July 2020

Extreme Fabric Automation 2.2.0 Release Notes

Part number: 9036661-00 Rev AB

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/.

Contents

Document History

Version	Summary of Changes	Publication Date
1.0	Initial Release for 2.2.0	June 2020

Preface

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, call (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 want to receive notifications.
   Note: You can change your product selections or unsubscribe at any time.
4. Select **Submit**.

Extreme Resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at www.extremenetworks.com. Product documentation for all supported releases is available to registered users at https://www.extremenetworks.com/support/documentation-home/.

Introduction

Extreme Fabric Automation (EFA) is a microservices-based application that manages the life cycle of IP Fabric Clos and Small Data Center deployments. All of the microservices support REST APIs that are detailed by OpenAPI.

EFA offers unique flexibility in supporting multiple IP Fabric topologies based on a BGP underlay with a BGP/EVPN overlay:

- Small Data Center Fabric (non-Clos topology from one switch pair up to four switch pairs)
- 3-stage Clos (Leaf / Spine)
- 5-stage Clos (Leaf / Spine / Super Spine)

Tenant Network onboarding services are supported on all the topologies, which allows you to create connectivity for devices that are connected to the Fabric, such as compute (servers), storage, and any other connectivity such as external routers or gateways.

Lifecycle management of the Fabric allows you to add or delete devices to the Fabric after Day 0. Similarly, you can add or remove Tenants as necessary. Key ecosystem integrations streamline Tenant and network provisioning by way of VMware vCenter, Microsoft System Center for Virtual Machine Management (SCVMM), and OpenStack with ML2 and L3 service plugins.

Key Features

The key features of EFA are as follows:

- Conformance to the EVD (Extreme Validated Design for IP Fabric) - https://www.extremenetworks.com/resources/extreme-validated-design/extreme-ip-fabric-architecture/
- Ease and speed of provisioning and troubleshooting
- Seamless installation and deployment mechanism
- High performance and low resource utilization with minimal touch points
- Programmable containerized services through an industry-standard Open API (https://www.openapis.org/)-based programmable interface
- Easy-to-use CLI commands to manage devices in an IP Fabric and in Tenant Networks
- Ecosystem support for OpenStack, VMware vCenter, and Microsoft Hyper-V

Supported Platforms, Software, and Topologies Matrix

Platforms	SLX-OS Release	Leaf	Spine	Super Spine	Border Leaf	Small DC Fabric
SLX 9140	18s.1.01, 18s.1.01a, 18s.1.01c, **18s.1.03**	✓				✓
SLX 9240	18s.1.01, 18s.1.01a, 18s.1.01c, **18s.1.03**	✓	✓	✓		
SLX 9030	18x.1.00, 18x.1.00a, **18x.1.00b**	✓				
SLX 9540	18r.1.00aa, 18r.1.00b,18r.1.00c, 18r.1.00cc				✓	
SLX 9540	**20.1.x**	✓			✓	
SLX 9640	**20.1.x**				✓	
SLX 9850	18r.1.00aa, 18r.1.00b, **18r.1.00c**		✓	✓		
SLX 9150	**20.1.x**	✓				✓
SLX 9250	**20.1.x**	✓	✓			✓

Supported Deployment Models

Table 1 Bare Metal Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
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	External Server (Bare metal)	More than 24	Yes	16.04, 18.04	CPU: 4 cores Storage: 50 GB RAM: 8 GB

Table 2 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 3 TPVM Software Support Matrix

Version	SLX-0S 20.1.1	SLX-OS 20.1.2x	Ubuntu Version	EFA Version
TPVM 3.0	Yes	Yes	16.04	2.1.0
TPVM 4.0	No	Yes	18.04	2.2.0

Table 4 OVA Deployment Models

Version	Deployment	Managed SLX Devices	Multi-Fabric Support	Ubuntu Version	Server Requirements
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 (Secure mode)	External Server (OVA)	More than 24	Yes	18.04	CPU: 4 cores Storage: 50 GB RAM: 8 GB

Notes:

- EFA 2.1.0 is not supported on Ubuntu 18.04.
- Support of EFA 2.2.0 on TPVM 3.0 is not planned due to security enhancements in EFA and SLX-OS.
- EFA 2.2.0 is not supported on Ubuntu 19.04/20.04 LTS.

EFA Core Services and Integrations

EFA comprises several core containerized services that interact with each other and with other infrastructure services to provide the core functions of IP Fabric automation.

Service	Description
Inventory Service	Inventory service manages the devices in fabric. Provides the secure credential store and deep discovery of physical and logical assets of the managed devices. Publishes the Asset refresh and change events to other services.
Fabric Service	Helps orchestrate and visualize the BGP and EVPN-based 3- and 5-stage CLOS networks.
Tenant Service	Helps manage the Tenants, Tenant Networks, and EPG, fully leveraging the knowledge of Assets and the underlying fabric.
Notification Service	Sends events, alerts, and task updates to external entities.
RASlog Service	Acts as a syslog server to process syslog messages from devices. Also acts as an SNMP trap receiver to process traps from devices.
Authentication and Authorization Service	Enforces a security boundary between northbound clients and the downstream operations between EFA and SLX.

EFA also provides a microservice for each ecosystem integration. This architecture permits rapid development and integration of different ecosystem integrations. Each operates independently to externally integrate while using the same underlying services to interact with the IP Fabric.

Ecosystem	Description
VMware vCenter	The vCenter integration provides connectivity between EFA and vCenter using a REST API as documented in the VI SDK. EFA does not connect to individual ESXi servers. All integration is done through vCenter.
OpenStack	OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter.
Microsoft Hyper-V	The Hyper-V integration supports networking configuration for Hyper-V servers in a datacenter, manual and automated configuration updates when VMs move, and visibility into the VMs and networking resources that are deployed in the Hyper-V setup.

New Features

Feature	Description
Notification service	Sends events, alerts, and task updates to external entities.
Layer 3 tenant network services	Provides support for IPv6 Anycast gateways; VRF backup routing, static routing, BFD static routing, local ASNs, and graceful restart; CEP reload delay, LACP timeout for port channels, distributed routing, and sharing L2 resources across tenants.
Security enhancements	Provides support for SSO, RBAC, CIS-CAT auditing, risk assessment, vulnerability analysis, and security hardening.
Data consistency	EFA is the data owner and Single Source of Truth (SSOT) for fabric configuration. This feature ensures that SLX devices have the correct configuration before allowing traffic.
Redundant Management Network	Provides fault tolerance for the management path. This is done using Linux bonding by pairing the physical Management port of the chassis with any one of the physical front panel User Ports.
RASlog service	Acts as a syslog server to process syslog messages from devices. Also acts as an SNMP trap receiver to process traps from devices.
New inventory 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

Prerequisites for SLX Devices

Prerequisite	Description
Supported devices	See Supported Platforms, Software, and Topologies Matrix.
Operating System	See Supported Platforms, Software, and Topologies Matrix.
IP addresses	Management IP addresses configured on all devices
Security	Proper certificate installation and audit log requires NTP configuration on SLX For example: "efa inventory device **execute-cli** --ip <switch IPs> --command " ntp server <ntp server ip> -config"

Prerequisites for the vCenter Controller and Compute Nodes

Minimum System Requirements	Description
vCenter Controller	Windows 2016 with vSphere 6.5.x and 6.7.x
Compute Nodes	ESXi 6.x or later

Prerequisites for the SCVMM Controller and Compute Nodes

Minimum System Requirements	Description
SCVMM Controller	Windows SCVMM 2016
Compute Nodes	Windows Hyper-V 2016

Prerequisites for the OpenStack Controller and Compute Nodes

Minimum System Requirements	Description
Operating System	Ubuntu 16.04
OpenStack Release	Pike
Extreme Neutron plugin	2.2.0 (on controller) Extreme Neutron plugin is delivered as a Debian package.

