	20.1.1
EFA 2.2.0	SLX 9150 TPVM SLX 9250 TPVM	Up to 24	Yes	18.04	20.1.x

Table 7: TPVM Deployment Models (continued)

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Minimum SLX-OS version
EFA 2.3.0	SLX 9150 TPVM SLX 9250 TPVM	Up to 24	Yes	18.04	20.1.2x
EFA 2.3.0	SLX 9740 TPVM	Up to 24	Yes	18.04	20.2.2

Table 8: TPVM Software Support

TPVM Version	SLX-OS 20.1.1	SLX-OS 20.1.2x	SLX-OS 20.2.1x	SLX-OS 20.2.2	Ubuntu Version	EFA Version
3.0	Yes	Yes	No	No	16.04	2.1.0
4.0.0-5	No	Yes	No	No	18.04	2.2.0
4.0.0-7, 4.0.1	No	Yes	Yes	No	18.04	2.2.0, 2.3.0
4.1.1-0	No	No	No	Yes	18.04	2.3.0

Table 9: IP Fabric Topology Matrix

Platforms	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9030	18x.1.00b	✔				
SLX 9140	18s.1.03	✔				✔
SLX 9240	18s.1.03	✔	✔	✔		
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		✔	✔	✔	✔
SLX 9850	18r.1.00c		✔	✔		

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

Prerequisites for SLX Devices

Table 11: Prerequisites

Prerequisite	Description
Supported devices	See Supported Platforms and Deployment Models on page 12.
Operating system	See Supported Platforms and Deployment Models on page 12.
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 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	Supported Version
Operating system	Ubuntu 16.04, at minimum
OpenStack release	Pike, at minimum
SLX platforms	SLX 9640 as border leaf SLX 9150 as spine and leaf SLX 9150T as spine and leaf SLX 9250 as spine and leaf For a list of supported SLX-OS versions for these devices, see Supported Platforms and Deployment Models on page 12.

EFA Requirements

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: 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
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
30432	Postgres database
30500	Logstash (non-TPVM)
30601	Kibana (non-TPVM)
30672	Rabbitmq
30920	Elasticsearch (non-TPVM)
30930	Elasticsearch (non-TPVM)
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 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.

EFA Installation and Upgrade

For complete information on the following topics, see the *Extreme Fabric Automation Deployment Guide, 2.3.0*.

- 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

Only upgrades from EFA 2.2.0 to EFA 2.3.0 are supported.

Resolved Defects

Parent Defect ID:	EFA-2900	**Issue ID:**	EFA-2900
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	**Reported in Release:**	EFA 2.2.0
Symptom:	"member vlan add <vlan-range>" and "member bd add <bd-range>" configuration doesnt get removed from the "cluster" instance on the switch,		
Condition:	1. "member vlan all" and "member bd all" configuration from the "cluster" instance is removed manually on the switch, 2. EPG creation is performed resulting in "member vlan add <vlan-range>" and "member bd add <bd-range>" configuration under the "cluster" instance on the switch. 3. EPG deletion of the EPG created in step 2 is performed.		
Recovery:	"member vlan add <vlan-range>" and "member bd add <bd-range>" configuration can be manually removed from the "cluster" instance on the switch,		

Parent Defect ID:	EFA-3500	**Issue ID:**	EFA-3500
Severity:	S4 - Low		
Product:	Extreme Fabric Automation	**Reported in Release:**	EFA 2.2.0
Symptom:	Request processing takes longer time and CLI remains hung till the request is processed. EPG delete fails after a wait of more than 16 minutes.		

Condition:	Delete EPG having many (> 256) bridge-domain based networks.
Workaround:	Incrementally delete networks from the EPG and finally delete EPG Example as below: efa tenant epg update --tenant <> --name <> --operation ctag-range-delete --ctag-range 1001-1256 efa tenant epg delete --tenant <> --name <>

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.		
Workaround:	There is no workaround.		

