

Extreme Fabric Automation Command Reference

2.7.1

9037425-01 Rev AA
June 2022



Copyright © 2022 Extreme Networks, Inc. All rights reserved.

Legal Notice

, 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 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

and the logo are trademarks or registered trademarks of , 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 trademarks, see:

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.....	8
Text Conventions.....	8
Documentation and Training.....	9
Help and Support.....	10
Subscribe to Product Announcements.....	10
Send Feedback.....	10
About this Document.....	12
What's New in this Document.....	12
Using the EFA Application CLI.....	13
User Accounts.....	13
Default Account Credentials.....	13
Log in to EFA.....	14
Completing CLI Commands.....	14
CLI Keyboard Shortcuts.....	14
Unsupported Input Characters.....	15
Extreme Fabric Automation Commands.....	16
efa auth apikey.....	22
efa auth client.....	23
efa auth execution.....	24
efa auth ldapconfig.....	25
efa auth rolemapping.....	29
efa auth settings token.....	30
efa certificate device install.....	31
efa certificate server.....	32
efa certificate server renew.....	33
efa fabric clone.....	34
efa fabric configure.....	35
efa fabric create.....	37
efa fabric debug clear-config.....	38
efa fabric debug config-gen-reason.....	39
efa fabric debug device drift.....	40
efa fabric debug service lock.....	41
efa fabric debug service smartdrc.....	42
efa fabric delete.....	43
efa fabric device add.....	44
efa fabric device add-bulk.....	46
efa fabric device remove.....	48
efa fabric error show.....	49
efa fabric execution.....	50
efa fabric setting.....	51

efa fabric show.....	56
efa fabric show-config.....	57
efa fabric show summary	58
efa fabric topology show overlay.....	59
efa fabric topology show underlay.....	60
efa fabric topology show physical.....	61
efa inventory admin-state.....	62
efa inventory config-backup.....	64
efa inventory config-replay.....	66
efa inventory debug devices-lock.....	68
efa inventory debug devices-unlock.....	69
efa inventory debug unblock-from-fwdl.....	70
efa inventory device clear route-all.....	71
efa inventory device compare.....	72
efa inventory device delete.....	73
efa inventory device discovery-time list.....	74
efa inventory device discovery-time update.....	75
efa inventory device execute-cli.....	76
efa inventory device firmware-download commit.....	77
efa inventory device firmware-download execute.....	78
efa inventory device firmware-download prepare add.....	82
efa inventory device firmware-download prepare list.....	87
efa inventory device firmware-download prepare remove.....	88
efa inventory device firmware-download restore.....	91
efa inventory device firmware-download show.....	92
efa inventory device health status.....	94
efa inventory device interface list.....	95
efa inventory device interface list-breakout.....	96
efa inventory device interface redundant-management.....	97
efa inventory device interface set-admin-state.....	98
efa inventory device interface set-breakout.....	101
efa inventory device interface set-fec.....	104
efa inventory device interface set-link-error-disable.....	105
efa inventory device interface set-mtu.....	106
efa inventory device interface set-speed.....	109
efa inventory device interface unset-breakout.....	112
efa inventory device interface unset-fec.....	114
efa inventory device interface unset-link-error-disable.....	115
efa inventory device list.....	116
efa inventory device lldp list.....	117
efa inventory device ntp disable-server.....	118
efa inventory device ntp server create.....	119
efa inventory device ntp server delete.....	120
efa inventory device ntp server list.....	121
efa inventory device running-config persist.....	122
efa inventory device register.....	123
efa inventory device reload.....	124
efa inventory device setting show.....	125
efa inventory device setting update.....	126

efa inventory device snmp community create.....	128
efa inventory device snmp community delete.....	129
efa inventory device snmp community list.....	130
efa inventory device snmp host create.....	131
efa inventory device snmp host delete.....	132
efa inventory device snmp host list.....	133
efa inventory device snmp user create.....	135
efa inventory device snmp user delete.....	137
efa inventory device snmp user list.....	138
efa inventory device timezone debug-show.....	139
efa inventory device timezone list.....	140
efa inventory device timezone set.....	141
efa inventory device timezone unset.....	142
efa inventory device tpvm-upgrade execute.....	143
efa inventory device tpvm-upgrade show.....	144
efa inventory device update.....	145
efa inventory drift-reconcile.....	146
efa inventory execution.....	148
efa inventory firmware-host delete.....	149
efa inventory firmware-host list.....	150
efa inventory firmware-host register.....	151
efa inventory firmware-host update.....	152
efa inventory kvstore.....	153
efa inventory rma.....	154
efa login.....	155
efa logout.....	156
efa mgmt route create.....	157
efa mgmt route delete.....	158
efa mgmt route show.....	159
efa mgmt subinterface create.....	160
efa mgmt subinterface delete.....	161
efa mgmt subinterface show.....	162
efa mgmt subinterface staticip add.....	163
efa mgmt subinterface staticip remove.....	164
efa mgmt subinterface staticip show.....	165
efa notification subscribers add-https.....	166
efa notification subscribers add-syslog-relp.....	167
efa notification subscribers delete.....	169
efa notification subscribers get.....	170
efa notification subscribers list.....	171
efa openstack debug.....	172
efa openstack execution.....	174
efa openstack network show.....	175
efa openstack network-interface show.....	176
efa openstack router show.....	177
efa openstack router-interface show.....	178
efa openstack router-route show.....	179
efa openstack subnet show.....	180
efa openstack sync start.....	181

efa policy prefix-list create.....	182
efa policy prefix-list update.....	183
efa policy prefix-list delete.....	185
efa policy prefix-list list.....	186
efa policy route-map create.....	187
efa policy route-map update.....	188
efa policy route-map delete.....	190
efa policy route-map list.....	191
efa policy route-map-match create.....	192
efa policy route-map-match delete.....	193
efa rbac execution.....	194
efa rbac role show.....	195
efa scvmm delete.....	197
efa scvmm links physical.....	198
efa scvmm links virtual.....	199
efa scvmm list.....	200
efa scvmm register.....	201
efa scvmm settings show.....	202
efa scvmm settings update.....	203
efa scvmm update.....	204
efa snmp subscriber.....	205
efa show-running-config.....	207
efa status.....	208
efa system backup.....	209
efa system backup-list.....	211
efa system cleanup.....	212
efa system feature show.....	213
efa system feature update.....	214
efa system restore.....	215
efa system settings reset.....	216
efa system settings show.....	217
efa system settings update.....	218
efa system supportsave.....	220
efa tenant create.....	221
efa tenant debug.....	223
efa tenant delete.....	225
efa tenant epg configure.....	226
efa tenant epg create.....	227
efa tenant epg delete.....	232
efa tenant epg detach.....	233
efa tenant epg error show.....	236
efa tenant epg show.....	237
efa tenant epg update.....	239
efa tenant execution.....	243
efa tenant po configure.....	244
efa tenant po create.....	245
efa tenant po delete.....	247
efa tenant po show.....	248
efa tenant po update.....	251

efa tenant service bgp peer configure.....	253
efa tenant service bgp peer create.....	254
efa tenant service bgp peer delete.....	259
efa tenant service bgp peer operational show.....	260
efa tenant service bgp peer show.....	266
efa tenant service bgp peer update.....	270
efa tenant service bgp peer-group configure.....	273
efa tenant service bgp peer-group create.....	274
efa tenant service bgp peer-group delete.....	276
efa tenant service bgp peer-group show.....	277
efa tenant service bgp peer-group update.....	279
efa tenant service mirror session create.....	281
efa tenant service mirror session delete.....	283
efa tenant service mirror session configure.....	284
efa tenant service mirror session show.....	285
efa tenant show.....	286
efa tenant update.....	287
efa tenant vrf create.....	289
efa tenant vrf delete.....	295
efa tenant vrf error show.....	296
efa tenant vrf show.....	297
efa tenant vrf update.....	300
efa vcenter debug.....	305
efa vcenter delete.....	309
efa vcenter links.....	310
efa vcenter list.....	312
efa vcenter register	313
efa vcenter update.....	314
efa version.....	315
efa-change-hostname.....	316
efa-change-ip.....	317
efa-change-vip.....	318
efactl.....	319
Openstack Controller EFA Commands.....	326
efa-blpair-mapping.....	327
efa-health show.....	328
efa-journal clear.....	330
efa-journal list.....	332
efa-journal list deps.....	333
efa-journal reset.....	334
efa-sync execute.....	335
openstack network efa-bl-pair-map create.....	337
openstack network efa-topology-link-map create.....	338
openstack network efa-topology-link-map delete.....	339
openstack network efa-topology-link-map list.....	340
openstack router create.....	341



Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

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

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

Table 1: Notes and warnings






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

Table 2: Text

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

Table 3: Command syntax

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

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

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

[Extreme Optics Compatibility](#)

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

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

Help and Support

If you require assistance, contact 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 customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by 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.

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

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

Subscribe to Product Announcements

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

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

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

Send Feedback

The Information Development team at has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

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

To send feedback, do either of the following:

- Access the feedback form at .
- Email us at .

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.



About this Document

[What's New in this Document](#) on page 12

What's New in this Document

This version of the Command Reference contains editorial changes to syntax for some commands.



Using the EFA Application CLI

- [User Accounts on page 13](#)
- [Default Account Credentials on page 13](#)
- [Log in to EFA on page 14](#)
- [Completing CLI Commands on page 14](#)
- [CLI Keyboard Shortcuts on page 14](#)
- [Unsupported Input Characters on page 15](#)

The command line provides a powerful means for configuring, managing, and monitoring packet traffic through the EFA device.

The following topics describe accessing and using the EFA command-line interface (CLI), including syntax, command completion, shortcuts, and other helpful subjects.

User Accounts

A user account specifies a user's level of access to the device CLI.

EFA uses role-based access control (RBAC) as the authorization mechanism. A *role* is a container for rules that specify which commands can be run and with which permissions. When you create a user account, you specify a role for that account. In general, *user* (as opposed to *user-level*) refers to any account to which an admin or user role can be assigned.

For more information about user accounts and roles, see the [Extreme Fabric Automation Security Guide, 2.7.0](#).

Default Account Credentials

As a best practice, log on as the administrator and change the default passwords immediately after the EFA is installed.



Note

The third-party virtual machine (TPVM) on SLX devices has one user, with the following case-sensitive credentials:

- user account: `extreme`
- password: `password`

Log in to EFA

Use of the EFA command line requires a valid, logged-in user.

Procedure

1. Verify the status of the EFA deployment using one of the following methods.
 - Run the SLX **show efa status** command.
 - Run the EFA **efactl status** script (or the **efa status** command, as an alternative).
2. Log in to EFA.

```
$ efa login --username <username>
Password: <password>
```

The <username> variable is optional. If you do not provide a user name, log-in defaults to the current (Unix) user.

With a successful log-in, the command prompt shows the logged-in user in green text. If the log-in is not successful, the command prompt is displayed in red text.

3. To log out of EFA, run the **efa logout** command.

Completing CLI Commands

To complete the spelling of commands or keywords automatically, begin typing the command or keyword and then press **Tab**. For example, at the CLI command prompt, type `efa a` and press **Tab**:

```
efa a
```

The CLI displays the following command.

```
efa auth
```

If there is more than one command or keyword associated with the characters typed, the CLI displays all choices. For example, at the CLI command prompt, type `efa auth` and press **Tab**.

```
apikey
client
execution
ldapconfig
rolemapping
settings
```

CLI Keyboard Shortcuts

The following table lists CLI keyboard shortcuts.

Keystroke	Description
Ctrl+A	Moves the cursor to the beginning of the command line.
Ctrl+B (or the left arrow key)	Moves the cursor back one character.
Ctrl+C	Escapes and terminates command prompts and ongoing tasks (such as lengthy displays), and displays a fresh command prompt.
Ctrl+E	Moves the cursor to the end of the command line.

Keystroke	Description
Ctrl+F (or the right arrow key)	Moves the cursor forward one character.
Ctrl+N (or the down arrow key)	Displays commands in the history buffer with the most recent command displayed last.
Ctrl+P (or the up arrow key)	Displays commands in the history buffer with the most recent command displayed first.
Ctrl+U	Deletes all characters from the cursor to the beginning of the command line.
Ctrl+W	Deletes the last word you typed.
Ctrl+Z	Returns to privileged EXEC mode. Using Ctrl+Z in privileged EXEC mode runs partial commands.
Esc B	Moves the cursor back one word.
Esc F	Moves the cursor forward one word.

Unsupported Input Characters

If unsupported input characters are used for user-defined objects, an error message is displayed.

Characters dependent on combinations of the **AltGr** key and another key are not supported.



Note

The **AltGr** key is the **Alt** key to the right of the space bar.