OpenStack EFA Neutron Plugin

1. Single-segment Neutron Network creation and deletion:
    - Internal VLAN network
    - VLAN provide network (default)

2. Create and delete networks and VMs from Dashboard, Horizon, Neutron CLI, and OpenStack CLI.
3. VM operations such as pause, suspend, shelve, resize, shut, lock, and reboot.

Note: SLX device configurations that are pushed from OpenStack should not be modified either through EFA or direct CLI access to the device. Doing so results in the EFA Database going to an inconsistent state. Configurations should only be modified using relevant OpenStack commands.

EFA Installation Prerequisites

For details on various EFA installation types, TCP and UDP port availability, and other prerequisites, see the "Deploying Extreme Fabric Automation" section of the *Extreme Fabric Automation Administration Guide, 2.2.0*.

EFA Installation Modes

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

- **Secure mode**: Traffic to EFA uses the HTTPS protocol. All non-HTTP requests are redirected to the secure port.
- **Standard mode**: Traffic to EFA uses the HTTP protocol.

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

Fresh Install of EFA on TPVM – Single Node

Prerequisite: For details about installing and upgrading TPVM, see the *Extreme SLX-OS Software Upgrade Guide*.

EFA on TPVM is supported only on the SLX 9150 and SLX 9250 platforms.

1. On the SLX where TPVM is planned to be run, verify that TPVM is set up for EFA deployment:
    - Validate that the TPVM is running with 4.0.0 and SLX-OS is running with at least minimum requirements.
      ```
      # show tpvm status
      # show version
      # lsb_release -a (on TPVM should show Ubuntu 18.04)
      ```
    - Validate that TPVM has an assigned IP address.
      ```
      # show tpvm ip-address
      ```
    - Validate that the SSH keys are uploaded.
      ```
      # show tpvm status
      ```
    - Validate that passwordless access is configured.
      ```
      # show tpvm status
      ```
    - Configure NTP on TPVM.
      ```
      # tpvm config ntp add server <ip>
      ```
    - Validate that TPVM and NTP are synchronized.
      ```
      # show tpvm config ntp
      ```
    - If required, log in into TPVM and configure the NTP time zone.
      ```
      # sudo timedatectl set-timezone your_time_zone
      ```

2. Enter SLX Linux mode.
   ```
   # start-shell
   # cd /efaboot
   ```
3. Copy the EFA tar file to the SLX device.
   ```
   # scp <efa-bundle>
   ```

4. Deploy EFA on TPVM from the SLX shell, select single-node and the required security mode (secure or non-secure) from the deployment menu.
   ```
   # efa deploy
   ```
5. Verify the status of the deployment.
   ```
   # show efa status
   ```

Fresh Install of EFA on a Server – Single Node

1. Download the image (*.tar.gz) and untar it.
   ```
   # tar -xfz efa-v2.2.0.tar.gz
   ```
2. Verify the PGP signature as described in Verify PGP Signatures.
3. Change directory to a single-node deployment.
   ```
   # cd efa
   ```
4. Configure NTP with Ubuntu commands.
   ```
   # sudo vim /etc/systemd/timesyncd.conf
   # sudo service systemd-timesyncd restart
   # systemctl status systemd-timesyncd
   # sudo timedatectl set-timezone your_time_zone
   ```
5. Check for prerequisites.
   - CPU: 4 cores
   - Storage: 50 GB
   - RAM: 8 GB
   - OS: Ubuntu 16.04/18.04
6. Run the application installation script, select single-node, and select the required security mode (secure or non-secure) from the deployment menu.
   ```
   # source deployment.sh
   ```

Fresh Install of EFA on a VM using OVA

Open Virtual Appliance (OVA) is an OVF file packaged with base image Ubuntu 18.04 (storage = 50 G, RAM = 8 GB) and preinstalled with EFA Secure mode. OVA is also compatible with VMware ESXi servers, so it can be deployed with VMware products.

Prerequisites:

- The virtual machine (VM) on which you deploy the OVA requires a network adapter with a valid IP address and DNS. The IP address is required to configure the SLX devices to forward syslog entries back to the VM. The VM needs DNS configuration to resolve the URL during setup and forwarding of events to the notification subscriber.
- The VM must be able to access switches and the notification subscriber.
- For networks without DHCP, you must assign valid, static IP addresses and DNS. Then reboot the VM. Ensure that all services are up and running before running commands.

Warnings:

- Do not change the host name of the .ova VM after deploying the .ova image. Doing so prevents EFA PODs from coming online.

- Use the OVA image only for new installations. For existing deployments, see Upgrade EFA on a Server or Upgrade EFA on TPVM.
- EFA v2.2.0 OVA is not supported for Oracle VirtualBox because the syslog service requires port forwarding for port 514 on UDP. The source IP address of the syslog message will be changed from the SLX device to the host IP, which the syslog service ignores.

Notes:

- When the VM starts up, a start-up script checks whether the IP address of the primary interface eth0 has changed since it was last configured. If the IP address has changed, the script updates the EFA profile and configuration files appropriately and reapplies the k3s application deployment template. This operation takes a few minutes to complete. On subsequent VM reboots, if the IP address has not changed, no operation is performed by the start-up script. The logs are located under **/var/log/efa/installer**.
- On normal bootup, the EFA services do not come up instantly. The k3s service takes a few minutes to do health checks and then re-initialize containers. If you try to log in to EFA while the services are being initialized, a message indicates that the services are not operational yet.
- The user credentials for the OVA installation are:
  - admin/password
  - ubuntu/ubuntu
- Log in with "admin" user and then use sudo to run the commands (**sudo efa supportsave**, **sudo efa backup**, **sudo efa restore**).

EFA Upgrade Prerequisites

- Back up the EFA 2.1.0 database to TPVM **/apps** using the **/apps/efa/efa_backup.sh** command. The backup and restore scripts are packaged with EFA 2.1.0. Create a backup directory for this step.
- Ensure that the required TCP and UDP ports are available. the "Deploying Extreme Fabric Automation" section of the *Extreme Fabric Automation Administration Guide, 2.2.0*.