Parent Defect ID:	EFA-3743	Issue ID:	EFA-3743
Reason Code:	Working as Designed	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.1.0
Symptom:	Config Gen reason CLI does not show the detailed reasons if device is in refreshed state		
Workaround:	EFA fabric show CLI shows these reasons, user can refer to that instead of the debug config-gen reason CLI		

Parent Defect ID:	EFA-3970	Issue ID:	EFA-3970
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.1.0
Symptom:	EPG create with multiple ctags using a single bridge-domain results in multiple bridge-domain instances to be created on the switch, instead of a single bridge-domain instance on the switch.		

Condition:	EPG create operation is performed with multiple ctags using a single bridge-domain. Example as below: efa tenant epg create --name vmware31x --tenant t4 --po ncb11 --switchport-mode trunk-no-default-native --ctag-range 300-301 --bridge-domain 300:b2 --bridge-domain 301:b2
Workaround:	EPG update "ctag-range-add" operation can be used to incrementally add multiple ctags using a single bridge-domain. Example as below: efa tenant epg create --name vmware31x --tenant t4 --po ncb11 --switchport-mode trunk-no-default-native --ctag-range 300 --bridge-domain 300:b2 efa tenant epg update --name vmware31x --tenant t4 --operation ctag-range-add --ctag-range 301 --bridge-domain 301:b2

Parent Defect ID:	EFA-4109	Issue ID:	EFA-4109
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Creation of an L3 EPG with only anycast-ipv4 results in the configuration of both IPv4 and IPv6 address-family under router-bgp. Ideally, the IPv6 address-family should get configured only when the L3 EPG has anycast-ipv6		
Condition:	Create an L3 EPG with only anycast-ipv4.		
Recovery:	No recovery needed.		

Parent Defect ID:	EFA-4876	Issue ID:	EFA-4876
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Bridge domain logical-Interface(s) config under the EFA provisioned port-channel is not reconciled on the switch.		
Condition:	1. Remove the EFA provisioned port-channel from the switch. 2. Reload the switch in maintenance mode.		
Recovery:	1. Remove the port-channel interface from the EPGs. 2. Delete and recreate the port-channel interface via EFA. 3. Add the port-channel interface to the EPGs.		

Parent Defect ID:	EFA-5196	Issue ID:	EFA-5196
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	"cluster-client auto" configuration will not be reconciled for the EFA provisioned port-channel on the switch.		

Condition:	Configuration drift created on the switch due to scenario as follows: 1. Remove "cluster-client" config from the EFA provisioned port-channel on the switch. 2. Reload the switch in maintenance mode.
Recovery:	1. Remove the port-channel interface from the EPGs. 2. Delete and recreate the port-channel interface via EFA. 3. Add the port-channel interface to the EPGs.

Parent Defect ID:	EFA-5203	Issue ID:	EFA-5203
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	During device registration, if an error is encountered when installing OAUTH2 cert or HTTPS cert, a warning is generated instead of failing device registration.		
Condition:	This can happen because the system clock in the EFA host and SLX device is not synchronized and the switch will reject the certificate being installed.		
Workaround:	There is no workaround - the device will be registered.		
Recovery:	Ensure that the system time between the EFA host and the devices being used are in sync		

Parent Defect ID:	EFA-5648	Issue ID:	EFA-5648
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG create with multiple ctags using a single bridge-domain results in multiple bridge-domain instances to be created on the switch, instead of a single bridge-domain instance on the switch.		
Condition:	EPG create operation is performed with multiple ctags using a single bridge-domain. Example as below: efa tenant epg create --name vmware31x --tenant t4 --po ncb11 --switchport-mode trunk-no-default-native --ctag-range 300-301 --bridge-domain 300:b2 --bridge-domain 301:b2		
Workaround:	EPG update "ctag-range-add" operation can be used to incrementally add multiple ctags using a single bridge-domain. Example as below: efa tenant epg create --name vmware31x --tenant t4 --po ncb11 --switchport-mode trunk-no-default-native --ctag-range 300 --bridge-domain 300:b2 efa tenant epg update --name vmware31x --tenant t4 --operation ctag-range-add --ctag-range 301 --bridge-domain 301:b2		

Parent Defect ID:	EFA-5697	Issue ID:	EFA-5697
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EFA provisioned port-channel member interface admin state will not be reconciled on the switch.		