Extreme Fabric Automation Commands

[efa auth apikey](#) on page 22
[efa auth client](#) on page 23
[efa auth execution](#) on page 24
[efa auth ldapconfig](#) on page 25
[efa auth rolemapping](#) on page 29
[efa auth settings token](#) on page 30
[efa certificate device install](#) on page 31
[efa certificate server](#) on page 32
[efa certificate server renew](#) on page 33
[efa fabric clone](#) on page 34
[efa fabric configure](#) on page 35
[efa fabric create](#) on page 37
[efa fabric debug clear-config](#) on page 38
[efa fabric debug config-gen-reason](#) on page 39
[efa fabric debug device drift](#) on page 40
[efa fabric debug service lock](#) on page 41
[efa fabric debug service smartdrc](#) on page 42
[efa fabric delete](#) on page 43
[efa fabric device add](#) on page 44
[efa fabric device add-bulk](#) on page 46
[efa fabric device remove](#) on page 48
[efa fabric error show](#) on page 49
[efa fabric execution](#) on page 50
[efa fabric setting](#) on page 51
[efa fabric show](#) on page 56
[efa fabric show-config](#) on page 57
[efa fabric show summary](#) on page 58
[efa fabric topology show overlay](#) on page 59
[efa fabric topology show underlay](#) on page 60
[efa fabric topology show physical](#) on page 61
[efa inventory admin-state](#) on page 62
[efa inventory config-backup](#) on page 64
[efa inventory config-replay](#) on page 66
[efa inventory debug devices-lock](#) on page 68

[efa inventory debug devices-unlock](#) on page 69

[efa inventory debug unblock-from-fwdl](#) on page 70

[efa inventory device clear route-all](#) on page 71

[efa inventory device compare](#) on page 72

[efa inventory device delete](#) on page 73

[efa inventory device discovery-time list](#) on page 74

[efa inventory device discovery-time update](#) on page 75

[efa inventory device execute-cli](#) on page 76

[efa inventory device firmware-download commit](#) on page 77

[efa inventory device firmware-download execute](#) on page 78

[efa inventory device firmware-download prepare add](#) on page 82

[efa inventory device firmware-download prepare list](#) on page 87

[efa inventory device firmware-download prepare remove](#) on page 88

[efa inventory device firmware-download restore](#) on page 91

[efa inventory device firmware-download show](#) on page 92

[efa inventory device health status](#) on page 94

[efa inventory device interface list](#) on page 95

[efa inventory device interface list-breakout](#) on page 96

[efa inventory device interface redundant-management](#) on page 97

[efa inventory device interface set-admin-state](#) on page 98

[efa inventory device interface set-breakout](#) on page 101

[efa inventory device interface set-fec](#) on page 104

[efa inventory device interface set-link-error-disable](#) on page 105

[efa inventory device interface set-mtu](#) on page 106

[efa inventory device interface set-speed](#) on page 109

[efa inventory device interface unset-breakout](#) on page 112

[efa inventory device interface unset-fec](#) on page 114

[efa inventory device interface unset-link-error-disable](#) on page 115

[efa inventory device list](#) on page 116

[efa inventory device lldp list](#) on page 117

[efa inventory device ntp disable-server](#) on page 118

[efa inventory device ntp server create](#) on page 119

[efa inventory device ntp server delete](#) on page 120

[efa inventory device ntp server list](#) on page 121

[efa inventory device running-config persist](#) on page 122

[efa inventory device register](#) on page 123

[efa inventory device reload](#) on page 124

[efa inventory device setting show](#) on page 125

[efa inventory device setting update](#) on page 126

[efa inventory device snmp community create](#) on page 128

[efa inventory device snmp community delete](#) on page 129

[efa inventory device snmp community list](#) on page 130

[efa inventory device snmp host create](#) on page 131
[efa inventory device snmp host delete](#) on page 132
[efa inventory device snmp host list](#) on page 133
[efa inventory device snmp user create](#) on page 135
[efa inventory device snmp user delete](#) on page 137
[efa inventory device snmp user list](#) on page 138
[efa inventory device timezone debug-show](#) on page 139
[efa inventory device timezone list](#) on page 140
[efa inventory device timezone set](#) on page 141
[efa inventory device timezone unset](#) on page 142
[efa inventory device tpvm-upgrade execute](#) on page 143
[efa inventory device tpvm-upgrade show](#) on page 144
[efa inventory device update](#) on page 145
[efa inventory drift-reconcile](#) on page 146
[efa inventory execution](#) on page 148
[efa inventory firmware-host delete](#) on page 149
[efa inventory firmware-host list](#) on page 150
[efa inventory firmware-host register](#) on page 151
[efa inventory firmware-host update](#) on page 152
[efa inventory kvstore](#) on page 153
[efa inventory rma](#) on page 154
[efa login](#) on page 155
[efa logout](#) on page 156
[efa mgmt route create](#) on page 157
[efa mgmt route delete](#) on page 158
[efa mgmt route show](#) on page 159
[efa mgmt subinterface create](#) on page 160
[efa mgmt subinterface delete](#) on page 161
[efa mgmt subinterface show](#) on page 162
[efa mgmt subinterface staticip add](#) on page 163
[efa mgmt subinterface staticip remove](#) on page 164
[efa mgmt subinterface staticip show](#) on page 165
[efa notification subscribers add-https](#) on page 166
[efa notification subscribers add-syslog-relp](#) on page 167
[efa notification subscribers delete](#) on page 169
[efa notification subscribers get](#) on page 170
[efa notification subscribers list](#) on page 171
[efa openstack debug](#) on page 172
[efa openstack execution](#) on page 174
[efa openstack network show](#) on page 175
[efa openstack network-interface show](#) on page 176
[efa openstack router show](#) on page 177

[efa openstack router-interface show](#) on page 178
[efa openstack router-route show](#) on page 179
[efa openstack subnet show](#) on page 180
[efa openstack sync start](#) on page 181
[efa policy prefix-list create](#) on page 182
[efa policy prefix-list update](#) on page 183
[efa policy prefix-list delete](#) on page 185
[efa policy prefix-list list](#) on page 186
[efa policy route-map create](#) on page 187
[efa policy route-map update](#) on page 188
[efa policy route-map delete](#) on page 190
[efa policy route-map list](#) on page 191
[efa policy route-map-match create](#) on page 192
[efa policy route-map-match delete](#) on page 193
[efa rbac execution](#) on page 194
[efa rbac role show](#) on page 195
[efa scvmm delete](#) on page 197
[efa scvmm links physical](#) on page 198
[efa scvmm links virtual](#) on page 199
[efa scvmm list](#) on page 200
[efa scvmm register](#) on page 201
[efa scvmm settings show](#) on page 202
[efa scvmm settings update](#) on page 203
[efa scvmm update](#) on page 204
[efa snmp subscriber](#) on page 205
[efa show-running-config](#) on page 207
[efa status](#) on page 208
[efa system backup](#) on page 209
[efa system backup-list](#) on page 211
[efa system cleanup](#) on page 212
[efa system feature show](#) on page 213
[efa system feature update](#) on page 214
[efa system restore](#) on page 215
[efa system settings reset](#) on page 216
[efa system settings show](#) on page 217
[efa system settings update](#) on page 218
[efa system supportsave](#) on page 220
[efa tenant create](#) on page 221
[efa tenant debug](#) on page 223
[efa tenant delete](#) on page 225
[efa tenant epg configure](#) on page 226
[efa tenant epg create](#) on page 227

[efa tenant epg delete](#) on page 232
[efa tenant epg detach](#) on page 233
[efa tenant epg error show](#) on page 236
[efa tenant epg show](#) on page 237
[efa tenant epg update](#) on page 239
[efa tenant execution](#) on page 243
[efa tenant po configure](#) on page 244
[efa tenant po create](#) on page 245
[efa tenant po delete](#) on page 247
[efa tenant po show](#) on page 248
[efa tenant po update](#) on page 251
[efa tenant service bgp peer configure](#) on page 253
[efa tenant service bgp peer create](#) on page 254
[efa tenant service bgp peer delete](#) on page 259
[efa tenant service bgp peer operational show](#) on page 260
[efa tenant service bgp peer show](#) on page 266
[efa tenant service bgp peer update](#) on page 270
[efa tenant service bgp peer-group configure](#) on page 273
[efa tenant service bgp peer-group create](#) on page 274
[efa tenant service bgp peer-group delete](#) on page 276
[efa tenant service bgp peer-group show](#) on page 277
[efa tenant service bgp peer-group update](#) on page 279
[efa tenant service mirror session create](#) on page 281
[efa tenant service mirror session delete](#) on page 283
[efa tenant service mirror session configure](#) on page 284
[efa tenant service mirror session show](#) on page 285
[efa tenant show](#) on page 286
[efa tenant update](#) on page 287
[efa tenant vrf create](#) on page 289
[efa tenant vrf delete](#) on page 295
[efa tenant vrf error show](#) on page 296
[efa tenant vrf show](#) on page 297
[efa tenant vrf update](#) on page 300
[efa vcenter debug](#) on page 305
[efa vcenter delete](#) on page 309
[efa vcenter links](#) on page 310
[efa vcenter list](#) on page 312
[efa vcenter register](#) on page 313
[efa vcenter update](#) on page 314
[efa version](#) on page 315
[efa-change-hostname](#) on page 316
[efa-change-ip](#) on page 317

[efa-change-vip](#) on page 318

[efactl](#) on page 319

This document describes commands available in the Extreme Fabric Automation CLI.

efa auth apikey

Creates or displays the API Key for the OpenStack instance.

Syntax

```
efa auth apikey generate [--client-id | --force ]  
efa auth apikey show [--client-id ]
```

Parameters

--client-id

Generates the API Key for the indicated client ID.

--force

Forces key regenerate.

Usage Guidelines

Use this command to generate the API Key after you used the **efa auth client register** command to register the client.

Save the generated key to use for configuring the ML2 plug in.

Examples

This example generates an API Key for the indicated client ID.

```
# efa auth apikey generate --client-id d6d7430e-7cd0-11ea-b7a6-aaa8d3cb654e  
API Key is created successfully.  
eyJhbGciOiJSUzI1NiIsImtZCI6IjEuMCIsInR5cCI6IkpX...  
--- Time Elapsed: 506.458623ms ---
```

efa auth client

Deletes or registers OpenStack tenant clients and displays client details.

Syntax

```
efa auth client delete [ --name tenant-name ]  
efa auth client register [ --name tenant-name | --type {cli |  
  openstack} ]  
efa auth client show [ --name tenant-name ]
```

Parameters

--name *tenant-name*
Specifies the name of the tenant to register.

--type *client-type*
Specifies the type of client you are registering.

Usage Guidelines

Use the generated key when you configure the ML2 plugin.

Examples

This example registers a tenant called RegionOne.

```
# efa auth client register --name RegionOne -- type OpenStack  
Successfully registered client.  
Attribute      Value  
name           RegionOne  
type           OPENSTACK  
ID             d6d7340e-7cd0-11ea-b7a6-aaa8d3cb654e
```

efa auth execution

Deletes or displays the authorization execution logs.

Syntax

```
efa auth execution delete --days int32
```

Deletes execution entries that are older than the specified number of days (default 30).

```
efa auth execution show [ --id string | --limit int32 | --status string ]
```

Shows a list of executions.

Parameters

--id string

Filters the executions based on execution id. The `limit` and `status` flags are ignored when the `id` flag is specified.

--limit int32

Limits the number of executions to be listed. The value 0 lists all the executions. Default is 10.

--status string

Filters the executions based on the status (failed/succeeded/all). Default is all.

Examples

This example deletes entries older than 15 days.

```
efa auth execution delete 15
```


efa auth ldapconfig

Adds, updates, or deletes an external LDAP server, or shows the current LDAP configuration.

Syntax

```
efa auth ldapconfig add [ --name ldap-name | --primary value | --host
  hostname | --port port-num | --tls | --insecure-tls --cacert cert-loc
  | --timeout value | --bind-user-name dn | --bind-user-password pwd
  | --user-search-base dn | --user-object-class obj-class | --user-
login-attribute att-value | --user-role-attribute att-value | --user-
role-attribute-key att-value | --user-member-attribute att-value | --
group-search-base dn | --group-object-class obj-class | --group-
attribute att-value | --group-member-user-attribute att-value | --
group-member-mapping-attribute att-value ]
```

```
efa auth ldapconfig update [ --name ldap-name | --primary value | --host
  hostname | --port port-num | --tls | --insecure-tls --cacert cert-loc
  | --timeout value | --bind-user-name dn | --bind-user-password pwd
  | --user-search-base dn | --user-object-class obj-class | --user-
login-attribute att-value | --user-role-attribute att-value | --user-
role-attribute-key att-value | --user-member-attribute att-value | --
group-search-base dn | --group-object-class obj-class | --group-
attribute att-value | --group-member-user-attribute att-value | --
group-member-mapping-attribute att-value]
```

```
efa auth ldapconfig delete [--name ldap-name | --primary value | --host
  hostname | --port port-num | --tls | --insecure-tls --cacert cert-loc
  | --timeout value | --bind-user-name dn | --bind-user-password pwd
  | --user-search-base dn | --user-object-class obj-class | --user-
login-attribute att-value | --user-role-attribute att-value | --user-
role-attribute-key att-value | --user-member-attribute att-value | --
group-search-base dn | --group-object-class obj-class | --group-
attribute att-value | --group-member-user-attribute att-value | --
group-member-mapping-attribute att-value ]
```

```
efa auth ldapconfig show [--name ldap-name | --primary value | --host
  hostname ]
```

Parameters

--name *ldap-name*

Specifies the name of the LDAP connection.

--primary *value*

Specifies 1 when multiple LDAP connections are available.

--host *hostname*

Specifies the host name or IP address of the host server.

--port *port-num* **tls** | **insecure-tls**

Specifies the port at which the LDAP server listens for connections.

Specify **--tls** to use LDAP over SSL and TLS. Specify **--insecure-tls** to use LDAP without certification verification.

--cacert *cert-loc*

Specifies the location of the Certificate Authority certificate.

--timeout *value*

Specifies the number of seconds that must elapse before the LDAP server is considered unreachable. The default is 5 seconds.

--bind-user-name *dn*

Specifies the Distinguished Name (DN) of the user that you want to use to bind, search, and retrieve LDAP entries.

--bind-user-password *pwd*

Specifies the password of the bind user.

--user-search-base *dn*

Specifies the DN of the node in the directory tree from which searches for user objects will start.

--user-object-class *obj-class*

Specifies the name of the object class to use for user objects. The default is `inetOrgPerson`.

--user-login-attribute *att-value*

Specifies the attribute that matches the user name part of credentials that users enter while logging in. The default is `uid`.

--user-role-attribute *att-value*

Specifies the attribute from which the user role is read.

--user-role-attribute-key *att-value*

Specifies the attribute that reads the role value from the role attribute.

--user-member-attribute *att-value*

Specifies the attribute that reads the member of the group that the user is part of.

--group-search-base *dn*

Specifies the DN of the node in the directory tree from which searches for group objects begins.

--group-object-class *obj-class*

Specifies the name of the object class to use for group searches. The default is `groupOfNames`.