Upgrade EFA on TPVM

1. Verify that the TPVM is set up for EFA deployment:
   a. Validate that the TPVM is running with 4.0.0 and SLX-OS is running with at least the minimum requirements.
      ```
      # show tpvm status
      # show version
      # lsb_release -a (on TPVM should show Ubuntu 18.04)
      ```
   b. Validate that the TPVM has an assigned IP address.
      ```
      # show tpvm ip-address
      ```
   c. Validate that the SSH keys are uploaded.
      ```
      # show tpvm status
      ```
   d. Validate that passwordless access is configured.
      ```
      # show tpvm status
      ```
   e. Configure NTP on the TPVM.
      ```
      # tpvm config ntp add server <ip>
      ```

f. Validate that the NTP is synchronized.
```
# show tpvm config ntp
```
g. If required, log in to TPVM and configure the NTP timezone.
```
# sudo timedatectl set-timezone your_time_zone
```

2. Determine whether more than one EFA version is available in the SLX `/efaboot` directory.
   a. If no version is available, the installer stops.
   b. If more than one version is available, you have the option to pick a version.
   c. If only one version is available, the installer picks up that version.
3. Determine whether the TPVM already has a version of EFA installed.
   a. If the same version is already installed, the installer stops.
   b. If no EFA is installed, the installer continues with installation.
   c. If a different version is detected, the upgrade will continue, depending on the detected version.
4. Copy the EFA tar file to the SLX device.
```
# start-shell
# cd /efaboot
# scp <efa-bundle>
```
5. Deploy EFA on the TPVM from the SLX device.
```
# efa deploy
```
6. Perform the steps in Post-Upgrade Validation Process.

Upgrade EFA on a Server

1. Download the image (*.tar.gz) onto a new sub-folder and untar it.
```
# tar -xfz efa-v2.2.0.tar.gz
```
2. Verify the PGP signature as described in Verify PGP Signatures.
3. Change directory to a single-node deployment.
```
# cd efa
```
4. Configure NTP with Ubuntu commands.
```
# sudo vim /etc/systemd/timesyncd.conf
# sudo service systemd-timesyncd restart
# systemctl status systemd-timesyncd
# sudo timedatectl set-timezone your_time_zone
```
5. Run the application installation script and required modes from the menu: secure or standard (non-secure).
```
# source deployment.sh
```

   If the previous deployment stack is already running, the script presents the following options:
   - **Remove the current stack**: You can remove the entire stack with this option.
   - **Upgrade or Redeploy**: If you are running the `deployment.sh` script from the new tar-ball, you can upgrade the setup without wiping out the current database volume. You have the same option if you rerun the script from the same folder, in which case the stack is only redeployed.
   - **Quit**: No change in the current stack.
6. Perform the steps in Post-Upgrade Validation Process.

Upgrade EFA on an OVA

1. Log in to the OVA as admin.
2. Change to the root by running the following command.
   # sudo su
3. Copy the new tar file under **/opt/godcapp/**.
4. Extract the tar.
   # tar –xvf efa-2.2.0.tar.gz
   # cd efa
5. Run the deployment script.
   # source deployment.sh
6. Select the **Upgrade** option.
7. When the upgrade is complete, run the following command to set the environment variable.
   # source /etc/profile.d/efa_env.sh

Now, EFA commands can be run only by the root user. To run commands as an admin user, a user needs to be assigned the SystemAdmin role.

1. Log in as admin user and switch to root user via the **sudo su** command.
2. Run the following command to assign a valid role to an admin user.
   # efa auth rolemapping add --name admin --type user --role SystemAdmin
3. Validate the role by running the following command.
   # efa auth rolemapping show

Ensure that the System is Ready

1. After any of the following scenarios, wait 10 minutes for EFA microservices to be operational before you run EFA commands.
   - Powering on the OVA
   - Rebooting the OVA
   - Rebooting the TPVM
   - Rebooting the SLX (which also reboots the TPVM)
   - Rebooting the server on which the TAR is installed
2. Run the following command to verify that all PODs are in a running state:
   `# k3s kubectl get pods -n efa`
3. Perform the steps in Post-Upgrade Validation Process.

Post-Upgrade Validation Process

1. Verify the status of EFA.
   `# show efa status`
2. Verify the status of the Fabric.
   `# efa fabric show`
3. Verify the status of the tenant.
   `# efa tenant show`

4. Verify the status of the EPG (endpoint group).
   `# efa tenant epg show`
   If any network is in **cfg-refreshed** state, run the following commands:
   - Run this command for each device in the tenant.
     `efa tenant debug device drift --device-ip <ip-address> --reconcile`
   - `efa inventory device update --fabric <fabric-name>`

Migrate the SLX Connection from Non-Secure to Secure during EFA Upgrade

Devices that are running SLX-OS 20.1.2a and migrated from EFA 2.1.0 to 2.2.0, with a connection from EFA to SLX can be made secure with the following steps.

1. Update the inventory in EFA 2.2.0 for this Fabric to make EFA aware of the new SLX versions.
2. Install certificates and configure OAuth2 in SLX from EFA using the following command: **efa certificate device install --ips <>**
3. Verify the certificates in SLX using the following commands:

   - **show crypto ca certificates**
   - **show run aaa**

   These commands should display two certificates installed and OAuth2 configuration in each device.

4. Update the inventory in EFA for this Fabric to synchronize EFA with SLX. Verify the inventory device list in EFA and ensure secure mode status is updated with "Y" for these devices.
5. Proceed with Fabric and tenant operations in SLX. Observe from the output of the SLX **show logging auditlog** command that these operations are performed by using the logged in account in EFA.

Note: As a best practice, deploy EFA 2.2 in secure mode to make northbound interfaces also secure.

Verify PGP Signatures

For any firmware image that you download from Extreme Networks, Inc., validate the PGP signature of the downloaded data. The process for performing this verification is detailed in this topic. You can also access it online here: https://extremeportal.force.com/ExtrArticleDetail?n=000048172.

The PGP signature verification key appears in Appendix: PGP Signature Verification Key.

The overall process for signature verification is as follows:

1. Download the Extreme Networks, Inc., PGP signing verification key. Import the key into a local instance of PGP or GPG. In this document all examples use GPG from the Linux command line.
2. Download a firmware image from the Extreme Networks, Inc., firmware download site.
3. Download the corresponding SHA256SUMS, SHA256SUM.asc, SHA512SUMS, and SHA512SUMS.asc files.

4. Verify that the PGP signatures of the SHA256SUMS and SHA512SUMS files are correct.
5. Only if the signatures are valid, determine whether the computed hashes match the hash values in the SHA256 and SHA512 files, respectively.
6. If the hashes match, then the downloaded firmware has been signed correctly and installation can proceed. If the hashes do not match, then it is critical that you do NOT USE THE DOWNLOADED IMAGE, because it is not trustworthy.

The following steps provide more detail for the process, using the software image file **efa-2.2.0.tar.gz** in the examples:

1. Import the signing key. This example assumes the key is in the **EXTR_signing_key.asc** file.

```
$ gpp --import EXTR_signing_key.asc
$ gpg --list-keys
/home/someuser/.gnupg/pubring.kbx
-----------------------------
pub   rsa4096 2020-04-24 [SC] [expires: 2022-04-24]
      4EA2F6FB375DFBE3C1D398BC33195B440230543A
uid           [ unknown] Extreme Networks, Inc. (Software Signing
Key) <security@extremenetworks.com>
sub   rsa4096 2020-04-24 [E] [expires: 2022-04-24]
```

   This sample output indicates the Extreme signing key is available.