Condition:	When the switch drifts in configuration wrt intended configuration on the EFA, due to following conditions: 1. EFA provisioned port-channel member interface admin state is changed to shutdown on the switch. 2. Reload the switch in maintenance mode.
Recovery:	1. Remove the port-channel interface from the EPGs. 2. Delete and recreate the port-channel interface via EFA. 3. Add the port-channel interface to the EPGs.

Parent Defect ID:	EFA-5778	Issue ID:	EFA-5778
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	While using efa firmware-download commands to upgrade switch firmware, intermittently firmware download process gives error "Firmware Activate Failed" for one or more switches		
Condition:	Firmware download process consists of multiple steps, one of which is restart the switch. Upon restart this process tries to connect to the switch and verify if the firmware was successfully finished. Switch reload usually take 1-2 minutes, however very rarely if it take more time, this process waits for switch to restart for 10 minutes. If event after 10 minutes switch restart is not complete, then the it returns this error "Firmware Activate Failed" for that switch		
Workaround:	This is an intermittent issue and there is no workaround to this.		
Recovery:	Start Firmware download process for only the switch with this error from start.		

Parent Defect ID:	EFA-5826	Issue ID:	EFA-5826
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG create CLIs can fail with an error "L3VNI is already consumed in the fabric"		
Condition:	1. Configure EPGs for a tenant. 2. Take the backup of the current configuration of the EFA using "efa show-running-config". 3. Delete tenant with force to clean up all the tenant configuration in EFA and switches. 4. ASCII replay the config backup of "efa show-running-config" taken in step 2.		
Workaround:	Reorder "epg create" CLIs obtained from "efa show-config" in ascending order of l3-vni parameter.		

Parent Defect ID:	EFA-5832	Issue ID:	EFA-5832
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	"channel-group" config of EFA provisioned port-channel member interface will not be reconciled on the switch.		

Condition:	When the switch drifts in configuration wrt intended configuration on the EFA, due to following conditions: 1. Remove the "channel-group" config from the EFA provisioned port-channel member interface. 2. Reload the switch in maintenance mode.
Recovery:	1. Remove the port-channel interface from the EPGs. 2. Delete and recreate the port-channel interface via EFA. 3. Add the port-channel interface to the EPGs.

Parent Defect ID:	EFA-5858	Issue ID:	EFA-5858
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	VRFcreate fails with an error "VRF name <vrf-name> is already taken. Please use different VRF name."		
Condition:	The switch already has a stale or manually created VRF and user is trying to create a VRF (via EFA) with the same name as that of the VRF present on the switch.		
Recovery:	Delete the stale or manually created VRF from the switch before attempting VRF create via EFA.		

Parent Defect ID:	EFA-5869	Issue ID:	EFA-5869
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EFA raslog and notification service stops receiving logs from devices/switch which was restored from backed up data.		
Condition:	Occurs after efa restore		
Workaround:	Restart the raslog service using efactl command		

Parent Defect ID:	EFA-5876	Issue ID:	EFA-5876
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	"cluster member vlan|bd add" and "cluster-client" config reconciliation fails on the switch.		
Condition:	When the switch drifts in configuration wrt intended configuration on the EFA, due to following conditions: 1. EFA provisioned "cluster" is deleted and recreated with another name on the switch. 2. Reload the switch in maintenance mode.		
Recovery:	Remove and re-add all the port-channels to the EPGs.		

Parent Defect ID:	EFA-5945	Issue ID:	EFA-5945
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0

Symptom:	Few of the EPG's creation takes more time (~20 seconds) compared to the other EPGs.
Condition:	Create EPGs (with few networks in each EPG) without any delay between EPG create CLIs.
Workaround:	There is no workaround.