--group-attribute *att-value*

Specifies the attribute that defines the search filter on a group. The default is `cn`.

--group-member-user-attribute *att-value*

Specifies the name of the user attribute whose format matches the group members. The default is `entrydn`.

--group-member-mapping-attribute *att-value*

Specifies the name of the group attribute that contains the members of a group. The default is `member`.

Usage Guidelines

You configure an LDAP server for user validation and to fetch user groups.

When a user is assigned EFA roles in LDAP, ensure that you define the **user-role-attribute** parameter.

You can use key-value pairs to define one attribute value that assigns multiple roles to a user. Use the **user-role-attribute-key** parameter for such a scenario.

When you use LDAP groups to assign roles to users, ensure that you define the **user-member-attribute** parameter.

When LDAP groups are not in the same search base as the users in the groups, ensure that you define the following parameters.

- **group-search-base**
- **group-object-class**
- **group-attribute**
- **group-member-user-attribute**
- **group-member-mapping-attribute**

To configure LDAP for a deployment of EFA on a TPVM, see the "TPVM Management" section of the *Extreme SLX-OS Management Configuration Guide*.

Examples

This example configures the bind user name, the bind password, and the DN of the node from which searches start.

```
# efa auth ldapconfig add --name ldapconfig -- host 10.x.x.x
--bind-user-name cn=admin,dc=extrnet,dc=com --bind-user-password password
--user-search-base ou=people,dc=extrnet,dc=com
```

This example configures the **--user-role-attribute** parameter for a user that is assigned EFA roles in LDAP.

```
# efa auth ldapconfig add --name ldap1 --host 10.x.x.x
--bind-user-name cn=x,dc=y,dc=com --bind-user-password xxx
--user-search-base ou=people,dc=y,dc=com --user-role-attribute role
```

This example assigns multiple roles to a user with one key-value pair. The role attribute for the user entry in LDAP has the value of `datacenterowner:SystemAdmin,datacenterowner:FabricAdmin`.

```
# efa auth ldapconfig add --name ldap1 --host 10.x.x.x
--bind-user-name cn=x,dc=y,dc=com --bind-user-password xxx
--user-search-base ou=people,dc=y,dc=com --user-role-attribute role
--user-role-attribute-key datacenterowner
```

This example configures the **--user-member-attribute** for a user entry in LDAP that has an attribute of `memberOf`.

```
# efa auth ldapconfig add --name ldap1 --host 10.x.x.x
--bind-user-name cn=x,dc=y,dc=com --bind-user-password xxx
--user-search-base ou=people,dc=y,dc=com --user-member-attribute memberOf
```

This example configures the attributes required when LDAP groups are not in the same search base as the users in the groups.

```
# ldapconfig add --name ldap1 --host 10.x.x.x --bind-user-name cn=x,dc=y,dc=com
--bind-user-password xxx --user-search-base ou=people,dc=x,dc=com
--group-search-base ou=groups,dc=x,dc=in --group-member-user-attribute dn
--group-member-mapping-attribute memberUid --group-object-class posixGroup
```

This example maps an LDAP group to an EFA role.

```
# efa auth rolemapping add --name group1 --type GROUP --role SystemAdmin
```

This example configures LDAP Active Directory.

```
# efa auth ldapconfig add --name ldap1 --host 10.x.x.x
--bind-user-name cn=x,dc=y,dc=com
--bind-user-password xxx --user-search-base ou=people,dc=y,dc=com
--user-object-class user
--user-login-attribute sAMAccountName --user-member-attribute memberOf
```

efa auth rolemapping

Assigns EFA roles to a user or an LDAP group.

Syntax

```
efa auth rolemapping add [--name user-name | --role efa-role | --type user-type ]  
efa auth rolemapping show  
efa auth rolemapping remove --id id
```

Parameters

--name *user-name*

Specifies the user name or the LDAP group name.

--role *efa-role*

Specifies the role that you want to assign. It can be one of the following: FabricAdmin, SecurityAdmin, NetworkOperator, SystemDebugger, SystemAdmin, <Tenant>Admin. The Tenant Administrator is assigned dynamically when the tenant is created. The role name has the following format: <Tenant-name>Admin.

--type *user-type*

Specifies the type of user. Enter either *user* or *group*.

--id *id*

ID of the role mapping.

Examples

This example assigns the role of Fabric Admin to a user named fabricuser.

```
# efa auth rolemapping add --name fabricuser --role FabricAdmin --type user  
Successfully added the role mapping
```

This example assigns the role of NetworkOperator to a group named viewer.

```
# efa auth rolemapping add --name viewer --role NetworkOperator --type group  
Successfully added the role mapping.
```

This example displays all assigned mappings.

```
# efa auth rolemapping show  
ID Name      Role           Type  
1  efauser    SystemAdmin    USER  
2  fabricuser FabricAdmin     USER  
3  viewer     NetworkOperator GROUP
```

This example deletes the role for the user with ID 3.

```
# efa auth rolemapping remove --id 3  
Deleted role mapping successfully
```

efa auth settings token

Configures and displays the expiration settings for authentication tokens.

Syntax

```
efa auth settings token update [ --type { cli | access | refresh } | --
    hours num | --minutes num ]
efa auth settings token show
```

Parameters

--type { cli | access | refresh }

Specifies the type of token you want to update.

--hours num

Specifies the number of hours that should elapse before a token expires.

--minutes num

Specifies the number of minutes that should elapse before a token expires.

Examples

This example sets the CLI token to expire after two hours.

```
$ efa auth settings token update --type CLI --hours 2

Successfully updated the token expiry time.
+-----+-----+-----+
| Type | Hours | Minutes |
+-----+-----+-----+
| CLI  | 2     | 0       |
+-----+-----+-----+
--- Time Elapsed: 193.455466ms ---
```

This example displays the current expiration settings.

```
$ efa auth settings token show

+-----+-----+-----+
| Type   | Hours | Minutes |
+-----+-----+-----+
| ACCESS | 1     | 0       |
+-----+-----+-----+
| REFRESH | 8    | 0       |
+-----+-----+-----+
| CLI    | 2     | 0       |
+-----+-----+-----+
```

efa certificate device install

Installs HTTPS and OAuth2 certificates on SLX devices.

Syntax

```
efa certificate device install [ --ip ip-addr | --certType { https | token } | --https-certificate | --https-key ]
```

Command Default

By default, certificates are not installed.

Parameters

--ip *ip-addr*

Specifies a comma-separated list of the IP addresses of the SLX devices on which you want to install the certificate.

--certType { **https** | **token** }

Specifies the type of certificate you are installing. If you do not specify a type, this command creates both types.

--https-certificate

Specifies the file name of the certificate.

--https-key

Specifies the file name of the certificate key.

Usage Guidelines

During the registration of an SLX device in EFA, the following configuration changes are made on the device.

- The public certificate for verifying an EFA token is copied to the device as an OAuth2 certificate.
- EFA generates the HTTPS certificate for the SLX device. The certificate is copied to the device, HTTP mode is turned off on the device, and HTTPS is enabled on the device.
- OAuth2 is enabled as the primary mode of authentication. Fallback is set to "local login."

You can use the **efa inventory device list** command to verify the status of the certificates on the device. If the **Cert/Key Saved** column contains "N," then certificates are not installed. You can then use the **efa certificates device install** command to install the certificates.

Examples

This example installs the HTTPS certificate on two devices.

```
efa certificates device install --ip 1.1.1.1,2.2.2.2 --certType https
```

efa certificate server

Installs a third-party certificate that was acquired through a trusted certificate authority.

Syntax

```
efa certificate server [ --certificate cert-filename --key key-filename  
  --cacert cert-filename ]
```

Parameters

--certificate *cert-filename*
Specifies the file name of the certificate.

--key *key-filename*
Specifies the file name of the certificate key.

--cacert *cert-filename*
Specifies the file name of the third-party certificate.

Examples

This example installs a certificate titled `my_server_162.pem` with key-filename of `my_server_162.key` and cert-filename of `ca-chain.pem`.

```
$ efa certificate server --certificate=my_server_162.pem  
--key=my_server_162.key --cacert=ca-chain.pem  
  
Please wait as the certificates are being installed...  
  
Certificates were installed!  
  
--- Time Elapsed: 30.946303683s ---
```


efa certificate server renew

Renew self-signed certificate on EFA.

Syntax

efa certificate server renew

This is not applicable if third-party certificates are installed on the system.

Examples

```
(efa:xmcdev)xmcdev@xmcdev-virtual-machine1:/opt/checkouts/efa/efa$ efa certificate server
renew
Certificate renewal is successful
--- Time Elapsed: 33.516064167s ---
```

efa fabric clone

Clones source fabric.

Syntax

```
efa fabric clone [ --source source-fabric-name | --destination destination-fabric-name ]
```

Parameters

--source

Name of the fabric to be cloned.

--destination

New name of the cloned fabric.

Usage Guidelines

Cloning can expedite the deployment of fabrics across different sites or data centers. For fabrics in two different data centers to look exactly the same for disaster recovery purposes, create a clone for the source fabric.

This command clones all the fabric properties - type, stage, description, fabric settings - but not the devices on the fabric.

Examples

The following example clones BLR_FABRIC into PUN_FABRIC.

```
efa fabric clone --source BLR_FABRIC --destination PUN_FABRIC
```

efa fabric configure

Configures the underlay and overlay on all fabric devices.

Syntax

```
efa fabric configure [ --name fabric-name ] [ --force ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--force

Forces the configuration on the devices.

Usage Guidelines

If the **--force** option is used, all the devices are removed and added back to the fabric. This action can result in **config remove** and **add on** all the devices.

If the addition of devices to a fabric is successful, the underlay and overlay is configured on all the devices of the fabric using the `efa fabric configure` command.

Examples

```
(efa:extreme)extreme@tpvm:~$ efa fabric configure --name fabric-2 --force
Configure fabric force
WARNING: Fabric configure with 'force' option 'removes' all the devices from fabric after
clearing
existing configuration and 'adds' back all the devices to fabric. Do you want to proceed
[y/n]?
y
Fetching devices from fabric fabric-2
Removing device from fabric fabric-2
Remove Device(s) [Success]
Removal of device with ip-address = 10.24.80.159 [Succeeded]
Removal of device with ip-address = 10.24.80.158 [Succeeded]
Updating the devices in inventory for fabric fabric-2
Update of device with ip-address = 10.24.80.158 [Succeeded]
Update of device with ip-address = 10.24.80.159 [Succeeded]
Clearing the devices in the fabric fabric-2
Updating the devices in inventory for fabric fabric-2
Update of device with ip-address = 10.24.80.158 [Succeeded]
Update of device with ip-address = 10.24.80.159 [Succeeded]
Add the devices to fabric fabric-2
Add Device(s) [Success]
Addition of Leaf device with ip-address = 10.24.80.158 [Succeeded]
Addition of Leaf device with ip-address = 10.24.80.159 [Succeeded]
Validate Fabric [Success]
Configure Fabric [Success]
Please verify the fabric physical/underlay topology using 'efa fabric topology show
{physical | underlay}
```

```
' before attempting tenant configuration on the fabric.  
- Time Elapsed: 1m52.840570796s -
```

efa fabric create

Creates a fabric.

Syntax

```
efa fabric create [ --name fabric-name | --type { clos | non-clos } | --stage { 3 | 5 } | --description description ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--type { **clos** | **non-clos** }

Specifies the type of the fabric. The default is Clos.

--stage { **3** | **5** }

Specifies the type of fabric to create. The default is a 3-stage fabric. Stage 0 is considered the default stage for Clos and is ignored for non-Clos.

--description *description*

Describes the product.

Examples

This example creates a three-stage fabric.

```
efa fabric create --name clos-fabric --type clos --stage 3
```

efa fabric debug clear-config

Clears the underlay or overlay configuration from the device and recovers the device from erroneous conditions.

Syntax

```
efa fabric debug clear-config [ --device device ip | --reference-fabric fabric name ]
```

Parameters

--device

Device IP address.

--reference-fabric

Name of the fabric and the device it will eventually belong to.

Examples

```
$ efa fabric debug clear-config --device 10.24.4810.24.48.131,10.24.51.135,10.24.51.131,10.25.225.58,10.24.80.139.131,10.24.51.135,10.24.51.131,10.25.225.58,10.24.80.139
```

efa fabric debug config-gen-reason

Obtains the configuration generation reason for a particular fabric device.

Syntax

```
efa fabric debug config-gen-reason [ --device device ip | --name fabric name ]
```

Parameters

--device

The device IP address.

--name

Name of the fabric to which the device belongs.

Examples

```
efa fabric debug config-gen-reason --device 10.24.80.139 --name BLR_FABRIC
```

efa fabric debug device drift

Reconciles or displays the device configuration drift between the device and intended fabric configuration.

Syntax

```
efa fabric debug device drift [ --name string ip | --ip string | --reconcile ]
```

Parameters

- name** *string*
Specifies the name of the fabric.
- ip** *string*
Specifies the device IP.
- reconcile**
Reconciles the device configuration.

efa fabric debug service lock

Lists all lock statuses for fabric services.

Syntax

```
efa fabric debug service lock
```

efa fabric debug service smartdrc

Enables debug of smartdrc.

Syntax

```
efa fabric debug service smartdrc [ --enable | --help ]
```

Parameters

--enable { true | false }

Enables or disables debug for smartdrc. Use the following format:

--enable=true

--enable=false

--help

Help for smartdrc debug.

efa fabric delete

Deletes the fabric from inventory.

Syntax

```
efa fabric delete [ --name fabric-name ] [ --force ]
```

Parameters

--name *fabric-name*

Name of the fabric to be deleted.

--force

Forces the deletion of fabric even if the fabric has devices.

Usage Guidelines

Deletion of a fabric is not allowed if the fabric has one or more devices. You must delete all the devices from the fabric prior to deleting the fabric.

Forced deletion of a fabric removes the devices from the fabric but not from the inventory.