2. Download the **efa-2.2.0.tar.gz** firmware image from the Extreme firmware downloads site.

```
$ ls -l efa-2.2.0.tar.gz
-rwx------+ 1 someuser 1694247452 May  8 16:37 efa-2.2.0.tar.gz
```

3. Download the following files SHA256SUMS, SHA256SUM.asc, SHA512SUMS, and SHA512SUMS.asc.

```
$ ls -altr SHA*
-rw-r--r--+ 1 someuser someuser 187 May  8 16:41 SHA256SUMS
-rw-r--r--+ 1 someuser someuser 801 May  8 16:42 SHA256SUMS.asc
-rw-r--r--+ 1 someuser someuser 315 May  8 16:43 SHA512SUMS
-rw-r--r--+ 1 someuser someuser 801 May  8 16:43 SHA512SUMS.asc
```

4. Verify that the PGP signatures of the SHA256SUMS and SHA512SUMS files are correct.

```
$ gpg --verify SHA256SUMS.asc SHA256SUMS
gpg: Signature made Fri 08 May 2020 01:42:08 PM PDT
gpg:                using RSA key 33195B440230543A
gpg: Good signature from "Extreme Networks, Inc. (Software Signing
Key) <security@extremenetworks.com>" [unknown]
gpg: WARNING: This key is not certified with a trusted signature!
gpg:          There is no indication that the signature belongs to
the owner.
Primary key fingerprint: 4EA2 F6FB 375D FBE3 C1D3  98BC 3319 5B44
0230 543A
```

```
$ gpg --verify SHA512SUMS.asc SHA512SUMS
gpg --verify SHA512SUMS.asc SHA512SUMS
gpg: Signature made Fri 08 May 2020 01:43:57 PM PDT
gpg:                using RSA key 33195B440230543A
gpg: Good signature from "Extreme Networks, Inc. (Software Signing
Key) <security@extremenetworks.com>" [unknown]
gpg: WARNING: This key is not certified with a trusted signature!
gpg:          There is no indication that the signature belongs to
the owner.
Primary key fingerprint: 4EA2 F6FB 375D FBE3 C1D3  98BC 3319 5B44
0230 543A
```

Note*:* The warning above about the key not being certified with a trusted signature can be removed if the signing verification key is also itself signed in the local GPG key ring. It does not affect whether the signature of the software images itself is valid. It rather has more to do with whether the user trusts the source of the verification key.

5. Because the PGP signatures are valid as shown in step 4, verify that the SHA256 and SHA512 hashes of the **efa-2.2.0.tar.gz** image file match the values in the SHA256SUMS and SHA512SUMS files, respectively.

```
$ shasum --check SHA256SUMS
efa-2.2.0.tar.gz: OK
$ shasum --check SHA512SUMS
efa-2.2.0.tar.gz: OK
```

At this point, the PGP signature validates the hashes, and the hashes have validated the raw image data. This forms a strong cryptographic verification that the data is correct, and the image is now trusted.

Backup and Restore the EFA Database

Follow these steps to back up and restore EFA data.

Notes:

- Backup from 2.1.0 and restore to 2.2.0 is not supported.
- Backup and restore works only between 2.2 builds and across environments.
- New commands are in 2.2.0.

1. To back up the database, run the following command as a root user with administrative privileges.
   (efa:extreme)extreme@tpvm:~$ efa backup

   Backup file :: /apps/efa_logs/backup//EFA-2020.06.10-07.06.08.tar

   --- Time Elapsed: 12.739234959s ---
   (efa:extreme)extreme@tpvm:~$

2. To restore the database, run the following command as a root user with administrative privileges.
   (efa:extreme)extreme@tpvm:~$ efa restore --backup-tar EFA-2020.06.10-07.06.08.tar
   Backup file name: EFA-2020.06.10-07.06.08.tar

   Restore result :: Restore operation successful

   --- Time Elapsed: 3m58.8150124s ---
   (efa:extreme)extreme@tpvm:~$

Note: The backup tar (for example, EFA-2020.06.10-07.06.08.tar) should be located under the **/var/log/efa/backup** directory for SERVER and **/apps/efa_logs/backup** directory for TPVM.

Known Limitations

- Downgrade from 2.2.0 to 2.1.0 is not supported.
- After the Fabric is configured, changes to Fabric settings are not allowed.
- After the Fabric is configured, wait for at least 3 minutes before configuring a tenant.
- In a vCenter ecosystem environment, EFA can manage an existing host or a DVS (Distributed Virtual Switch) that is present at the time of vCenter registration. Any newly added or deleted ESXi host or DVS is not managed by EFA.
- In a vCenter ecosystem environment, EFA can manage a DVS Port Group that is present at the time of vCenter registration. Any newly added Port Group is not managed by EFA for CCEP (Cluster Client End Point) port-channels.
- With a Hyper-V ecosystem integration and certain models of older Intel NIC on the compute node, the MAC address is incorrectly discovered and will not work.
- For the EFA 2.2.0 release, the EFA LDAP client is tested only with the OpenLDAP server. It is not tested with the Windows AD LDAP server. Also, when using the host IP for LDAP over SSL, the LDAP server certificate should be configured to have Subject Alternative names.

Open Defects

Issue ID:	EFA-2857
Summary:	Able to add fabric links and mct links in tenant
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.
Symptom:	Addition of fabric ports and MCT ports to the tenant succeeds even though these ports cannot be used to connect to any endpoints.
Workaround:	Fabric ports and MCT ports to be not added to the tenant.

Issue ID:	EFA-2900
Summary:	Tenant Service : When the cluster doesn't have "member vlan all" config and an EPG is created, "member vlan add <vlan-range>" gets configured under cluster, but the EPG deletion doesn't remove the config
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.
Symptom:	"member vlan add <vlan-range>" and "member bd add <bd-range>" configuration doesn't get removed from the "cluster" instance on the switch,
Workaround:	N/A

Issue ID:	EFA-3500
Summary:	Deleting 4000 l2 bd networks taking more than 2 hours 30 minis
Condition	Delete EPG having many (> 256) bridge-domain based networks.
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.
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 <>

Issue ID:	EFA-3512
Summary:	Provided VNI: 1006 already consumed in 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.
Symptom:	EPG create might fail with VNI resource not being available in the fabric
Workaround:	Provide a wait of 30 seconds between the create and delete CLI on the same EPG.

Issue ID:	EFA-3694
Summary:	REST: epg update with po add taking more time(>12min) with bd option with scale
Condition	1. Create EPG with 100 networks. 2. Update EPG with 20 ports/port-channels from one switch using "port-group-add" operation.
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.
Workaround:	N/A

Issue ID:	EFA-3717
Summary:	L3 bd epg went to cfg-refresh-err
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.
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.
Workaround:	Cleanup the out-of-band created LIFs from the switch before EPG create,

Issue ID:	EFA-3743
Summary:	Fabric is going cfg-refreshed without any reason showing.
Condition	
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

Issue ID:	EFA-3841
Summary:	EPG creation failed for PO with particular sequence
Condition	Create Tenants and register vcenter. EPG's are created properly. Delete vcenter. Wait for the EPG's getting removed. Create PO's Register vcenter EPG creation failed with error Tried updating inventory for the whole fabric and even then EPG creation failed.
Symptom:	EPG Creation Issue with deletion/addition of PO/vCenter
Workaround:	N/A

Issue ID:	EFA-3970
Summary:	epg creation with same bridge-domain name is failing
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
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.
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

Issue ID:	EFA-4016
Summary:	Devices configured from legacy efa, migrating to efa 2.1, adding extra bgp configs on leaf nodes and configure failing
Condition	
Symptom:	We are not supporting import from EFA legacy to 2.2
Workaround:	N/A