Parent Defect ID:	EFA-5958	Issue ID:	EFA-5958
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG/Networks create fails with error string as below: failed with netconf error '"0000" is an invalid value.' Where pattern of "0000" can be from one "0" to many "00000"s.		
Condition:	1. "cluster" instance on the switch doesn't have "member vlan all" and "member bd all" config. 2. Execute EPG create with some ctags.		
Recovery:	1. Configure "member vlan all" and "member bd all' under "cluster" instance on the switch. 2. Execute inventory update CLI for the switch. 3. Ceate EPG with some ctags.		

Parent Defect ID:	EFA-5960	Issue ID:	EFA-5960
Severity:	S1 - Critical		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	When Switch pumps Raslogs continually (due to faulty SFP) or whatever reason inventory service keeps itself busy in updating device maps. Causing Periodic device discovery not being done. 1.Device discovery will not happen 2.if device is maintenance mode device will not be removed from maintenance mode		
Workaround:	Remove the faulty port or SFP or remove the factor which is causing the switch to pump raslogs continuously .		

Parent Defect ID:	EFA-5965	Issue ID:	EFA-5965
Severity:	S1 - Critical		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Firmware download prepare add command fails without giving any error.		
Condition:	When trying to add a device to firmware download, which is part of a non clos fabric, sometimes the mct pairs are configured incorrectly which causes failure of sanity checks in firmware download prepare add command.		

Workaround:	Intermittent issue, no workaround.
Recovery:	Delete the devices from fabric followed by deleting the fabric. Wait for 2 minutes. Recreate the fabric and then start using the firmware download commands.

Parent Defect ID:	EFA-5966	Issue ID:	EFA-5966
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Further port-group delete operations are not allowed on EPG if there is device connectivity issue while performing port-group-delete operation on EPG		
Condition:	1. No device connectivity from EFA 2. Run EPG Update to delete port-group 3. Operation failed with "Device connection error" 4. Device connectivity to EFA is restored		
Recovery:	1. Delete EPG 2. Re-create EPG without port-group which was intended to be removed previously		

Parent Defect ID:	EFA-5978	Issue ID:	EFA-5978
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Mac move detect not getting reconciled on device		
Condition:	when mac move limit is set mac move detect does not get reconciled on the device as both mac move limit and mac move detect are together pushed to the device		
Workaround:	Introduce a drift on device by removing mac move limit and then reconcile the device, this would configure both mac move limit and mac move detect.		

Parent Defect ID:	EFA-5987	Issue ID:	EFA-5987
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	If the first IP that shows up from the output of hostname -I isn't the IP to be used for registering the devices with Syslog server then an incorrect IP will be picked up.		
Condition:	Incorrect IP gets registered on the device and hence EFA doesnt receive syslog messages from the switches and hence delayed updates.		
Workaround:	Users should ensure that the HOST IP to be registered with the devices as syslog server should be the 1st in the list when executing hostname -I on the EFA server.		
Recovery:	Delete the syslog server that was configured and using efa `execute cli` configure the right IP to be use as the IP for syslog server on the devices.		

Parent Defect ID:	EFA-3743	Issue ID:	EFA-3743
Reason Code:	Working as Designed	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.1.0

Symptom:	Config Gen reason CLI does not show the detailed reasons if device is in refresshed state
Workaround:	EFA fabric show CLI shows these reasons, user can refer to that instead of the debug config-gen reason CLI

Parent Defect ID:	EFA-5830	Issue ID:	EFA-5830
Reason Code:	Not Reproducible	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Unable to update devices using command : inventory device update --fabric {fabric name}		
Condition:	When device update command is submitted immediately after fabric is created, intermittently it returns error with message "fabric" does not exist		
Workaround:	Use the device update command by specifying the ip instead of fabric name e.g. efa inventory device update --ip {comma separated list of ip addresses}		
Recovery:	This issue is intermittent and resolves in few minutes. If it persists, delete the devices from fabric followed by deleting the fabric. Wait for 2 minutes. Recreate the fabric.		