Examples

The following example deletes the fabric BLR_FABRIC.

```
efa fabric delete --name BLR_FABRIC
```

efa fabric device add

Adds a device to an existing fabric.

Syntax

```
efa fabric device add [ --name fabric-name | --ip device-ip | --role
  { leaf | spine | super-spine | border-leaf } | --leaf-type { single-
homed | multi-homed } | --hostname hostname | --asn local-asn | --vtep-
loopback id | --loopback port-num | --pod name | --username username |
--password password | --rack name ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--ip *device-ip*

Specifies the pair of device IP addresses for rack, or a single IP address for non-rack.

--role { **leaf** | **spine** | **super-spine** | **border-leaf** }

Specifies the device role.

--leaf-type { **single-homed** | **multi-homed** }

Specifies whether the leaf is single-homed or multi-homed.

--hostname *hostname*

Specifies the host name.

--asn *local-asn*

Specifies the local ASN.

--vtep-loopback *id*

Specifies the VTEP loopback ID.

--loopback *port-num*

Specifies the loopback port number.

--pod *name*

Specifies the name of the pod.

--username *username*

Specifies the user name for the device.

--password *password*

Specifies the password for the device.

--rack *name*

Specifies the name of the rack.

Usage Guidelines

A device must be registered with Inventory Service before being added to a fabric. The Fabric Service supports IP numbered configuration. Each interface on a link between leaf and spine is assigned an IP address. EBGP peering uses these IP addresses.

Device credentials must be provided as part of this command if the devices are not already registered with the inventory.

If you provide the user name and password, then the device is automatically registered with the inventory service.

If you do not provide the user name and password, then you must explicitly register the device with the inventory service.

Examples

This example adds a pair of devices to the specified fabric and the specified rack.

```
efa fabric device add --name extr-fabric --ip 10.24.80.134,10.24.80.135
--rack room1-rack1 --username admin --password password
```

efa fabric device add-bulk

Adds multiple devices to a fabric.

Syntax

```
efa fabric device add-bulk [ --name fabric-name | --ip device-ip | --leaf list | --border-leaf list | --three-stage-pod name | --five-stage-pod name | --spine list | --super-spine list | --username username | --password password | --rack name | --border-leaf-rack list | --border-leaf-ip device-ip ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--ip *device-ip*

Specifies the pair of device IP addresses for rack.

--leaf *list*

Specifies the comma-separated list of leaf IP addresses or host names.

--border-leaf *list*

Specifies the comma-separated list of border leaf IP addresses or host names.

--three-stage-pod *name*

Specifies the name of the leaf or spine pod.

--five-stage-pod *name*

Specifies the name of the super-spine pod.

--spine *list*

Specifies the comma-separated list of spine IP addresses or host names.

--super-spine *list*

Specifies the comma-separated list of super spine IP addresses or host names.

--username *username*

Specifies the user name for the list of devices.

--password *password*

Specifies the password for the list of devices.

--rack *name*

Specifies the rack name.

--border-leaf-rack *list*

Specifies the border leaf rack name.

--border-leaf-ip *device-ip*

Specifies the pair of border leaf device IPs for the rack.

Usage Guidelines

Ensure you perform the following operations before running this command:

- Run the **efa inventory device register --ip <list-of-device-ips>** command.
- Run the **efa inventory device interface list --ip <device-ip>** command. In the output of the command, verify that the states of the port links are as you expected (in the Admin Status and Oper Status fields). If not, manually check the physical cabling and fix any issues. Then continue with the **efa fabric device add-bulk** operation.

If you provide a user name and password, devices are automatically registered with the inventory service.

If you do not provide a user name and password, you must explicitly register the devices with the inventory service.

A single “three-stage-pod” and “five-stage-pod” can be provided each time you run the command.

If, during device registration, the device is found to be in maintenance mode, then an error is returned and registration fails. To take the device out of maintenance mode before reattempting device registration, you must perform a manual procedure.

Examples

This example adds two leaf IP addresses and two border leaf IP addresses to the specified fabric.

```
efa fabric device add-bulk --name BLR_FABRIC --leaf 10.24.48.131,10.24.51.135
--border-leaf 10.24.51.131,10.25.225.58 --spine 10.24.80.139 --username admin --password
password
```

efa fabric device remove

Removes an existing device from a fabric.

Syntax

```
efa fabric device remove [ --name fabric-name | --ip device-ips ] [ --no-device-cleanup]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--ip *device-ips*

Specifies a comma-separated list of device IP addresses to remove.

--no-device-cleanup

Specifies that you do not want to clean up the configuration on the devices.

Usage Guidelines

If the `no-device-cleanup` option is used, the configuration pushed by the automation engine is not cleaned up from the fabric devices. Removal of a device from the fabric does not delete the device from inventory. You must explicitly delete the device from inventory.

Examples

This example removes five devices and does not clean up the configuration.

```
efa fabric device remove --ip 10.24.48.131,10.24.51.135,10.25.225.58,  
10.24.51.131,10.24.80.139 --name BLR_FABRIC --no-device-cleanup
```


efa fabric error show

Displays the name of the fabric, the error types, and the reasons for the errors.

Syntax

```
efa fabric error show [ --name fabric-name | --export string ]
```

Parameters

- name** *fabric-name*
Specifies the name of the fabric.
- export** *string*
Exports error details to a CSV file.

Usage Guidelines

Topology validation occurs during the addition of a device and during fabric configuration. Use this command to display any errors that occur during validation and, optionally, to export them to a CSV file.

Errors that occur during the add, validate, and configure phases are persisted in the database.

Examples

```
efa fabric device error show --name BLR_FABRIC
```

efa fabric execution

Displays or deletes the execution history for the fabric service.

Syntax

```
efa fabric execution delete [ --days number-of-days ]  
efa fabric execution show [ --id execution-id | --limit number-of-executions | --status { failed | succeeded | all } ]  
efa fabric execution show-event [ --device ip-address | --execution-id UUID ]
```

Parameters

- days** *number-of-days*
Deletes execution entries older than the specified number of days (default 30).
- device** *ip-address*
Filters on the IP address.
- execution-id** *UUID*
Filters on the execution UUID.
- id** *execution-id*
Filters the executions based on the execution id. The **--limit** and **--status** options are ignored when the **--id** option is used.
- limit** *number-of-executions*
Limits the number of executions to be listed. A value of 0 lists all the executions. The default is 10.
- status** { **failed** | **succeeded** | **all** }
Filters the executions based on the status. The default is **all**.

Examples

```
efa fabric execution show
```

efa fabric setting

Displays or updates the IP fabric settings.

Syntax

```
efa fabric setting show [ --name fabric-name | --advanced options ]

efa fabric setting update [--name fabric-name | --p2p-link-range fabric-
name | --loopback-ip-range ip-pool | --rack-l3-backup-ip-range ip-
pool | --loopback-port-number id | --vtep-loopback-port-number id |
--spine-asn-block asn-pool | --super-spine-asn-block asn | --leaf-
asn-block asn-pool | --border-leaf-asn-block asn-pool | --rack-asn-
block asn-pool | --rack-border-leaf-asn-block asn-pool | --anycast-
mac-address mac | --ipv6-anycast-mac-address mac | --mac-aging-
timeout timeout | --mac-aging-conversation-timeout timeout | --mac-
move-limit limit | --duplicate-mac-timer time | --duplicate-mac-
timer-max-count max| --configure-overlay-gateway { yes | no }| --bfd-
enable { yes | no }| --bfd-tx interval | --bfd-rx interval | --bfd-
multiplier multiplier| --bgp-multihop tth | --max-paths paths| --
allow-as-in path | --mtu size | --ip-mtu size | --leaf-peer-group
group-name | --optimized-replication-enable { yes | no }| --spine-
peer-group group-name | --super-spine-peer-group group-name | --rack-
underlay-ebgp-peer-group group-name | --rack-overlay-ebgp-peer-group
group-name | --mctlip-range pool | --lacp-timeout { long |
short }| --mct-port-channel id | --single-rack-deployment { yes | no }
| --control-vlan id | --control-ve id | --vni-auto-map { yes | no } |
--backup-routing-enable { yes | no }| --backup-routing-ipv4-range
range | --backup-routing-ipv6-range range | --optimized-replication-
enable { yes | no }| --mdtgroup-range range | --default-mdtgroup
ipv4-address | --md5-password password | --md5-password-enable { yes
| no }| --bgp-dynamic-peer-listen-limit peers]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--p2p-link-range *ip-pool*

Specifies the IP address pool used for P2P link configuration.

--loopback-ip-range *ip-pool*

Specifies the IP address pool used for P2P link configuration.

--rack-l3-backup-ip-range *ip-pool*

Specifies the IP address pool for Layer 3 backup.

--loopback-port-number *id*

Specifies the loopback ID on the device to be used as the donor IP interface for the link between leaf and spine. Valid values range from 1 through 255.

- vtep-loopback-port-number** *id*
Specifies the loopback ID on the device to be used as a VTEP IP interface. Valid values range from 1 through 255.
- spine-asn-block** *asn-pool*
Specifies the ASN pool for spine nodes or one ASN.
- super-spine-asn-block** *asn*
Specifies one ASN.
- leaf-asn-block** *asn-pool*
Specifies the ASN pool for leaf nodes.
- border-leaf-asn-block** *asn-pool*
Specifies the ASN pool for border leaf nodes.
- rack-asn-block** *asn pool*
Specifies the ASN pool for rack nodes. Valid values are numbers from 4200000000 through 4200065534.
- rack-border-leaf-asn-block** *asn pool*
Specifies the ASN pool for rack border leaf nodes. Valid values are numbers from 4200065535 through 4200065635.
- anycast-mac-address** *asn pool*
Specifies the IPv4 anycast MAC address.
- ipv6-anycast-mac-address** *mac*
Specifies the IPv6 anycast MAC address.
- mac-aging-timeout** *timeout*
Specifies the MAC aging timeout. Valid values are 0 (for no timeout) or 60 through 86400.
- mac-aging-conversation-timeout** *timeout*
Specifies the MAC conversational aging timeout in seconds. Valid values are 0 (for no timeout) or 60 through 100000.
- mac-move-limit** *limit*
Specifies the MAC move detect limit. Valid values range from 5 through 500.
- duplicate-mac-timer** *time*
Specifies the number of seconds for detecting a duplicate MAC address.
- duplicate-mac-timer-max-count** *max*
Specifies the maximum number of duplicate MAC addresses that can be detected.
- configure-overlay-gateway** { **yes** | **no** }
Configures an overlay gateway.
- bfd-enable** { **yes** | **no** }
Turns bidirectional forwarding detection (BFD) on or off.
- bfd-tx** *interval*
Specifies the BFD minimum transmit interval in milliseconds. Valid values range from 50 through 30000.

--bfd-rx *interval*

Specifies the BFD minimum receive interval in milliseconds. Valid values range from 50 through 30000.

--bfd-multiplier *multiplier*

Specifies the number of times that a packet is missed before BFD declares the neighbor is down. Valid values range from 3 through 50.

--bgp-multihop *ttl*

Determines the TTL value for EBGP neighbors that are not on directly connected networks. Valid values range from 1 through 255.

--max-paths *paths*

Specifies the maximum number of paths over which packets can be forwarded. Valid values range from 1 through 64.

--allow-as-in *path*

Turns off the AS_PATH check of the routes learned from the AS. Valid values range from 1 through 10.

--mtu *size*

- For SLX-OS 21.1 and later versions: The MTU size in bytes. Valid values range from 1500 through 9216.
- For other SLX-OS versions: The MTU size in bytes. Valid values range from 1548 through 9216.

--ip-mtu *size*

Specifies the IPv4 or IPv6 MTU size. Valid values range from 1300 through 9194.

--leaf-peer-group *group-name*

Specifies the name of the leaf peer group. Valid name length is from 1 through 63 characters.

--spine-peer-group *group-name*

Specifies the name of the spine peer group. Valid name length is from 1 through 63 characters.

--super-spine-peer-group *group-name*

Specifies the name of the super spine peer group. Valid name length is from 1 through 63 characters.

--rack-underlay-ebgp-peer-group *group-name*

Specifies the name of the EBGP peer group of the rack underlay. Valid name length is from 1 through 63 characters.

--rack-overlay-ebgp-peer-group *group-name*

Specifies the name of the EBGP peer group of the rack overlay. Valid name length is from 1 through 63 characters.

--mctlink-ip-range *pool*

Specifies the IP address pool to be used for MCT peering.

--lACP-timeout { **long** | **short** }

Specifies LACP timeout for MCT member ports. Accepted values are long and short.

--mct-port-channel *id*

Specifies the port-channel interface ID to be used as an MCT peer-interface. Valid values range from 1 through 64.

--single-rack-deployment { **yes** | **no** }

Specifies whether single-rack deployment is used. The default is no.

--control-vlan *id*

Specifies the ID of the VLAN to be used as the MCT cluster control VLAN. Valid values range from 1 through 4090.

--control-ve *id*

Specifies the ID of the VE to be used as the MCT cluster control VE. Valid values range from 1 through 4090.

--vni-auto-map { **yes** | **no** }

Specifies whether to automatically map the VTEP VLAN (or bridge domain) to VNI mode.

--backup-routing-enable { **yes** | **no** }

Turns backup routing on or off.

--backup-routing-ipv4-range *range*

Specifies the IPv4 backup routing range.

--backup-routing-ipv6-range *range*

Specifies the IPv6 backup routing range.

--optimized-replication-enable { **Yes** | **No** }

Enables optimized replication.

--mdtgroup-range *range*

Specifies the IPv4 multicast address range in IP prefix format. For example: 239.0.0.0/8

--default-mdtgroup *ipv4-address*

Specifies the default IPv4 multicast address. The address must be from the MDT group range.

---md5-password *password*

Specifies the BGP MD5 password to be used on fabric links. If the password is a word, the following rules apply: 1-80 for plain text and 1-164 for encrypted string. The default is "<prompt>".