Issue ID:	EFA-4045
Summary:	"cfg refresh error" after moving MCT interface cables
Condition	
Symptom:	MCT configuration of two leaves on already Configured fabric was modified manually by replacing the existing MCT ports with 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

Issue ID:	EFA-4055
Summary:	VE creation failed after upgrade with existing VRF not having IPv6 AFI
Condition	1. Deploy EFA 2.1.0 version and create an L3 EPG resulting in VRF creation on the switch with IPv4 address-family but not Ipv6 address-family, because the ipv6 is not supported in EFA 2.1.0 version. 2. Upgrade EFA from 2.1.0 version to 2.2.0 version. 3. Try to create an EPG with the same VRF (used in step 1) by providing an ipv6-anycast-address.
Symptom:	EPG creation will fail as it will try to push anycast IPv6 on the VE but the IPv6 address-family is not present under the VRF on the switch.
Workaround:	N/A

Issue ID:	EFA-4056
Summary:	IPV6 config not getting pushed for bd-label without ipv6
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.
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.
Workaround:	Provide a unique anycast-ipv6 for a given BD name across the L3 EPG during EPG create and update operations.

Issue ID:	EFA-4109
Summary:	While creating ipv4 epg vrf, efa is pushing ipv6 config as well
Condition	Create an L3 EPG with only anycast-ipv4.
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
Workaround:	N/A

Issue ID:	EFA-4111
Summary:	When vrf with anycast-ip is present, updating of anycast-ipv6 is not provisioned though command is accepted and vice versa
Condition	Create L3 EPG with a VRF and only anycast-ipv4, then try to add the same VRF into the same EPG using EPG update "vrf-add" operation with an anycast-ipv6.
Symptom:	The EPG update "vrf-add" operation succeeds but anycast-ipv6 is not configured on the switch.
Workaround:	Remove the VRF from all the associated EPGs using "vrf-delete" operation and then re-add the VRF to the corresponding EPGs by providing both anycast-ip and anycast-ipv6 together.

Issue ID:	EFA-4832
Summary:	EFA is failed to communicate with the device if device credentials are removed
Condition	When a device is registered with device credentials in EFA and the same credentials are deleted or updated in SLX.
Symptom:	On any operation on EFA services, the error 'Invalid credentials for device 10.x.x.x.' is displayed.
Workaround:	Any update in device credentials on SLX(deleting the user, changing password etc) has to be updated in EFA using EFA CLI.

Issue ID:	EFA-4876
Summary:	DR: port-channel is not getting created in switch after rconcile CLI
Condition	1. Remove the EFA provisioned port-channel from the switch. 2. Reload the switch in maintenance mode.
Symptom:	Bridge domain logical-Interface(s) config under the EFA provisioned port-channel is not reconciled on the switch.
Workaround:	N/A

Issue ID:	EFA-5024
Summary:	When logged in TPVM with LDAP user, "efa login" command fails and the user is not allowed to execute any CLI
Condition	User sees the error as 'no roles are available for the user' on login to EFA.
Symptom:	When LDAP is configured for TPVM, user logs in as an LDAP user to the host, and tries to login to EFA with the same credentials.
Workaround:	Configure same LDAP details on EFA using 'efa ldapconfig command'

Issue ID:	EFA-5026
Summary:	efa login with remote user fails when LDAP over TLS is configured in EFA with server certificates are configured not to have Subject Alternate names
Condition	In TLS/SSL mode, and when the following error is seen. Could not authenticate against LDAP. LDAP Result Code 200 \"Network Error\": x509: cannot validate certificate for 10.x.x.x. because it doesn't contain any IP SANs
Symptom:	EFA is unable to authenticate users against configured LDAP in SSL mode when LDAP server certificates are configured not to have Subject Alternate names.
Workaround:	When using the host IP for SSL, the LDAP server certificate should be configured to have Subject Alternate names.

Issue ID:	EFA-5064
Summary:	Both suppress-arp and suppress-nd is enabled all the time
Condition	Configure L3 EPG with only anycast-ip and no anycast-ipv6.
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.
Workaround:	N/A

Issue ID:	EFA-5132
Summary:	Reconcile is success without reconciling ip/ipv4 address under "interface ve"
Condition	Drift created on the switch due to scenario as follows: 1. Remove IP address config(s) for Ve interface on switch 2. Add IP anycast-address config(s) for same Ve interface on switch 3. Reload the switch in maintenance mode.
Symptom:	IP address config for Ve Interface under EFA provisioned EPG is not reconciled on the switch
Workaround:	N/A

Issue ID:	EFA-5196
Summary:	Reconcile: "cluster-client auto" configuration under port-channel not reconciled after reboot
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.
Symptom:	"cluster-client auto" configuration will not be reconciled for the EFA provisioned port-channel on the switch.
Workaround:	N/A

Issue ID:	EFA-5203
Summary:	Device registration should fail if certificates are not pushed to device
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.
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.
Workaround:	There is no workaround - the device will be registered.

Issue ID:	EFA-5257
Summary:	vrf_route_target_mapping error while creating epg (after vrf delete)
Condition	Issue can be observed when vrf-add and vrf-delete operation is executed on Endpoint Group in quick succession multiple times.
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.
Workaround:	Workaround is to wait for a few minutes before executing the vrf-add again on Endpoint Group

Issue ID:	EFA-5286
Summary:	Rollback Error operation is not properly handled for rollback failure case for VRF Update
Condition	Issue is observed when VRF parameter update is requested, error occurs and rollback that is triggered also encounters error.
Symptom:	Responses to REST requests for VRF Update do not contain details of specific errors that occurred during rollback of errored configuration.
Workaround:	The final error encountered is visible as part of the final error returned to the REST request, in string form.

Issue ID:	EFA-5287
Summary:	RFE: Capture of child container error from device and reporting is missing for VRF components
Condition	Issue is observed when configuration under VRF has drifted on the switch due to various reasons.
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.
Workaround:	Workaround is to use display of 'efa tenant epg show' to determine if there is a drift in configuration.

Issue ID:	EFA-5290
Summary:	Drift-Reconcile - reload-delay & lacp timeout not getting reconciled in case of partial/conflict configuration.
Condition	When the switch drifts in configuration wrt intended configuration on the EFA, due to following scenario: 1. "reload-delay" provisioned (via EFA) on the CEP port is deleted from the switch. 2. "lacp timeout" provisioned (via EFA) is changed (from long to short) on the port-channel member interface from the switch. 3. Reload the switch in maintenance mode.
Symptom:	"reload-delay" config and "lacp timeout" configurations are not reconciled on the switch.
Workaround:	N/A

Issue ID:	EFA-5590
Summary:	"efa inventory rma execute --ip 10.20.54.65 --co" stuck for some time and triggered RMA
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
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
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.