Parent Defect ID:	EFA-5859	Issue ID:	EFA-5859
Reason Code:	Already Reported	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	"efa tenant epg show" CLI execution time increases on each session with the number of parallel executions.		
Condition:	Execute "efa tenant epg show" CLI in parallel from multiple sessions.		
Workaround:	Do not execute "efa tenant epg show" CLI in parallel from multiple sessions. The CLI executions will be queued and processed sequentially.		

Parent Defect ID:	EFA-5941	Issue ID:	EFA-5941
Reason Code:	Already Reported	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	If large number of Endpoint Groups are created and many of them use same VRF, the 'tenant epg show' output displays EPGs in 'cfg-refresh-err' state for some time after Endpoint Group creation. The state moves to 'cfg-in-sync' few minutes after all Endpoing Groups are created.		
Condition:	This issue is seen when multiple Endpoint Groups are created and commands are run one after the other. The time taken to move Endpoint Group status to cfg-in-sync is impacted by number of EFA config and show commands that are being run.		
Workaround:	Workaround is to wait for about 30 minutes after all EFA config and show commands have executed.		
Recovery:	No recovery required		

	EFA-5949	Issue ID:	EFA-5949

Parent Defect ID:			
Reason Code:	Already Reported	Severity:	S2 - High
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	"efa tenant epg show" CLI remains hung for a long time.		
Condition:	Execute "efa tenant epg show" CLI for a tenant having more than 100 EPGs		
Workaround:	Execute "efa tenant epg show" CLI for a given EPG instead of executing it for a given tenant.		

Parent Defect ID:	EFA-5954	Issue ID:	EFA-5954
Reason Code:	Insufficient Information	Severity:	S3 - Medium
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	firmware-download prepare add/remove command fails with error message "Failed to find fabric with name: {fabric name}" even when fabric with this name exists		
Condition:	When firmware-download prepare add/remove command with fabric name as parameter, is submitted immediately after the fabric is created, intermittently it returns error with message "fabric" does not exist		
Workaround:	Use the firmware-download prepare command by specifying the ip instead of fabric name e.g. efa inventory device firmware-download prepare remove --ip {comma separated list of ip addresses}		
Recovery:	This issue is intermittent and resolves in few minutes. If it persists, delete the devices from fabric followed by deleting the fabric. Wait for 2 minutes. Recreate the fabric.		

Open Defects

Parent Defect ID:	EFA-2857	Issue ID:	EFA-2857
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-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.
Workaround:	Provide a wait of 30 seconds between the create and delete CLI on the same EPG.

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.		
Workaround:	Cleanup the out-of-band created LIFs from the switch before EPG create,		
Recovery:	1. Delete the EPG which is in cfg-refresh-err state 2. Clean up the out-of-baned created LIFs from the switch. 3. EFA inventory update to reflect the modified config from the switch. 4. Create the EPG again.		

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.		
Workaround:	Provide a unique anycast-ipv6 for a given BD name across the L3 EPG during EPG create and update operations.		
Recovery:	Remove the subsequent input incorrect anycast-ipv6 from the EPG and provide a unique anycast-ipv6 for a given BD name across the L3 EPG during L3 EPG create and update operations.		

Parent Defect ID:	EFA-4832	Issue ID:	EFA-4832
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	On any operation on EFA services, the error 'Invalid credentials for device 10.x.x.x.' is displayed.		

Condition:	When a device is registered with device credentials in EFA and the same credentials are deleted or updated in SLX.
Workaround:	Any update in device credentials on SLX(deleting the user, changing password etc) has to be updated in EFA using EFA CLI.

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.		
Recovery:	No recovery needed.		