---md5-password-enable { **yes** | **no** }

Enables or disables the BGP MD5 password on fabric links.

---bgp-dynamic-peer-listen-limit *peers*

Specifies the maximum number of dynamic BGP peers that can be operational across the VRFs in the SLX. Valid values range from 1 through 2400. The default is 100.

Examples

These examples update optimized replication, the MDT group, and the default MDT group.

```
# efa fabric setting --update --mdtgroup-range 239.0.0.0/8 --name clos_fabric
# efa fabric setting --update --default-mdtgroup 239.1.1.1 --name clos_fabric
```

In this example, MD5 authentication is created or updated:

```
efa fabric setting update --name fabric1 --md5-password-enable yes
Please supply a password for BGP MD5 authentication on fabric links:
efa fabric configure --name fabric1
```

efa fabric show

Displays the details of the fabric.

Syntax

```
efa fabric show [--name fabric-name ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

Usage Guidelines

Displays the details of all fabrics when the `name` option is not provided.

Displays the details of the specified fabric when the `name` option is provided.

Examples

This example displays details for the BLR_FABRIC.

```
efa fabric show --name BLR_FABRIC
```


efa fabric show-config

Displays the configuration of the specified fabric.

Syntax

```
efa fabric show-config [ --name fabric-name | --device-role { leaf | spine | super-spine | border-leaf } | --ip ip-address ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

--device-role { **leaf** | **spine** | **super-spine** | **border-leaf** }

Specifies the role of devices for which configuration is displayed.

--ip *ip-address*

Specifies the IP address of the devices for which configuration is displayed. Used with the **--device-role** option.

Examples

```
efa fabric show-config --name BLR_FABRIC --device-role border-leaf  
--ip 10.25.225.58
```

efa fabric show summary

Displays a summary of the specified fabric.

Syntax

```
efa fabric show summary [--name fabric-name ]
```

Parameters

--name *fabric-name*

Specifies the name of the fabric.

Usage Guidelines

Displays the details of all fabrics when the `name` option is not provided.

Displays the details of the specified fabric when the `name` option is provided.

Examples

This example displays a summary for the BLR_FABRIC fabric.

```
efa fabric show summary --name BLR_FABRIC
```

efa fabric topology show overlay

Displays the overlay connectivity (VxLAN tunnels) of the devices in a fabric.

Syntax

```
efa fabric topology show overlay [--name name ]
```

Parameters

--name *name*

Specifies the name of the fabric.

Examples

This example displays the overlay topology information for the fabric.

```
# efa fabric topology show overlay --name clos_fabric
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| OVERLAY ECAP TYPE | TUNNEL TYPE | SOURCE LEAF IP | DESTINATION LEAF IP | SOURCE VTEP IP |
| DESTINATION VTEP IP | OVERLAY ADMIN STATE | OVERLAY OPER STATE | | |
+-----+-----+-----+-----+-----+
| vxlan            | unicast     | 10.24.80.158   | 10.24.80.159       | 172.31.254.7   | | |
| 172.31.254.9     | | up         | | up               | |               |
| vxlan            | unicast     | 10.24.80.159   | 10.24.80.158       | 172.31.254.9   |
| 172.31.254.7     | | up         | | up               | |               |
| vxlan            | multicast   | 10.24.80.158   | |                   | 172.31.254.7   |
| 239.1.1.1        | | up         | | up               | |               |
| vxlan            | multicast   | 10.24.80.159   | |                   | 172.31.254.9   |
| 239.1.1.1        | | up         | | up               | |               |
+-----+-----+-----+-----+-----+
```

efa fabric topology show underlay

Displays the underlay connectivity (BGP peers) of the devices in a fabric for the default VRF.

Syntax

```
efa fabric topology show underlay [--name name ]
```

Parameters

--name *name*

Specifies the name of the fabric.

Examples

This example displays the underlay topology information for the fabric.

```
efa fabric topology show underlay --name nc
+-----+-----+-----+-----+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
+-----+-----+-----+-----+
| 10.20.61.90      | 10.20.61.91          | 172.31.254.1           | 10.20.20.3   |
| 10.20.61.91      | 10.20.61.90          | 172.31.254.3           | 10.20.20.2   |
+-----+-----+-----+-----+
| SOURCE DEVICE ASN | DESTINATION DEVICE ASN | NEIGHBOR AFI STATE | NEIGHBOR SAFI | UNDERLAY STATE |
+-----+-----+-----+-----+
| 4200000000        | 4200000000            | ipv4                 | unicast        |                  |
| 4200000000        | 4200000000            | ipv4                 | unicast        |                  |
+-----+-----+-----+-----+
```

efa fabric topology show physical

Displays the physical connectivity (LLDP neighbors) of the devices in a fabric.

Syntax

```
efa fabric topology show physical [ --name name ]
```

Parameters

--name *name*

Specifies the name of the fabric.

Examples

This example displays a portion of the physical topology information for the fabric.

```
efa fabric topology show physical --name clos_fabric
+-----+-----+-----+-----+-----+
| SOURCE      | SOURCE      | DESTINATION  | DESTINATION  | SOURCE      |
| NODE IP     | NODE ROLE   | NODE IP      | NODE ROLE    | NODE INTERFACE |
+-----+-----+-----+-----+-----+
| 10.24.80.135 | LEAF        | 10.25.225.163 | LEAF         | 0/39          |
| 10.24.80.135 | LEAF        | 10.25.225.163 | LEAF         | 0/40          |
| 10.24.80.135 | LEAF        | 10.24.80.134  | LEAF         | 0/46          |
| 10.24.80.134 | LEAF        | 10.24.80.135  | LEAF         | 0/46          |
+-----+-----+-----+-----+-----+
| DESTINATION  | SOURCE DEVICE| DESTINATION  | DESTINATION  |
| NODE INTERFACE | MULTI HOMED  | MULTI HOMED  | MULTI HOMED  |
+-----+-----+-----+-----+-----+
| 0/39         | TRUE         | FALSE        | FALSE        |
| 0/40         | TRUE         | FALSE        | FALSE        |
| 0/46         | TRUE         | TRUE         | TRUE         |
| 0/46         | TRUE         | TRUE         | TRUE         |
```

efa inventory admin-state

Changes, displays, and deletes device state, state change history, or state change record.

Syntax

```
efa inventory admin-state up [ --ip ip-addr ]  
efa inventory admin-state down [--ip ip-addr ]  
efa inventory admin-state detail [--uuid uuid ]  
efa inventory admin-state delete [ --key ip-addr ]  
efa inventory admin-state history [--ip ip-addr ]  
efa inventory admin-state show [--ip ip-addr ]
```

Parameters

--ip *ip-addr*

Specifies the IP address of the device.

--key *ip-addr*

Specifies the IP address of the device or config-backup UUID to be deleted.

--uuid *uuid*

Specifies the ID of the configuration backup.

Usage Guidelines

When a device changes to `admin down` state, the device goes into maintenance mode.

When a device changes to `admin up` state, the device is taken out of maintenance mode.

For more information, see "Administered Partial Success" in the [Extreme Fabric Automation Administration Guide, 2.7.0](#).

Examples

This example changes the device state to `admin up` and generates a UUID to use in the **efa inventory admin-state detail** version of the command.

```
$ efa inventory admin-state up --ip 10.24.80.158  
AdminStateUp [success]  
Admin State Up execution UUID: 8d9fa0cf-dc76-42cc-ac7a-57902a47c1b2
```

This example changes the device state to `admin down` and generates a UUID to use in the **efa inventory admin-state detail** version of the command.

```
$ efa inventory admin-state down --ip 10.24.80.158  
AdminStateDown [success]  
Admin State Down execution UUID: 28eb0845-7a7a-4851-b453-b3020c6900f2
```

This example displays the history of the admin status for the specified device.

```
# efa inventory admin-state history --ip 10.24.80.158
-----
UUID                               Device IP      Admin State  Status
-----
8d9fa0cf-dc76-42cc-ac7a-57902a47c1b2  10.24.80.158  up           success
-----
28eb0845-7a7a-4851-b453-b3020c6900f2  10.24.80.158  down        success
-----
Device State Change Records
```

This example shows details of admin state changes for the specified UUID.

```
$ efa inventory admin-state detail --uuid
28eb0845-7a7a-4851-b453-b3020c6900f2
-----+
NAME          VALUE
-----+
UUID          28eb0845-7a7a-4851-b453-b3020c6900f2
-----+
Device IP     10.24.80.158
-----+
Admin State Action down
-----+
Status        success
-----+
Maintenance Mode Enable Status    success
-----+
Start Time    2020-07-30 10:04:33.982447 +0000 UTC
-----+
Last Modified 2020-07-30 10:04:38.437425 +0000 UTC
-----+
Duration      4.454976211s
-----+
```

This example shows details of the state of the specified device.

```
$ efa inventory admin-state show --ip 10.24.80.158
-----+
NAME          VALUE
-----+
Device IP     10.24.80.158
-----+
Admin State   up
-----+
Health Check Status  Disable
-----+
```

This example deletes the state change instance for the specified IP address.

```
$ efa inventory admin-state delete --key 10.24.80.159
-----
USERKEY      STATUS      ERROR
-----
10.24.80.159  deleted
-----
```

efa inventory config-backup

Configures the backup for device configuration.

Syntax

```
efa inventory config-backup history [ --ip ip-address ]  
efa inventory config-backup execute [ --ip ip-address ]  
efa inventory config-backup delete [ --key ip-address ]  
efa inventory config-backup detail [ --uuid cb-uuid | --show-config | --file-dump filename ]
```

Parameters

--ip *ip-address*

Specifies the IP address of the device.

--key *ip-address*

Specifies the IP address of the device or config-backup UUID to be deleted.

--uuid *cb-uuid*

Specifies the ID of the configuration backup.

--show-config

Displays configuration text.

--file-dump *filename*

Saves configuration text to the specified file.

Usage Guidelines

Use the **efa inventory config-backup history** command to display the history of configuration backups.

Use the **efa inventory config-backup execute** command to run the configuration backup for the specified device.

Use the **efa inventory config-backup delete** command to delete the configuration backup record.

Use the **efa inventory config-backup detail** command to display the details of the configuration backup and to, optionally, show the text of the configuration and save the text to a file.

Examples

The following examples show **efa inventory config-backup** commands.

```
# efa inventory config-backup execute --ip 10.24.14.133  
# efa inventory config-backup history --ip 10.24.14.133  
# efa inventory config-backup detail --uuid 1111-1111-1111 --show-config
```



```
# efa inventory config-backup detail --uuid 1111-1111-1111 --show-config
--file-dump fileABC
# efa inventory config-backup delete --key 10.24.14.133
# efa inventory config-backup delete --key 1111-1111-111
```

efa inventory config-replay

Displays, runs, and deletes the configuration replay for the specified device.

Syntax

```
efa inventory config-replay history [--ip ip-addr ]  
efa inventory config-replay execute [--ip ip-addr | --uuid cr-uuid | --  
  ssid ssid | --startup-config | --no-reboot ]  
efa inventory config-replay delete [--key ip-addr ]  
efa inventory config-replay detail [--uuid cr-uuid ]
```

Parameters

--ip *ip-addr*
Specifies the IP address of the device.

--uuid *cr-uuid*
Specifies the config replay ID or the UUID of config backup job.

--ssid *ssid*
Specifies the SSID of the config-backup job.

--startup-config
Copies the backup file to the startup-config file.

--no-reboot
Configures replay without rebooting the device.

--key *ip-addr*
Specifies the IP address of the device or config-replay UUID to be deleted.

Usage Guidelines

The device configuration backup and replay feature enables backup of the device configuration based on inventory device settings, commands, and REST APIs.

Use the **efa inventory config-replay history** command to display the configuration replay history for the specified device.

Use the **efa inventory config-replay execute** command to run the configuration backup for the specified device.

Use the **efa inventory config-replay delete** command to delete the configuration replay record for the specified device.

Use the **efa inventory config-replay detail** command to display detailed replay history for the specified device.

Examples

The following examples show **efa inventory config-replay** commands.

```
# efa inventory config-replay execute --ip 10.24.14.133
--uuid 1111-1111-111

# efa inventory config-replay history --ip 10.24.14.133
# efa inventory config-replay detail --uuid 1111-1111-1111
# efa inventory config-replay delete --key 10.24.14.133
# efa inventory config-replay delete --key 1111-1111-111
# efa inventory config-replay execute --ip 10.24.14.133 --uuid 1111-1111-111
--startup-config

# efa inventory config-replay execute --ip 10.24.14.133 --uuid 1111-1111-111
--no-reboot
```

efa inventory debug devices-lock

Displays lock status of the devices.

Syntax

```
efa inventory debug devices-lock
```

Usage Guidelines

Diagnostic commands are developed and intended for specialized troubleshooting. Work closely with Extreme Networks technical support when running **debug** or **show system internal** commands and interpreting their results.

Examples

```
efa inventory debug devices-lock
+-----+
|           IP Address           | Locked |
+-----+
| 10.139.44.156                   | false |
+-----+
| 10.139.44.155                   | false |
+-----+
| Global DeviceWalkListLock       | false |
+-----+
| Global DeviceLockCollectionMapLock | false |
+-----+
| Global FabricEventLock          | false |
+-----+
| Global MapLock of Events        | false |
+-----+
| GoSwitch Repository             | false |
+-----+
| Service Lock                    | false |
+-----+
```

efa inventory debug devices-unlock

Unlocks devices forcefully if they were locked during the configuration backup process.

Syntax

```
efa inventory debug devices-unlock [--ip ip-list | --fabric fabric-name |
  --servicelock ]
```

Parameters

- ip** *ip-list*
List of IPs to unlock.
- fabric** *fabric-name*
Fabric with which devices are associated.
- servicelock**
Specify to unlock all service locks.