Issue ID:	EFA-5592
Summary:	config-reply, RMA, drift and reconcile gets completed even though there is no certificate on switch/certificates are not copied on devices
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.
Symptom:	Certificates need to be manually imported on replaced equipment in-order to perform RMA.
Workaround:	import certificate manually efa certificates device install --ips x,y --certType

Issue ID:	EFA-5623
Summary:	Logs for Drift reconcile history is not getting forwarded to notification service
Condition	
Symptom:	Drift and reconcile logs are not notified to the Notification engine.
Workaround:	We can use the inventory service log file to get the Details of drift and reconcile

Issue ID:	EFA-5648
Summary:	EPG creation using same --bridge-domain <ctag:bd> for both CTAGs - creates 2 Different BDs vs expected 1 BD
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
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.
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

Issue ID:	EFA-5675
Summary:	vrf delete from epg and re-adding vrf to epg fails intermittently
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.
Symptom:	EPG update "vrf-add" operation can fail with the reason as vrf to be added has conflicting vrf on the switch.
Workaround:	Provide a wait of 30 seconds between the EPG update vrf-add and vrf-delete operations on the same EPG.

Issue ID:	EFA-5689
Summary:	Backup routing configs are pushed to Freedom switches and is not working
Condition	L3 EPG create or update operation with the member ports residing on SLXOS-9140.
Symptom:	VRF Backup routing configuration on the SLXOS-9140 will be inadequate and hence the backup routing functionality will not work on SLXOS-9140.
Workaround:	Backup routing needs to be disabled at the fabric setting level if the fabric has SLXOS-9140 devices.

Issue ID:	EFA-5697
Summary:	[REST DRIFT]: drift-reconcile history is showing success but port-channel member port is not coming up after device reload
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.
Symptom:	EFA provisioned port-channel member interface admin state will not be reconciled on the switch.
Workaround:	N/A

Issue ID:	EFA-5708
Summary:	[DRIFT]: After changing periodic discovery time and introducing a drift, efa fabric show app state is still displaying cfg in-sync after device rediscovery
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
Symptom:	Device status remains in cfg-in-sync when devices have mismatch in selective configs w.r.t to intended configs
Workaround:	Drift and Reconcile events fix the problem

Issue ID:	EFA-5732
Summary:	Fabric device remove is allowing for device which firmware download is in progress
Condition	If fabric delete command is submitted when firmware download is in progress, it fails.
Symptom:	When firmware download is in progress, fabric delete command is accepted without an error.
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}

Issue ID:	EFA-5745
Summary:	EFA tenant po create failed with error- it says the SLX switches are not MCT Pairs.
Condition	This issue is seen while creating PO using tenant service cli command
Symptom:	When creating fabric repetitively, intermittently when creating a non fabric, the MCT pair IP's are incorrectly configured.
Workaround:	This is an intermittent issue hence no workaround can be documented.

Issue ID:	EFA-5778
Summary:	efa firmware download across fabric is failing with Firmware Activate Failed error
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
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
Workaround:	This is an intermittent issue and there is no workaround to this.

Issue ID:	EFA-5794
Summary:	EFA inventory log file should mask/remove the SLX admin password given during inventory register
Condition	Log files for internal services like fabric, tenant etc inherit SLX details from inventory and may contain SLX login account password in plain text format.
Symptom:	EFA may write the password of SLX access account in plain text in some of the internal service log files.
Workaround:	N/A

Issue ID:	EFA-5821
Summary:	EPG creation failed without any error message when EPG is created with VLANS 2-4000 with BD-enable
Condition	Create EPG having more than 256 networks.
Symptom:	EPG create fails after a wait of more than 16 minutes.
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

Issue ID:	EFA-5822
Summary:	EPG delete taking 1hr to delete EPG with VLANS 2-4000 with BD-enable
Condition	Delete EPG having more than 256 networks.
Symptom:	EPG delete fails after a wait of more than 16 minutes.
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 <>

Issue ID:	EFA-5826
Summary:	efa showing-running is failing coz BR VNI conflicts
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.
Symptom:	EPG create CLIs can fail with an error "L3VNI is already consumed in the fabric"
Workaround:	Reorder "epg create" CLIs obtained from "efa show-config" in ascending order of l3-vni parameter.

Issue ID:	EFA-5830
Summary:	efa inventory device update fails with error "fabric doesn’t exist" after efa-firmware download
Condition	When device update command is submitted immediately after fabric is created, intermittently it returns error with message "fabric" does not exist
Symptom:	Unable to update devices using command : inventory device update --fabric {fabric name}
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}

Issue ID:	EFA-5832
Summary:	Drift-Reconcile : Removing channel-group from PO interface not reconciling
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.
Symptom:	"channel-group" config of EFA provisioned port-channel member interface will not be reconciled on the switch.
Workaround:	N/A

Issue ID:	EFA-5834
Summary:	EPG creation with 4K ctags can result in switch going un-responsive
Condition	EPG create operation with 4K ctags.
Symptom:	EPG creation with 4K ctags can result in switch going un-responsive
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

Issue ID:	EFA-5841
Summary:	L2/L3 tenant creation should not be allowed while efa firmware-download is in progress
Condition	If tenant commands are submitted when firmware download is in progress, it results in erroneous configuration and some configurations may miss.
Symptom:	When firmware download is in progress, tenant create command is accepted without an error.
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}

Issue ID:	EFA-5851
Summary:	After restore, command execution fails in same tab with error Token is not valid
Condition	When EFA backup taken from one system is restored on a different EFA.
Symptom:	After restore, all commands fail with the following error.(efa:extreme)extreme@tpvm:~$ efa fabric show Show [Failed] Token used in this session is not valid. Please logout and login with a valid user.
Workaround:	Either open a new session, or run 'source /etc/profile' and re-login to the application

Issue ID:	EFA-5858
Summary:	VRF Creation Failed: VRF name t1_vrf2 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.
Symptom:	VRFcreate fails with an error "VRF name <vrf-name> is already taken. Please use different VRF name."
Workaround:	N/A

Issue ID:	EFA-5869
Summary:	Notification service stops receiving logs from registered switches after efa restore
Condition	Occurs after efa restore
Symptom:	EFA raslog and notification service stops receiving logs from devices/switch which was restored from backed up data.
Workaround:	Restart the raslog service using efactl command

Issue ID:	EFA-5874
Summary:	Old EFA ip is present under syslog server configuration in switch after efa restore on different server
Condition	
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 device and remove the logging entry for the old EFA IP.

Issue ID:	EFA-5876
Summary:	Tenant Service Auto Reconcile failed when EFA created cluster is removed, new cluster is created followed by maintenance mode reload
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.
Symptom:	"cluster member vlan|bd add" and "cluster-client" config reconciliation fails on the switch.
Workaround:	N/A

Issue ID:	EFA-5897
Summary:	EFA does not bring all the breakout interface admin up while configuring breakout for multiple ports
Condition	Scale setup with 80 ports in breakout mode.
Symptom:	With scale, breakout ports configured may not be administratively up.
Workaround:	Ports will be brought administratively up when they are added to EPG.