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.		
Workaround:	Workaround is to wait for a few minutes before executing the vrf-add again on Endpoint Group		
Recovery:	No recovery required		

Parent Defect ID:	EFA-5286	Issue ID:	EFA-5286
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Responses to REST requests for VRF Update do not contain details of specific errors that occurred during rollback of errored configuration.		
Condition:	Issue is observed when VRF parameter update is requested, error occurs and rollback that is triggered also encounters error.		
Workaround:	The final error encountered is visible as part of the final error returned to the REST request, in string form.		
Recovery:	No recovery required		

	EFA-5287	Issue ID:	EFA-5287

Parent Defect ID:			
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	Drifted configuration under VRF is not reflected in ‘efa tenant vrf show’ output. The VRF is shown as being in ‘cfg-in-sync’ state.		
Condition:	Issue is observed when configuration under VRF has drifted on the switch due to various reasons.		
Workaround:	Workaround is to use display of ‘efa tenant epg show’ to determine if there is a drift in configuration.		
Recovery:	No recovery required		

Parent Defect ID:	EFA-5590	Issue ID:	EFA-5590
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:	(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		
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-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.		

Workaround:	Provide a wait of 30 seconds between the EPG update vrf-add and vrf-delete operations on the same EPG.
Recovery:	-

Parent Defect ID:	EFA-5689	Issue ID:	EFA-5689
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-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-refresshed		
Workaround:	Drift and Reconcile events fix the problem		
Recovery:	Execute drift and reconcile		

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		

	EFA-5821	Issue ID:	EFA-5821

Parent Defect ID:			
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG create fails after a wait of more than 16 minutes.		
Condition:	Create EPG having more than 256 networks.		
Workaround:	Incrementally add networks to the EPG. Example as below: efa tenant epg create --tenant <> --name <> --ctag-range 2-256 efa tenant epg update --tenant <> --name <> --operation ctag-range-add --ctag-range 257-512		
Recovery:	Verify the EPG/Networks state using "efa tenant epg show" after CLI completes and correct the EPG/Networks state using "efa tenant epg update" CLI.		

Parent Defect ID:	EFA-5822	Issue ID:	EFA-5822
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.2.0
Symptom:	EPG delete fails after a wait of more than 16 minutes.		
Condition:	Delete EPG having more than 256 networks.		
Workaround:	Incrementally delete networks from the EPG and finally delete EPG Example as below: efa tenant epg update --tenant <> --name <> --operation ctag-range-delete --ctag-range 1001-1256 efa tenant epg delete --tenant <> --name <>		
Recovery:	Verify the EPG is deleted using "efa tenant epg show" after CLI completes.		

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.		
Workaround:	Incrementally add networks to the EPG. Example as below: efa tenant epg create --tenant <> --name <> --ctag-range 2-1000 efa tenant epg update --tenant <> --name <> --operation ctag-range-add --ctag-range 1001-2000		

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-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-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.
Recovery:	Manually delete the stale Overlay-Gateway and Cluster instances from the switch.

Parent Defect ID:	EFA-5952	Issue ID:	EFA-5952
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-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.		
Workaround:	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-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. - Create VRF/EPG having route target, static route and bgp configuration. - Introduce drift in VRF route target or static route or bgp configuration on SLX switch. - Update device from efa command "efa inventory device update --ip <device ip>" - Check device drift using efa command as "efa inventory drift-reconcile execute --ip <device ip>" - 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-6832	Issue ID:	EFA-6832
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	If any of the delete nature CLI of EPG/PO has failed then that EPG/PO will be stuck in delete-pending state and no other operation will be allowed on that EPG/PO except delete.		
Condition:	This behaviour is correct if delete operation has failed because of device connectivity or device de-configuration issues. But if it is failed because validations then EPG/PO should remain in their original state and not in the delete-pending state.		
Workaround:	None		
Recovery:	There are 2 recovery options: 1. Fix the reason for failure on the device and re-run the same CLI which will complete the operation and EPG/PO will back to the there expected state. 2. Use --force operation if EPG/PO needs to be completely deleted from EFA.		

Parent Defect ID:	EFA-6835	Issue ID:	EFA-6835
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-7062	Issue ID:	EFA-7062
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		

Workaround:	Re-issue the 'efa inventory device settings update --maint-mode-enable yes ' command to execute the required setting or validate the setting on SLX device using 'show system maintenance'
Recovery:	Re-issue the 'efa inventory device settings update --maint-mode-enable yes ' command to execute the required setting or validate the setting on SLX device using 'show system maintenance'

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.		
Workaround:	Tenant commands should be attempted when the system is in a stable state i.e. when the EFA tenant is not busy deriving the scaled configuration drift of the tenant entities.		