Examples

```
efa-client inventory debug devices-unlock --ip 10.20.246.5,10.20.246.6
+-----+-----+
| IP Address | Status |
+-----+-----+
| 10.20.246.6 | Success |
+-----+-----+
| 10.20.246.5 | Success |
+-----+-----+
Force Unlock Devices
--- Time Elapsed: 100.5453ms ---
```

```
efa-client inventory debug devices-unlock --servicelock
+-----+-----+
| IP Address | Status |
+-----+-----+
| Service Lock | Success |
+-----+-----+
Force Unlock Devices
--- Time Elapsed: 39.3574572s ---
```

```
efa-client inventory debug devices-unlock --fabric fabric1
+-----+-----+
| IP Address | Status |
+-----+-----+
| 10.20.246.6 | Success |
+-----+-----+
| 10.20.246.5 | Success |
+-----+-----+
Force Unlock Devices
```

efa inventory debug unblock-from-fwdl

Unlocks devices forcefully if they were locked during a firmware download.

Syntax

```
efa inventory debug unblock-from-fwdl [--ip ip-list | --fabric fabric-name ]
```

Parameters

--ip *ip-list*

List of IPs to unlock.

--fabric *fabric-name*

Fabric with which devices are associated.

Usage Guidelines

During firmware download, a lock is activated to prevent other services from doing any operations. This debug command allows you to unlock the lock in case, for some reason, the system does not unlock on its own.

Examples

```
efa-client inventory debug unblock-from-fwdl --ip 10.20.246.1,10.20.246.2
+-----+-----+
| IP Address | Status |
+-----+-----+
| 10.20.246.2 | Success |
+-----+-----+
| 10.20.246.1 | Failed |
+-----+-----+
Force Removal of FWDL
Status
--- Time Elapsed: 42.0891034s ---
```

```
efa-client inventory debug unblock-from-fwdl --fabric non-clos
+-----+-----+
| IP Address | Status |
+-----+-----+
| 10.20.246.2 | Success |
+-----+-----+
| 10.20.246.1 | Success |
+-----+-----+
Force Removal of FWDL
```

efa inventory device clear route-all

Clears all routes on devices.

Syntax

```
efa inventory device clear route-all [ --ip device-ips | --fabric fabric-name | --ipv4-only | --ipv6-only | -h, --help ]  
efa inventory device clear route-all
```

Parameters

--ip *device-ips*

Specifies comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Specifies the name of the fabric.

--ipv4-only

Optional flag to only clear IPv4 routes.

--ipv6-only

Optional flag to only clear IPv6 routes.

-h, --help

Help for device clear all routes.

Examples

The following example clears all routes on device with IP address 10.20.30.40.

```
efa inventory device clear route-all --ip 10.20.30.40
```

The following example clears only IPv4 routes on device with IP address 10.20.30.40.

```
efa inventory device clear route-all --ip 10.20.30.40 --ipv4-only
```

efa inventory device compare

Compares the device configuration with the configuration details that are saved in the application.

Syntax

```
efa inventory device compare [ --ip ]
```

Parameters

--ip

Specifies the IP address of the device to be compared.

efa inventory device delete

Deletes the device.

Syntax

```
efa inventory device delete [ --ip ip-addr | --fabric fabric-name ]
```

Parameters

--ip *ip-addr*

Specifies comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Fabric name for which devices needs to be deleted.

Usage Guidelines

If an inventory device is deleted while Drift and Reconcile is in process for the same device and the device is currently in maintenance mode, then the device remains in maintenance mode after deletion. To take the device out of maintenance mode before device registration, you must perform a manual procedure.

efa inventory device discovery-time list

Displays the configured device discovery interval for devices or for a fabric.

Syntax

```
efa inventory device discovery-time list [ --ips ip-addr | --fabric fabric-name]
```

Command Default

The default discovery interval is one hour. You can modify this interval by using the **efa inventory device discovery-time update** command.

Parameters

--ips *ip-addr*

Specifies a comma-separated list of the IP addresses for which you want to see the configured interval. For example, 10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Specifies the name of the fabric.

Examples

This example displays the device discovery interval for each device in the fabric named myFabric.

```
efa inventory device discovery-time list --fabric myFabric
```

This example displays the device discovery interval for two device IP addresses.

```
efa inventory device discovery interval --ips 1.1.1.1,2.2.2.2
```

efa inventory device discovery-time update

Configures the interval for periodic discovery of devices.

Syntax

```
efa inventory device discovery-time update [ --ip ip-addr | --fabric fabric-name ] [ --min minutes | --hour hours]
```

Command Default

The default discovery interval is one hour.

Parameters

--ip *ip-addr*

Specifies the IP address of the device for which you are configuring the interval.

--fabric *fabric-name*

Specifies the name of the fabric for which you are configuring the interval. All devices in this fabric will have the same interval.

--min *minutes*

Specifies the interval in minutes.

--hour *hours*

Specifies the interval in hours.

Usage Guidelines

Tenant and Fabric Services use periodic discovery to detect out-of-sync configurations on the devices.

These services act on the published events and update the database to reflect the status of the devices as in-sync and out-of-sync.

Examples

This example configures an interval of 20 minutes for a specific IP address.

```
efa inventory device discovery-time update --ip 1.1.1.1 --min 20
```

This example configures an interval of two hours for a fabric named myFabric.

```
efa inventory device discovery-time update --fabric myFabric --hour 2
```

efa inventory device execute-cli

Runs a specified command on one or more devices.

Syntax

```
efa inventory device execute-cli [ --ip ip-addr | --fabric fabric-devices | --role role-devices | --command command | --config { config-term | exec-mode } ]
```

Parameters

--ip *ip-addr*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-devices*

Specifies all devices in the specified fabric.

--role *role-devices*

Specifies devices in the fabric, based on their role.

--command *command*

Specifies a comma-separated or semicolon-separated list of commands to run on the devices.

--config { **config-term** | **exec-mode** }

Indicates whether commands are for config-term or exec-mode. The default value is `exec-mode`.

Examples

```
# efa inventory device execute-cli --ip 10.18.120.187  
--command "Interface ethernet 0/1, no reload-delay enable" --config
```

efa inventory device firmware-download commit

Commits new firmware where the firmware download was started with the NoAutoCommit flag.

Syntax

```
efa inventory device firmware-download commit [ --fabric fabric-name | --  
ip ip-addr ]
```

Parameters

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *ip-addr*

Specifies a comma-separated list of device IP addresses.

```
efa inventory device firmware-download commit --fabric fabric1  
Firmware Download Commit for fabric fabric1 - [success]  
+-----+-----+-----+-----+  
| DeviceIP      | Device FWDL Status                | Error |  
+-----+-----+-----+-----+  
| 10.20.246.6   | Firmware Committed Successfully   |       |  
+-----+-----+-----+-----+  
--- Time Elapsed: 1m25.765540754s ---
```

efa inventory device firmware-download execute

Clears out the prepared list for successfully completed devices. Any failed devices remain in the prepared list to allow for the run to be retried.

Syntax

```
efa inventory device firmware-download execute --fabric fabric-name |  
  [ --noMaintMode ] | [ --noAutoCommit] | --drc | --group-execution  
  { continue-on-error | stop-on-error } | --help
```

Parameters

--fabric *fabric-name*

Specifies the name of the fabric.

--noMaintMode

Ignores Maintenance Mode operations during firmware download. The option defaults to FALSE, which means Maintenance Mode is turned off during firmware download. If you add the `noMaintMode` option, Maintenance Mode operations are not performed during the firmware download process.

--noAutoCommit

Configures auto commit in the firmware download. By default, the firmware download process automatically commits the firmware following successful completion. If you add the `noAutoCommit` option, you must manually run a new **efa inventory device firmware-download commit** command to commit the downloaded firmware on the device.

--drc

During firmware download workflow DRC is controlled by user input flag "drc". If the flag is set to "true" and maintenance mode is enabled, then DRC is triggered. By default "drc" flag is "false". The following table describes the different scenarios for DRC during firmware download:

noMaint	drc	enable_on_reboot	FWDL Reboot Behavior	Comments
False	False	True/ False	<ul style="list-style-type: none"> Maintenance mode will be enabled before reboot NO DRC Maintenance mode will be disabled after reboot 	Default Case FWDL is driving maint-mode so enable_on_reboot has no effect
False	True	True/ False	<ul style="list-style-type: none"> Maintenance mode will be enabled before reboot DRC will be triggered Maintenance mode will be disabled on successful DRC 	FWDL is driving maint-mode so enable_on_reboot has no effect
True	False	True	<ul style="list-style-type: none"> enable_on_reboot mode will bring switch in maintenance mode NO DRC Maintenance mode will be disabled after reboot 	
True	False	False	<ul style="list-style-type: none"> No maintenance mode NO DRC 	

noMaint	drc	enable_on_reboot	FWDL Reboot Behavior	Comments
True	True	True	<ul style="list-style-type: none"> enable_on_reboot mode will bring switch in maintenance mode DRC will be triggered Maintenance mode will be disabled on successful DRC 	
True	True	False	Invalid use case; appropriate error will be raised	

--group-execution [continue-on-error | stop-on-error]

Specifies the group run policy when a firmware download process fails on one or more devices defined in a group. When all devices defined in the group have completed the firmware download process and there is one or more device that has failed, the group run policy rules if the firmware download work flow continues to process the remaining groups or not.

There are two options:

- Under the `continue-on-error` group run policy (the default), in case of a device upgrade failure, the work flow continues to the next group, until all groups have been processed.
- Under the `stop-on-error` group run policy, in case of a device upgrade failure, the work flow does not continue to the next group.

Usage Guidelines

You may receive a general warning for traffic loss for single-homed servers if any leaf devices or non-Clos devices are prepared for firmware download.

Run only one instance of this command per fabric.

```

efa inventory device firmware-download execute --noAutoCommit
--fabric fabric
Firmware Download Execute [success]
Monitor firmware download execution progress using:

efa inventory device firmware-download show --fabric fabric

Please don't execute other commands on these devices until firmware download
is in progress

Firmware download execution initiated with '--noAutoCommit' flag, please
commit/restore after completion using:

```



```
efa inventory device firmware-download commit --fabric fabric
efa inventory device firmware-download restore --fabric fabric
```

This example shows only a portion of all possible output.

```
efa inventory device firmware-download show --fabric fabric
+-----+-----+-----+-----+-----+-----+
+-----+
| IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.5 | NHF-Leaf1 | 3009 | SLX9150-48Y | 0 | Leaf |
20.2.2slxos20.2.2b_201208_0600 |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.6 | NHF-Leaf2 | 3009 | SLX9150-48Y | 0 | Leaf |
20.2.2slxos20.2.2b_201211_1800 |
+-----+-----+-----+-----+-----+-----+
+-----+
```

Firmware Download Show Details

Firmware Download Show Overall Status [Completed]

If Firmware download execution initiated with '--noAutoCommit' flag,
please commit/restore after completion using:

```
efa inventory device firmware-download commit --fabric fabric
efa inventory device firmware-download restore --fabric fabric
```

```
efa inventory device firmware-download execute --fabric fabric1
--noMaintMode
Firmware Download Execute [success]
Monitor firmware download execution progress using:
efa inventory device firmware-download show --fabric fabric1
```

Please don't execute other commands on these devices until firmware download
is in progress

```
--- Time Elapsed: 137.250437ms ---
```

efa inventory device firmware-download prepare add

Prepares a device for a firmware download. Firmware host sanity validations are done at this time. If the validations are successful, the device is prepared.

Syntax

```
efa inventory device firmware-download prepare add [ --fabric fabric name
--ip device ip address | [--group group number ] --force --firmware-
host firmware download host --firmware-directory path to target
firmware build ]
```

Command Default

Traffic loss expected for non-redundant devices (single non-MCT leaf, spine, or super spine).

Parameters

--fabric *fabric name*

Name of the fabric. Using the **--fabric** parameter automatically generates a group-based prepared list. The **--group** parameter is not permitted with the **--fabric** parameter.

--ip *device ip address*

Specifies a comma-separated list of device IP addresses.

- Using the **--ip** and optionally **--group** parameters adds the specified devices in the prepared list group. Group 1 is implicitly used if not specified with the **--ip** parameter.
- Using the **--ip** and optionally **--force** parameters bypasses all validations and forcefully prepare the device in the default group 1 or in any group number if the **--group** parameter is specified. This can be used if the user does not care about traffic loss and wants to perform the firmware download simultaneously on all devices. The firmware restart-ability should allow for both EFA hosted node's SLX firmware to be upgraded simultaneously.

--group *group number*

Group number. Using the **--ip** and optionally **--group** parameters adds the specified devices in the prepared list group. Group 1 is implicitly used if not specified with the **--ip** parameter.

--force

Bypasses all validations and forcefully prepares the device. The **--force** parameter can be used to download the same existing firmware version on the device. Normally, this is not allowed. However, using the **--force** parameter implicitly performs a full installation to carry out the same firmware version download.

--firmware-host *firmware download host*

Firmware download host IP address.

--firmware-directory *path to target firmware build*

Path to the target firmware build.

Usage Guidelines



Note

- The firmware host must be registered.
- Device IPs belong to the same fabric (for Clos topologies).
- Do not allow devices to be prepared if a firmware download is in progress.
- Allow a device to be prepared after a firmware download has completed or in an unprepared state.
- Do not allow both MCT leaf pairs to be prepared together.
- Do not allow all spines in the same pod of the same fabric to be prepared together.
- Do not allow all super-spines in the fabric to be prepared together.
- EFA performs a sanity check using the firmware-host and firmware-directory that are specified in the command to ensure that the switch takes the firmware from that host and directory.