Issue ID:	EFA-5902 5953
Summary:	System Management Commands like efactl, efa supportsave, efa backup need to be run with sudo access
Condition	User is logged in as admin and run the above-mentioned commands.
Symptom:	Throws permission denied error while running efactl, efa supportsave, efa backup commands from admin user.
Workaround:	Run commands with sudo - (efa)admin@efa:~$ sudo efa supportsave [sudo] password for admin: Version : 2.2.0 Build: 42 Time Stamp: 20-06-03:03:15:47 INFO[0002] Using host IP : -h10.20.62.213 INFO[0024] Support Save File: /var/log/efa/efa_2020-06-05T13:53:07.logs.zip Support Save File: /var/log/efa/efa_2020-06-05T13:53:07.logs.zip — Time Elapsed: 24.839415769s — or switch to root user using sudo su

Issue ID:	EFA-5927
Summary:	Scale Config: Drift-Reconcile failed with an error "drift and reconcile failed waiting for status from tenant "
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
Symptom:	Configuration reconciliation fails with an error "drift and reconcile failed waiting for status from tenant." because of the timeout.
Workaround:	N/A

Issue ID:	EFA-5928
Summary:	EFA displays "validate fabric failed with missing links" even though links are connected - while adding the nodes to the fabric
Condition	Added devices immediately after setting to default startup config
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
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>

Issue ID:	EFA-5936
Summary:	Cluster, overlay-gateway configurations not deleted from Switch if devices from fabric and inventory are deleted before deleting the tenant
Condition	Configure EPG/Networks on a fabric device and delete device from the fabric.
Symptom:	Overlay-Gateway/Cluster Instance is remaining on the switch.
Workaround:	N/A

Issue ID:	EFA-5941
Summary:	After creating 500 EPGs, some EPGs are in "cfg-refresh-err" state
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.
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 Endpoint Groups are created.
Workaround:	Workaround is to wait for about 30 minutes after all EFA config and show commands have executed.

Issue ID:	EFA-5945
Summary:	[Scale] BD based L3 EPG creation takes 20+ seconds for few EPG -2K EPG
Condition	Create EPGs (with few networks in each EPG) without any delay between EPG create CLIs.
Symptom:	Few of the EPG's creation takes more time (~20 seconds) compared to the other EPGs.
Workaround:	N/A

Issue ID:	EFA-5952
Summary:	[Auto-Persistence]"port-channel 1" showing "cfg-refreshed" in reconciliation status even if no drift on the devices
Condition	This happens when drift is executed while SLX is in maintenance mode.
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.
Workaround:	N/A

Issue ID:	EFA-5954
Summary:	efa firmware download prepare remove with fabric option not working
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
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 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}

Issue ID:	EFA-5958
Summary:	"efa tenant epg create" failed with netconf error '"0000" is an invalid value.',"
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.
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.
Workaround:	N/A

Issue ID:	EFA-5960
Summary:	[DR] Switch doesn't come out of maintenance mode
Condition	
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 .

Issue ID:	EFA-5961
Summary:	Tenant name is not completely visible in Audit log for OpenStack
Condition	Requests are sent from openstack ML2 plugin.
Symptom:	Tenant name and user name are shown partially in the audit log. 1018 AUDIT, 2020/06/04-03:24:20 (GMT), [SEC-3022], INFO, SECURITY, 47e2aab76d19484e8bdcb608d7a14fad-demo-R/admin/10.133.131.166/ssh/CLI,, SLX9540, Event: logout, Status: success, Info: Successful logout by user [47e2aab76d19484e8bdcb608d7a14fad-demo-R].
Workaround:	N/A

Issue ID:	EFA-5965
Summary:	efa firmware-download prepare add is aborting without 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.
Symptom:	Firmware download prepare add command fails without giving any error.
Workaround:	Intermittent issue, no workaround.

Issue ID:	EFA-5966
Summary:	Shared-port - Error: Operation "epg-event-pgupdate-delete-pg" not allowed for "ts-state-npepg" seen while deleting shared port
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
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
Workaround:	N/A

Issue ID:	EFA-5978
Summary:	EFA is not adding "mac-address-table mac-move detect" configuration after drift and reconcile
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
Symptom:	Mac move detect not getting reconciled on 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.

Issue ID:	EFA-5979
Summary:	After drift and reconcile - BFD config for backup routing for default-vrf is missing
Condition	Post default config running drift and reconcile
Symptom:	BFD config for router bgp does not get pushed after reconciling from default config
Workaround:	This config need to be pushed manually as this does not get reconciled

Issue ID:	EFA-5987
Summary:	IP registered for Syslog server on SLX on upgrade or fresh install could be incorrect.
Condition	Incorrect IP gets registered on the device and hence EFA doesn’t receive syslog messages from the switches and hence delayed updates.
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.
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.

Issue ID:	EFA-5993
Summary:	efa "system backup" and "system restore" fails after upgrading from 2.1 to 2.2
Condition	"efa system" cli commands are executed as follows- 1. Install EFA 2.1 2. Perform upgrade to EFA 2.2 3. Execute "efa system backup" and/or "efa system restore" commands
Symptom:	After upgrade from EFA 2.1 to EFA 2.2, "efa system backup" and "efa system restore" commands do not work
Workaround:	Restart "gosystem-service" pod to use "efa system backup" and/or "efa system restore" commands

Issue ID:	EFA-5998
Summary:	bgp ipv6 peer add failed after upgrade from 2.1 to 2.2
Condition	Create/Update operation on BGP Service with IPv6 neighbor failed since - ipv6 address-family vrf was not configured under router bgp instance in 2.1.
Symptom:	1. Deploy EFA 2.1.0 version and create an L3 EPG resulting in VRF creation on the switch with IPv4 address-family Ipv6 address-family not supported in EFA 2.1.0. 2. Upgrade EFA from 2.1.0 version to 2.2.0 version. 3. Configure ipv6 neighbor using BGP Service's create or update operation.
Workaround:	1. Delete and re-create all the L3 EPGs having the VRF association, which results in the VRF being deleted and re-created on the switch. OR 2. Perform EPG update "vrf-delete" and "vrf-add" operation on all the L3 EPGs having the VRF association, which results in the VRF being deleted and re-created on the switch.

Appendix 1: Upgrade SLX, TPVM, and EFA

Take the following steps to upgrade an SLX device running 20.1.1, with TPVM version 3.0.0 and EFA 2.1.x, to SLX 20.1.2a with TPVM version 4.00 and EFA 2.2.0.

1. Back up the EFA database to TPVM **/apps** using the **/apps/efa/efa_backup.sh** command. The backup and restore scripts are packaged with EFA 2.1.0.
    a. Create a backup directory for this step.
2. Tar the backup directory and copy the database backup to **/efaboot** of the SLX device, using SCP from TPVM.
3. Upgrade the SLX 9250 device running TPVM to SLX-OS 20.1.2a (normal firmware upgrade).
4. Stop the TPVM 3.0.0 application.
5. Uninstall TPVM 3.0.0.
6. Remove the TPVM3.0 image from **rm -rf /tftpboot/SWBD2900/*** using admin access.
7. Copy TPVM build **tpvm-4.0.0-5**. (The 20.1.2 TPVM GA).
8. Deploy TPVM with DHCP and static IP addresses.
    a. If you encounter any errors in this step, run the following: **do atpvm stop**, **tpvm uninstall**, and **tpvm deploy.**
9. Configure NTP on TPVM: **tpvm config ntp add server <ip_addr>**.
10. Copy the EFA 2.2.0 build.
11. Deploy the EFA 2.2.0 build, ensuring that you select secure or non-secure (standard) mode.

    Data migration from EFA 2.1.0 to 2.2.0 occurs from **/apps**.
    Note: As a best practice, install EFA 2.2.0 in secure mode.