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.		
Workaround:	None		
Recovery:	To recover please follow below steps, - Delete "pseudowire-profile" from SLX switch. - Execute drift reconcile as "efa inventory drift-reconcile execute --ip <SLX switch IP> --reconcile"		

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.
Workaround:	None
Recovery:	There are no side effects of this issue except "update" operations on that BGP service will be allowed which is not expected behavior after the "service bgp peer delete" operation has failed. There are 2 recovery options: 1. Fix the reason for failure and again try to delete BGP service. 2. Delete BGP service using "--force" option. In this case, configs will remain on admin down the device.

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:	- Remove the VLAN configuration from the device - Reload the device with "maintenance mode on reboot enabled" - Description configuration is missing for reconciled VLANs		
Recovery:	Add VLAN description on the device		

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.		
Workaround:	EFA CLI completion can be disabled. This does not affect EFA CLI operation or limit the commands available; it prevents convenient discovery of extra sub-commands and flags by using the tab key from within the bash shell. To disable, do the following: - remove /etc/bash_completion.d/efa - comment out or remove the line "source /etc/bash_completion.d/efa" from /etc/bash.bashrc - start a new shell for EFA CLI use		

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
Workaround:	The Drift and Reconcile can be run manually using the below command: 'efa inventory drift-reconcile execute '

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.		
Workaround:	No workaround available to make the operation faster on EFA		
Recovery:	efa-sync execute becomes faster after the SLX nodes becomes 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.		
Workaround:	Workaround: 1. After the device is reloaded and switch is taken out of maintenance mode, issue a "no shutdown" on the MCT PO interface on the device either through CLI or EFA exec mode. 2. Issue a manual drift and reconcile on the devices where this MCT PO is configured.		
Recovery:	No recovery required		

Parent Defect ID:	EFA-7164	Issue ID:	EFA-7164
Severity:	S3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	For Static Route/Static Route-BFD delete operation, if deletion has failed for any device then configs will be deleted from all other devices but they will remain in EFA for all the devices, and "efa tenant vrf show" CLI will continue to show all of them to the user and VRF will remain in the delete-pending state.		

Condition:	SR/SR-BFD delete operation in vrf update has failed because of bellow conditions 1. device connectivity issues 2. Failed to deconfigure the peer on device 3. Targeted topology falls into the APS category.
Workaround:	None
Recovery:	Delete SR/SR-BFD using the same "VRF update" command after fixing the issue which was causing the failer.

Parent Defect ID:	EFA-7181	Issue ID:	EFA-7181
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Reconcilation 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.		
Recovery:	To reconcile the device having 300+ L3 Networks, first replay the last saved good configuration and execute reconcile command for the device. 1. Replay Config: efa inventory config-replay execute --ip <device-ip> --uuid <saved-config> 2. Reconcile: efa inventory drift-reconcile execute --ip <device-ip>		

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.		
Workaround:	None		
Recovery:	Delete missing configs from EFA also and recreate them.		

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.		

Workaround:	None
Recovery:	Remove the port-channel from device and execute manual reconcile.

Parent Defect ID:	EFA-7237	Issue ID:	EFA-7237
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-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.		
Workaround:	None		
Recovery:	Manual execution of reconcile will sync the device and EFA database. This will in turn resolve the display issue.		

Parent Defect ID:	EFA-7264	Issue ID:	EFA-7264
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 reconcilation, 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-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-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.		
Workaround:	The user should choose secure if the current install is a secure-install. The user should choose non-secure if the current install is a non-secure install.		