Examples

This example prepares half of the devices in Group 1. Group 1 is default when a `--group` parameter is not provided.

```
$ efa inventory device firmware-download prepare add --ip
10.20.245.1,10.20.245.3,10.20.245.5,10.20.245.9 --firmware-host 10.20.241.101 --firmware-
directory /build/slx/slxos20.2.3a
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware | Firmware Host | Firmware Directory | Target Firmware | Last Update
Time |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.1 | L1 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.3 | L3 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.5 | L5 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.9 | S1 | 3012 | SLX9250-32C | 0 | Spine |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
```

```
Firmware Download Prepare Add Details
Validate Firmware Download Prepare Add [success]
```

This example prepares the remaining half of devices in Group 3.

```
$ efa inventory device firmware-download prepare add --ip
10.20.245.2,10.20.245.4,10.20.245.6,10.20.245.10 --group 3 --firmware-host 10.20.241.101
--firmware-directory /build/slx/slxos20.2.3a
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware | Firmware Host | Firmware Directory | Target Firmware | Last Update
Time |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.2 | L2 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.4 | L4 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.6 | L6 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.10 | S2 | 3012 | SLX9250-32C | 0 | Spine |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
Firmware Download Prepare Add Details
Validate Firmware Download Prepare Add [success]
```

This example lists all groups of devices that are prepared in the fabric.

```
$ efa inventory device firmware-download prepare list --fabric fab_3stg
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware | Firmware Host | Firmware Directory | Target Firmware | Last Update
Time |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.1 | L1 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.3 | L3 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.5 | L5 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.9 | S1 | | 3012 | SLX9250-32C | 0 | Spine |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:30 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.2 | L2 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.4 | L4 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.6 | L6 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.10 | S2 | | 3012 | SLX9250-32C | 0 | Spine |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:14:38 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
Firmware Download Prepare List Details
Validate Firmware Download Prepare List [success]

```

This example prepares the fabric and lets EFA automatically generate the group-based prepared list for the fabric.

```

$ efa inventory device firmware-download prepare add --fabric fab_3stg --firmware-host
10.20.241.101 --firmware-directory /build/slx/slxos20.2.3a
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware | Firmware Host | Firmware Directory | Target Firmware | Last Update
Time |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.1 | L1 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:21 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.3 | L3 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:21 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.5 | L5 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:21 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2 | 10.20.245.2 | L2 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:22 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2 | 10.20.245.4 | L4 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:22 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2 | 10.20.245.6 | L6 | | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:22 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3 | 10.20.245.9 | S1 | | 3012 | SLX9250-32C | 0 | Spine |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:23 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 4 | 10.20.245.10 | S2 | | 3012 | SLX9250-32C | 0 | Spine |
20.2.3 | 10.20.241.101 | /build/slx/slxos20.2.3a | 20.2.3a | 2021-03-19
20:36:24 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
Firmware Download Prepare Add Details
Validate Firmware Download Prepare Add [success]

```

efa inventory device firmware-download prepare list

Displays the current listing of prepared devices for a firmware download operation in the fabric.

Syntax

```
efa inventory device firmware-download prepare list [ --fabric fabric-name ]
```

Parameters

--fabric *fabric-name*
Specifies the name of the fabric.

Usage Guidelines

The command displays all prepared devices per group. The group column is the first column displayed in the output. The output is sorted in ascending order by group, then by IP address.

Examples

This example prepares the fabric and lets EFA automatically generate the group-based prepared list for the fabric. This example shows only a portion of all possible output.

```
$ efa inventory device firmware-download prepare list --fabric fab_3stg
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.1 | L1       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.3 | L3       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.5 | L5       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.9 | S1       | 3012 | SLX9250-32C | 0   | Spine |
20.2.3
+-----+-----+-----+-----+-----+-----+-----+-----+
| 3     | 10.20.245.2 | L2       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3
+-----+-----+-----+-----+-----+-----+-----+-----+

Firmware Download Prepare List Details
Validate Firmware Download Prepare List [success]
```

efa inventory device firmware-download prepare remove

Removes the specified devices from the prepared list.

Syntax

```
efa inventory device firmware-download prepare remove [ --fabric fabric-name | --ip ip-address | --group group-number ]
```

Parameters

--fabric *fabric-name*

Specifies the name of the fabric. Using this parameter completely deletes the prepared list for the fabric, including all groups.

--ip *ip-address*

Specifies a comma-separated list of device IP addresses. Using this parameter deletes the prepared devices, even if they are prepared in separate groups.

--group *group-number*

Specifies a group number. Using this parameter removes all devices prepared in the given group.

Usage Guidelines

The *fabric*, *group*, and *ip* parameters cannot be used together.

Examples

Firmware Download Prepare Add (Fabric-based)

Prepare the fabric and let EFA automatically generate the group-based prepared list for the fabric. This example shows only a portion of all possible output.

```
$ efa inventory device firmware-download prepare add --fabric fab_3stg
--firmware-host 10.20.241.101 --firmware-directory /build/slx/slxos20.2.3a
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.1 | L1       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.3 | L3       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1     | 10.20.245.5 | L5       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 2     | 10.20.245.2 | L2       | 3012 | SLX9250-32C | 0   | Leaf |
20.2.3  |
```



```

+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2      | 10.20.245.4 | L4      | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3  |
Firmware Download Prepare Add Details
Validate Firmware Download Prepare Add [success]

```

Firmware Download Prepare Delete (IP-based)

Given the prepared list from the previous example, delete an MCT leaf pair prepared across groups. This example shows only a portion of all possible output.

```

$ efa inventory device firmware-download prepare remove
--ip 10.20.245.5,10.20.245.6
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1      | 10.20.245.5 | L5      | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2      | 10.20.245.6 | L6      | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
Firmware Download Prepare Remove Details
Validate Firmware Download Prepare Remove [success]

```

Firmware Download Prepare Delete (Group-based)

Continuing from the previous example, delete devices prepared in Group 2. This example shows only a portion of all possible output.

```

$ efa inventory device firmware-download prepare remove --group 2
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2      | 10.20.245.2 | L2      | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2      | 10.20.245.4 | L4      | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3  |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
Firmware Download Prepare Remove Details
Validate Firmware Download Prepare Remove [success]

```

Firmware Download Prepare Delete (Fabric-based)

Continuing from the previous example, delete the remaining devices in the fabric. This example shows only a portion of all possible output.

```

$ efa inventory device firmware-download prepare remove --fabric fab_3stg
+-----+-----+-----+-----+-----+-----+-----+
+-----+

```

```
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current  
Firmware |  
+-----+-----+-----+-----+-----+-----+-----+-----  
+-----+  
| 1 | 10.20.245.1 | L1 | 3012 | SLX9250-32C | 0 | Leaf |  
20.2.3 |  
+-----+-----+-----+-----+-----+-----+-----+-----  
+-----+  
| 1 | 10.20.245.3 | L3 | 3012 | SLX9250-32C | 0 | Leaf |  
20.2.3 |  
+-----+-----+-----+-----+-----+-----+-----+-----  
+-----+  
| 3 | 10.20.245.9 | S1 | 3012 | SLX9250-32C | 0 | Spine |  
20.2.3 |  
+-----+-----+-----+-----+-----+-----+-----+-----  
+-----+  
| 4 | 10.20.245.10 | S2 | 3012 | SLX9250-32C | 0 | Spine |  
20.2.3 |  
+-----+-----+-----+-----+-----+-----+-----+-----  
+-----+  
Firmware Download Prepare Remove Details  
Validate Firmware Download Prepare Remove [success]
```

efa inventory device firmware-download restore

Restores the installed firmware, when the firmware download was started with the `NoAutoCommit` flag.

Syntax

```
efa inventory device firmware-download restore [ --fabric fabric-name | --ip ip-addr ]
```

Parameters

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *ip-addr*

Specifies a comma-separated list of device IP addresses.

The restore operation involves a reload of the device. Careful selection of which devices to restore together should be done to minimize traffic loss.

```
efa inventory device firmware-download restore --fabric fabric1
Firmware Download Restore for fabric fabric1 - [success]
+-----+-----+-----+
| DeviceIP    | Device FWDL Status                | Error |
+-----+-----+-----+
| 10.20.246.6 | Firmware Restored Successfully    |       |
+-----+-----+-----+
--- Time Elapsed: 1m25.765540754s ---
```

efa inventory device firmware-download show

Shows the progress and status of the running firmware download.

Syntax

```
efa inventory device firmware-download show [ --fabric fabric-name | --ip
device-ip-address ]
```

Parameters

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *device-ip-address*

Specifies a comma-separated list of device IP addresses.

Examples

This example shows a running firmware download where Group 1 has completed and Group 2 is in progress.

```
$ efa inventory device firmware-download show --fabric fab_3stg
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name | ASN | Role | Current
Firmware | Target Firmware | Update State | Status | Last Update
Time |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.1 | L1 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3a | 20.2.3a | Completed | Firmware Committed |
2021-03-19 21:18:51 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.3 | L3 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3a | 20.2.3a | Completed | Firmware Committed |
2021-03-19 21:21:42 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 1 | 10.20.245.5 | L5 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3a | 20.2.3a | Completed | Firmware Committed |
2021-03-19 21:19:21 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2 | 10.20.245.2 | L2 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 20.2.3a | In Progress | Firmware Download Started |
2021-03-19 21:22:56 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2 | 10.20.245.4 | L4 | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3 | 20.2.3a | In Progress | Firmware Download Started |
```

```

2021-03-19 21:23:27 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 2      | 10.20.245.6 | L6      | 3012 | SLX9250-32C | 0 | Leaf |
20.2.3   | 20.2.3a     | In Progress | Firmware Download Started |
2021-03-19 21:23:03 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 3      | 10.20.245.9 | S1      | 3012 | SLX9250-32C | 0 | Spine |
20.2.3   | 20.2.3a     | Prepared   |
2021-03-19 21:06:23 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 4      | 10.20.245.10 | S2      | 3012 | SLX9250-32C | 0 | Spine |
20.2.3   | 20.2.3a     | Prepared   |
2021-03-19 21:06:24 +0000 UTC |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
Firmware Download Show Details
Firmware Download Show Overall Status [In Progress]

```

efa inventory device health status

Displays device health status.

Syntax

```
efa inventory device health status [ --ip ]
```

Parameters

--ip

Specifies IP address of the device.

efa inventory device interface list

Displays the list of interfaces and details for the specified IP address, including the application state that indicates whether the device configuration is synchronized with EFA or has drifted (refreshed or deleted).

Syntax

```
efa inventory device interface list [--ip ip-addr | --type interface-type
  { physical | loopback | ve | po | all } | --admin-state state { up |
  down | all } | --oper-state oper-state { up | down | all } | --rme
  list ]
```

Parameters

--ip *ip-addr*

Specifies the IP address of the device for which you want to see a list of interfaces. For example: 10.1.1.13,10.1.1.50,10.1.1.101.

--type { **physical** | **loopback** | **ve** | **po** | **all** }

Specifies the type of interfaces on the devices. The default is `physical`.

--admin-state { **up** | **down** | **all** }

Specifies the administrative state of interfaces of the device. The default is `all`.

--oper-state { **up** | **down** | **all** }

Specifies the operational state (line protocol) of interfaces of the device. The default is `all`.

--rme *list*

Lists redundant management-enabled interface. The default is `all`.

Examples

The following example shows a portion of typical output for the command.

```
$ efa inventory device interface list --ip 10.20.62.145
+-----+-----+-----+-----+-----+
| DeviceIP | Name | Interface Type | Admin Status | Oper Status |
+-----+-----+-----+-----+-----+
| 10.20.62.145 | 0/1 | ethernet | up | up |
+-----+-----+-----+-----+-----+
| | 0/2 | ethernet | up | down |
+-----+-----+-----+-----+-----+
| | 0/3 | ethernet | up | down |
+-----+-----+-----+-----+-----+

Interface Details
--- Time Elapsed: 79.1761ms ---
```

efa inventory device interface list-breakout

Lists breakout ports, including the application state that indicates if the configuration on the device is in sync or has drifted (refreshed or deleted) with respect to EFA.

Syntax

```
efa inventory device interface list-breakout [ --ip device-ip ]
```

Parameters

--ip *device-ip*

Comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

Examples

The following example shows the list breakout for devices on a specific IP.

```
$ efa inventory device interface list-breakout --ip 10.20.246.18
+-----+-----+-----+
| IP Address | Name | AppState |
+-----+-----+-----+
| 10.20.246.18 | 0/52:1 | cfg-refreshed |
+-----+-----+-----+
| | 0/52:2 | cfg-refreshed |
+-----+-----+-----+
| | 0/52:3 | cfg-refreshed |
+-----+-----+-----+
| | 0/52:4 | cfg-refreshed |
+-----+-----+-----+
| | 0/53:1 | cfg-in-sync |
+-----+-----+-----+
| | 0/53:2 | cfg-in-sync |
+-----+-----+-----+
| | 0/53:3 | cfg-in-sync |
+-----+-----+-----+
| | 0/53:4 | cfg-in-sync |
+-----+-----+-----+
```


efa inventory device interface redundant-management

Configures redundant management Ethernet (RME) on the SLX interface.

Syntax

```
efa inventory device interface redundant-management [--ip device-ip | --  
if-name | --enable { true | false }]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--if-name

Specifies one interface name. A list of names is not supported.

--enable { **true** | **false** }

Enables redundant management.

Usage Guidelines

The configuration you set is persisted in the EFA database.

DRC and idempotency are supported.

The command is supported only on the 9150, 9250, and 9740 SLX platforms where this feature is supported.

Internally, EFA sets PPS (packets per second) to some value when RME is enabled. For SLX 9150 and 9250, PPS will be set to 8000.

On SLX 9740, BPS (bits per second) is set to 20000 Kbps (kilobits per second) when RME is enabled.

Examples

```
efa inventory device interface redundant-management --ip 10.20.246.10  
--if-name 0/17 --enable true
```

This example turns off RME:

```
efa inventory device interface redundant-management --ip 10.20.246.10  
--if-name 0/17 --enable false
```

efa inventory device interface set-admin-state

Brings an interface administratively up or down.

Syntax

```
efa inventory device interface set-admin-state [--ip device-ip | --if-  
type eth | --if-name if-list | --state { up | down } ]
```

Command Default

By default, all ports on the SLX device are down.

Parameters

--ip *device-ip*

Specifies a comma-separated list of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--if-type *eth*

Specifies that the interface type is Ethernet, which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. Example: 0/50-52,0/45,0/5:4,0/5:1-2.

--state { *up* | *down* }

State. Can be up or down.

Usage Guidelines

The configuration you set is persisted.

DRC and idempotency are supported.

In SLX-OS, you can use the **show interface ethernet** command to see the status of the Ethernet interfaces on your device.

Examples

This example changes the Admin Status on one IP address.