Appendix 2: PGP Signature Verification Key

```
-----BEGIN PGP PUBLIC KEY BLOCK-----

mQINBF6jOAcBEAC0xe9dRipRMPEXpwVe2CLN5wk6QLbfjpK6FJUF1DonGL8+6TjI
UycIf15RcW8an9QDqIn7t254fmme+BvD7Nub+6rLa0NrQ2x6zWkk3VGKJBa/Q9Fz
+mKTfFFYZB/D55ipLgixQ6xxPNl8vsbbZS/KxdpL+29t3CY/OmGp3jpdWWAiEnBX
78Prt5BqmTxwOrTpZ0hcP1nKwmUqaEm/YOvgcENIPiRhMQoqAn7KY9tWKEdiCUYe
JkwITUDabcCq1HZ8bGmTsq+Z+BJTnYb/35QZea7MMUgBcfAb7ZCMv8IsR21sxC4w
wfGe/GNBVaGmh7rWr5fNf8TSeUhBIU/6wD0iK1B3F3zLnMJAEyhcS4RzEZhNuTcR
dJxE0oCMcNv5EqTGJoieZF6RPIK4rAg5sEevj/fNtkgNQtA9196+deO0CFfN07P7
/zxRmCaOr1ajEI5FxRIh4DIjRsKXV1YxzQlIhlT9ZU1GEQh9D1Mp6/odT//j6ShK
IcsDPS0BeyFLsRyjYMEjXGN081txFOMIiefXhvHw6jYCktL5oxKpyN6fLIusWD+s
lcCYHDu5Hn4s6Whh76ZGB8AX3BB0dU7hPeMtq25mK69K27gyfe3QPNBBFBR8Opz0
3VadQGd+dXT9TPbugLElL5tzD+cpo5AWnu1nTYsykyT1ZfNe8/v1xynJBwARAQAB
tExFeHRyZW1lIE5ldHdvcmtzLCBJbmMuIChTb2Z0d2FyZSBTaWduaW5nIEtleSkg
PHNlY3VyaXR5QGV4dHJlbWVuZXR3b3Jrcy5jb20+iQI+BBMBAgAoBQJeozgHAhsD
BQkDwmcABgsJCAcDAgYVCAIJCgsEFgIDAQIeAQIXgAAKCRAzGVtEAjBUOouqEACa
BHpbInqTZZ5+jPGX+TbClBwMwR5nFhrb8wjTFDQCW4IG+quvt40lTvukqkR5yhnL
SbETTpE/+zSrBrEepZkgfpRqE/V9/67xOiE+hybWwzlQwjdlgU1uYTVmqqMVscCr
skBWYCITPz/z7Wf/CBIfXvcxjheEYW2PAadnqCqZLD2dhteqdScLc6tq95uv0TnG
6Ti9QUV6PCRdGgwRbPnmgA0YfWLkzzfYLBIJfNmugSst6GHBSsU/MJGEgucfaG36
8Qblg5iNmwKFuBVHI3Gcn8ahlnloYRW7pmXbXFm+c8bJxTaT9EWmNvapyedpPtGa
txObjYZtWhkzEQR550+Pcw7qwvTuEmtbfTeHWe/jgsh/dI3I473r26QuvaPBj1Mq
TN6yd9a+gzbs1GqbrVsduTyE6w5B/4U1kV87ALJ56wqbbxhxpRc+1mC1aC/JqHwu
Fo7knopWrnTAMKkMJhKyjOSPdBVCuacaETuuOYb5GsTo6vZOIg5fXE5164eTjT1U
Rhw9hnphUp42QW58kX8WiN4PPP8GA5AVK5QNaG3xt8wY2Mp8hFPI2DC5NcnRGhLK
GAF+87JWXM29cjKYIA16/rsBHpvWfzRx1py5uc4ZDE8rdMqLv+oVu0V38VrJrWGW
1tf2fFUj4BvtOMlxMov08sDXvtA25FhttCu2veGd6rkCDQReozgHARAAyWXRmn9N
dE515/1j4HzWufQH/J9llhoYKouezTSyY3V84wIKIdWNh/4Yk4X6AXcdYyA6vkBj
6mzPZ4afR5rzrRx2hGi5CYMO8UOt77Q/V4nB2dFFd1SNHOAlXOx4Jequ0+XW/7Je
2gNP+Dfag//R7VmqH9Flgb7JqSzHtGfhAeK/+TtRrq2jJE79Nqjh7jrWm4eS4VwM
SgCZyqmoLq3O26G7sZn3d0y6vz8rAsWLUS3X9AiYwWAQpfOH5iwhtz+8bWgpJWFY
Tx6yYBd/1SMpbyF6RXn4D0pKHW3mG4dAVR+5vXzM1+DEZyHqC/8netVW03SqYr7t
D+IZKs34Zh8Jw//9mcC+Pbz1Frjyx7ERa1LkGqFzQGOsPnFNbfHYS2LH5wiGi+Rd
9q9BfUxlutiVvlMFlf/2ahxrvi8Wd4pFEgQ/6aPVWKmvoiUUE7Mf5nMvu9qfWSzi
oo2BTAk+Vf5Xa5sKE/MijGs/ozQTJveuUbcVyLQIPM7AfeHCCULIcNiayj2sIexk
oj2/+QE1+A9nPF6bUJs6i6SCmjorh9VBnE6ElxA6QYeeWENNoSj/mSIsWXizZG6Q
gXXW1ztAW5kZvZPtafaEgI4IVO4qWyIZaR6mXUwJFX+2J/ORwGem3MQgWnb27kQa
8LkhWCWOv0tOV/b96iB0UFMlbmzsHwOthK8AEQEAAYkCJQQYAQIADwUCXqM4BwIb
DAUJA8JnAAAKCRAzGVtEAjBUOtvAEACFf10qXq+/uYyuAZsT9Ie5B52GGtisjpzi
PXIf2lJNPeIX8CqfkFsFJqxE3NjVO8efI7eL2JqIqcrLFNKoSOX3DZU4PuTMiHPK
sHhKH8cIFucQrdQUMW6uPCe05wlQYxF9UlXrGyMU4KtPY9f0w2B2R5U2TeM3v0a1
z15ZllMFNhxKn1fC/XHjJmUmAEE+v3zWfELRvV92t+Hl6bTZKdxSCeaNFwlzb6FJ
pl43Sp91JKvY0wjmqku26LvIGQoCmcLJeliJDq1VclngL7YjVXIjNEtGSvx+PYAu
QjVsUTH4c2BbbIvDFeBdkXUNbv1UBIHyEniOz6N5ieGrsnai4M7ftQk2xMsMX1hY
GB5dqOfqB88bhrzd/WuilurloXGTU/pqJwaBddpCdRutCQM+7BuFdpoGLqCorGAb
ZAw2w538X9PGbavmtVf4lmdX4FNi+2Ws8nEIWI+DW+sJm6NbGRoRNdOUt/96rgoD
sFsiPU1Dtz2AIQToqIZQFi8O8YwExl0PaQ9FwfhxEZb0oLUWB5Ow0yIyOZI/lxOp
```

```
6jEVCfV1bzZJMEviORTAo76fFM2GDcNrNf+p5VFaSouwfGAIJFNYz0ucZGOvFPA3
XHWBNLwBg96zeVfkxSG4cChekJiTL1XDk4bCMfdRAzN1hyLkRId/lHXwTmvwzlhQ
IHBwMESrRg==
=Eh/m
-----END PGP PUBLIC KEY BLOCK-----
```