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.
Recovery:	Update device setting parameters once again using CLI. $ efa inventory device setting update --ip <Device-ip> --maint-mode-enable-on-reboot <yes> --health-check-enable <yes> --health-check-interval <7m> --health-check-heartbeat-miss-threshold <3> --config-backup-periodic-enable <yes> --config-backup-interval <4m> --number-of-config-backups <10>

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.		
Workaround:	Recommended way to perform an SLX firmware upgrade from the EFA HA cluster being hosted by the same SLX devices: Upgrade the standby EFA node SLX device. Perform HA failover by stop/start active TPVM Upgrade the old active or new standby EFA node SLX device.		
Recovery:	workaround to allow the firmware download status details to be read from the device and updated in inventory DB. This can be triggered by the command - "efa inventory device update'		

Parent Defect ID:	EFA-7326	Issue ID:	EFA-7326
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-7341	Issue ID:	EFA-7341
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0

Symptom:	"efactl status" or any "k3s kubectl" commands respond with the following error message: The connection to the server 127.0.0.1:6443 was refused - did you specify the right host or port?
Condition:	The k3s datastore cluster may fail to restore when the standby node's management interface is shut down for sometime and brought back up again. efa-monitor incorrectly recovered the standby node as a new standalone cluster. Once the management interface was restored, the active node tries to join this new cluster and fails to join because the transactions are ahead of the standby node. The standby node should have remained down in a fault state and joined the current active once the management interface was restored.
Recovery:	The k3s datastore is backed by a MariaDB Galera Cluster. The database cluster can be restored/reformed following the prescribed recovery procedure. https://galeracluster.com/2016/11/introducing-the-safe-to-bootstrap-feature-in-galera-cluster/

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.		
Recovery:	Please see the section "Upgrade from EFA 2.1.0 to 2.2.0 to 2.3.0" in the *Extreme Fabric Automation Deployment Guide, 2.3.0*.		

Parent Defect ID:	EFA-7346	Issue ID:	EFA-7346
Severity:	S2 - High		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Except local-ip other configs are not removed from the "Admin-UP" device after EPG delete operation on partial success topology.		
Condition:	This issue can be observed if EPG deletion containing local-ips has failed because any one of the devices is admin down.		
Workaround:	None		
Recovery:	There are 2 recovery options 1. Bring Admin Down device to Admin UP state and then delete the EPG. 2. Delete EPG with --force option. In this case, configs will be deleted from admin up the device but they will remain on admin down the device.		

Parent Defect ID:	EFA-7351	Issue ID:	EFA-7351
Severity:	P3 - Medium		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Deletion of the EPG fails with an error " Error: Operation "epg delete" not allowed on an "epg in uninitialised state."		

Condition:	The following is the scenario in which the issue can happen: 1.EPG creation is in progress. 2.Tenant service restarts while (1) is in progress. 3.Deletion of EPG (created in step 1) is attempted after the tenant service restart.
Workaround:	None
Recovery:	If the user wants to delete the EPG stuck in "uninitialised state" then the user will need to execute "efa tenant delete --name <tenant-name> --force".

Parent Defect ID:	EFA-7355	Issue ID:	EFA-7355
Severity:	P5 - None		
Product:	Extreme Fabric Automation	Reported in Release:	EFA 2.3.0
Symptom:	Even though the configuration is in sync between the EFA and SLX, the switch doesn't have: 1. "cluster-track" configuration under the CEP port-channel. 2. "graceful-restart" configuration under the router bgp address-family.		
Condition:	The issue can occur in the following scenario: 1. EFA 2.1.0 version is installed. 2. CEP port-channels are configured. 3. L3 EPGs are configured with the CEP port-channel and anycast-ipv4. 4. EFA 2.1.0 is upgraded to EFA 2.2.0. 5. efa-db-upgrade-from-2-1-0.sh is executed. 6. EFA 2.2.0 is upgraded to EFA 2.3.0. 7. Manual drift and reconcile is executed using "efa inventory drift-reconcile execute --ip <switch-ip> --reconcile" CLI.		
Workaround:	None		
Recovery:	A temporary EPG can be created using the "CEP port-channel" and the "VRF" for which the "cluster-track" and "graceful-restart" configuration was missing respectively on the switch, so that the missing "cluster-track" and "graceful-restart" configurations get pushed to the switch during the EPG create.		