```
efa inventory device interface set-admin-state --ip 10.25.225.167  
--if-name 0/20-22 --state down
```

DeviceIP	ID	Name	Interface Type	Admin Status	Result	Reason
10.25.225.167	9	0/21	ethernet	down	Success	
	89	0/20	ethernet	down	Success	
	1	0/22	ethernet	down	Success	

```
Interface Details  
--- Time Elapsed: 19.5606145s ---
```

This example attempts to change the Admin Status on two IP addresses, however one IP address is invalid.

```
efa inventory device interface set-admin-state --ip 10.25.225.167,10.10.10.10
--if-name 0/20-22 --state down
```

DeviceIP	ID	Name	Interface Type	Admin Status	Result	Reason
10.10.10.10					Failed	Device does not exist with IP: 10.10.10.10
10.25.225.167	9	0/21	ethernet	down	Success	
	89	0/20	ethernet	down	Success	
	1	0/22	ethernet	down	Success	

Interface Details
--- Time Elapsed: 20.740265s ---

This example changes the Admin State on multiple valid IP addresses.

```
efa inventory device interface set-admin-state
--ip 10.25.225.167,10.24.48.131,10.24.51.135 --if-name 0/20-22 --state up
```

DeviceIP	ID	Name	Interface Type	Admin Status	Result	Reason
10.24.48.131	110	0/21	ethernet	up	Success	
	142	0/20	ethernet	up	Success	
	148	0/22	ethernet	up	Success	
10.24.51.135	16	0/21	ethernet	up	Success	
	48	0/22	ethernet	up	Success	
	86	0/20	ethernet	up	Success	
10.25.225.167	9	0/21	ethernet	up	Success	
	89	0/20	ethernet	up	Success	
	1	0/22	ethernet	up	Success	

Interface Details
--- Time Elapsed: 56.4964544s ---

This example changes the Admin Status for the specified fabric name.

```
efa inventory device interface set-admin-state --fabric nc_no_vni
--if-name 0/20-22 --state up
```

DeviceIP	ID	Name	Interface Type	Admin Status	Result	Reason
10.24.51.135	48	0/22	ethernet	up	Success	
	16	0/21	ethernet	up	Success	
	86	0/20	ethernet	up	Success	
10.24.48.131	148	0/22	ethernet	up	Success	

```
|          | 142 | 0/20 | ethernet      | up          | Success |          |
+-----+-----+-----+-----+-----+-----+-----+
|          | 110 | 0/21 | ethernet      | up          | Success |          |
+-----+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 37.9236481s ---
```

efa inventory device interface set-breakout

Breaks a port into multiple interfaces, such as breaking one 40G port into four 10G ports.

Syntax

```
efa inventory device interface set-breakout [--ip device-ip | --if-type
eth | --if-name if-list | --mode { 1x10g | 1x25g | 1x40g | 1x100g |
2x40g | 2x50g | 4x10g | 4x25g } ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated list of device IP addresses. For example: 10.1.1.13,10.1.1.50,10.1.1.101.

--if-type *eth*

Specifies that the interface type is Ethernet, which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. For example: 0/50-52,0/45.

--mode { 1x10g | 1x25g | 1x40g | 1x100g | 2x40g | 2x50g | 4x10g | 4x25g }

Configures breakout mode for the ports.

Usage Guidelines

In SLX-OS, you can use the **show running-config hardware** command to see whether breakout mode is configured for a device.

The new breakout interfaces you create are identified by the name of the original interface followed by a suffix.

The configuration you set is persisted.

DRC and idempotency are supported.

Examples

This example breaks three interfaces into four ports each.

```
efa inventory device interface set-breakout --ip 10.24.80.158
+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | Result |
+-----+-----+-----+-----+
| 10.24.80.158 | 73 | 0/2:2 | ethernet | Success |
+-----+-----+-----+-----+
| | 72 | 0/1:4 | ethernet | Success |
+-----+-----+-----+-----+
| | 74 | 0/3:2 | ethernet | Success |
+-----+-----+-----+-----+
| | 78 | 0/3:3 | ethernet | Success |
+-----+-----+-----+-----+
| | 75 | 0/3:4 | ethernet | Success |
```

```

+-----+-----+-----+-----+
| 70 | 0/1:1 | ethernet | Success |
+-----+-----+-----+-----+
| 71 | 0/1:3 | ethernet | Success |
+-----+-----+-----+-----+
| 80 | 0/2:1 | ethernet | Success |
+-----+-----+-----+-----+
| 79 | 0/1:2 | ethernet | Success |
+-----+-----+-----+-----+
| 76 | 0/2:3 | ethernet | Success |
+-----+-----+-----+-----+
| 69 | 0/3:1 | ethernet | Success |
+-----+-----+-----+-----+
| 77 | 0/2:4 | ethernet | Success |
+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 48.3801684s ---

```

This example attempts to break out ports on two devices, but one IP address is invalid.

```

efa inventory device interface set-breakout
--ip 10.24.80.158,10.10.10
+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | Result | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.10.10.10 | | | | Failed | Device does not exist with IP: 10.10.10.10 |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.158 | 94 | 0/8:3 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 91 | 0/9:3 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 83 | 0/9:4 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 86 | 0/6:3 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 59.3971958s ---

```

This example breaks out ports on multiple devices.

```

efa inventory device interface set-breakout
--ip 10.24.80.158,10.24.80
+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | Result | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.158 | 0 | 0/9 | ethernet | Failed | Interface [0/9] from the range [0/9-12] does not exist or port is already breakout for device IP 10.24.80.158. Specify a valid interface. |
+-----+-----+-----+-----+-----+-----+
| | 105 | 0/10:2 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 108 | 0/12:2 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 102 | 0/12:3 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 104 | 0/11:4 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 97 | 0/12:4 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+

```

```

|          | 106 | 0/10:4 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+-----+
|          | 98  | 0/11:2 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 100 | 0/11:3 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
| 10.24.80.159 | 120 | 0/10:1 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 114 | 0/9:4  | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 118 | 0/10:3 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 115 | 0/12:2 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 1m16.6280022s ---

```

This example configures break out for the specified fabric name.

```

efa inventory device interface set-breakout --if-name 0/19-20
--mode 4x25g --fabric fabric1
+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID  | Name | Interface Type | Result | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.158 | 188 | 0/20:2 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 184 | 0/20:3 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 190 | 0/20:4 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 191 | 0/19:1 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 187 | 0/19:2 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 189 | 0/19:3 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 185 | 0/19:4 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 186 | 0/20:1 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
| 10.24.80.159 | 196 | 0/19:2 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 194 | 0/19:3 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 197 | 0/19:4 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 199 | 0/20:1 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 193 | 0/20:2 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 195 | 0/20:3 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 192 | 0/20:4 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
|          | 198 | 0/19:1 | ethernet      | Success |          |
+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 1m6.2210288s ---

```

efa inventory device interface set-fec

Configures forwarding error correction (FEC) on the SLX interface.

```
efa inventory device interface set-fec [--ip device-ip | --if-type eth |  
--if-name if-list | --mode { fs-fec | rs-fec | auto | disabled }]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--if-type **eth**

Specifies that the interface type is **eth** (Ethernet), which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. For example: 0/50-52,0/45,0/5:4,0/5:1-2.

--mode { **fs-fec** | **rs-fec** | **auto** | **disabled** }

Specifies the FEC mode for the ports.

Usage Guidelines

The configuration you set is persisted in the EFA database.

The default value of FEC configured by SLX is **auto**, that is auto-negotiation.

DRC and idempotency are supported.

```
efa inventory device interface set-fec --ip 10.20.246.10  
--if-type ethernet --if-name  
0/20 -mode rs-fec
```


efa inventory device interface set-link-error-disable

Configures port dampening (link-error-disable) on the SLX interface. Minimizes excessive interface flapping.

Syntax

```
efa inventory device interface set-link-error-disable [--ip device-ip |  
--if-type eth | --if-name if-list | --toggle-threshold value | --  
sampling-time value | --wait-time value ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--if-type **eth**

Specifies that the interface type is Ethernet, which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. For example: 0/50-52,0/45,0/5:4,0/5:1-2.

--toggle-threshold *value*

Valid values are 1 through 50.

--sampling-time *value*

Valid values in seconds, are 1 through 65565. The default is 1.

--wait-time *value*

Valid values, in seconds, are 0 through 65565.

Usage Guidelines

The configuration you set is persisted in the EFA database.

DRC and idempotency are supported.

Examples

```
efa inventory device interface set-link-error-disable --ip 10.20.246.10  
--if-type eth --if-name 0/20 --toggle-threshold 10 --sampling-time 20  
--wait-time 10
```

efa inventory device interface set-mtu

Configures the MTU (maximum transmission unit) at the physical port level for Layer 2 and IPv4.

Syntax

```
efa inventory device interface set-mtu [--ip device-ip | --if-type eth |
  --if-name if-list | --mtu int mtu-value | --ip-mtu int ip-mtu-value |
  --ipv6-mtu int32 ipv6-mtu-value ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--if-type **eth**

Specifies that the interface type is Ethernet, which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. For example: 0/50-52,0/45.

--mtu int *mtu-value*

Specifies the global Layer 2 MTU on a device. For SLX-OS 21.1 and later, valid values range from 1500 through 9216. For other SLX-OS versions, valid values range from 1548 through 9216.

--ip-mtu int *ip-mtu-value*

Specifies the global IP MTU value on a device. For SLX IPv4, valid values range from 1300 through 9194.

--ipv6-mtu *ipv6-mtu-value*

Specifies the global IPv6 MTU on a device. For SLX IPv4 or IPv6, valid values range from 1300 through 9194.

Usage Guidelines

In SLX-OS, you can use the **show interface ethernet** command to see the MTU configuration for an interface.

The configuration you set is persisted.

DRC and idempotency are supported.

Examples

This example configures the MTU on one device.

```
efa inventory device interface set-mtu --ip 10.25.225.167
--if-name 0/20-22
--mtu 2000 --ip-mtu 2000 --ipv6-mtu 3000
+-----+-----+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | MTU | IP MTU | IPv6 MTU | Result |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

```

| 10.25.225.167 | 1 | 0/22 | ethernet | 2000 | 2000 | 3000 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 9 | 0/21 | ethernet | 2000 | 2000 | 3000 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 89 | 0/20 | ethernet | 2000 | 2000 | 3000 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
Interface MTU Details
--- Time Elapsed: 18.6886395s ---

```

This example attempts to configure the MTU for two IP addresses, however one IP address is invalid.

```

efa inventory device interface set-mtu --ip 10.25.225.167,10.10.10.10
--if-name 0/20-22 --mtu 2500 --ip-mtu 2500 --ipv6-mtu 2500
+-----+-----+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | MTU | IP MTU | IPv6 MTU | Result
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.10.10.10 | | | | | Failed | Device does not
| | | | | | | exist with IP:
| | | | | | | 10.10.10.10
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.25.225.167 | 89 | 0/20 | ethernet | 2500 | 2500 | 2500 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 1 | 0/22 | ethernet | 2500 | 2500 | 2500 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 9 | 0/21 | ethernet | 2500 | 2500 | 2500 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
Interface MTU Details
--- Time Elapsed: 18.6252821s ---

```

This example configures the MTU on multiple IP addresses.

```

efa inventory device interface set-mtu --ip 10.25.225.167,10.24.48.131,10.24.51.135
--if-name 0/20-22 --mtu 3600 --ip-mtu 3600 --ipv6-mtu 3600
+-----+-----+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | MTU | IP MTU | IPv6 MTU | Result
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.25.225.167 | 9 | 0/21 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 89 | 0/20 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 1 | 0/22 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.24.48.131 | 142 | 0/20 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 148 | 0/22 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 110 | 0/21 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.24.51.135 | 16 | 0/21 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 48 | 0/22 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 86 | 0/20 | ethernet | 3600 | 3600 | 3600 | Success
+-----+-----+-----+-----+-----+-----+-----+-----+
Interface MTU Details
--- Time Elapsed: 59.5462548s ---

```

This example configures the MTU for the specified fabric name.

```

efa inventory device interface set-mtu --fabric nc_no_vni
--if-name 0/20-22,0/55-58 --mtu 3200 --ip-mtu 3200 --ipv6-mtu 3200
+-----+-----+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | MTU | IP MTU | IPv6 MTU | Result |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

```
| 10.24.51.135 | 16 | 0/21 | ethernet | 3200 | 3200 | 3200 | Success |
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 48 | 0/22 | ethernet | 3200 | 3200 | 3200 | Success |
|
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 86 | 0/20 | ethernet | 3200 | 3200 | 3200 | Success |
|
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.24.48.131 | 110 | 0/21 | ethernet | 3200 | 3200 | 3200 | Success |
|
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 148 | 0/22 | ethernet | 3200 | 3200 | 3200 | Success |
|
+-----+-----+-----+-----+-----+-----+-----+-----+
| | 142 | 0/20 | ethernet | 3200 | 3200 | 3200 | Success |
|
+-----+-----+-----+-----+-----+-----+-----+-----+
Interface MTU Details
--- Time Elapsed: 37.3021602s ---
```

efa inventory device interface set-speed

Configures the speed for receiving and transmitting data on a physical port.

Syntax

```
efa inventory device interface set-speed [--ip device-ip | --if-type eth
  | --if-name if-list | --speed { 100mbps | 1gbps | 10gbps | 25gbps |
  40gbps | 100gbps | 1GbpsAN | AUTO }]
```

Parameters

--ip *device-ip*

Specifies a comma-separated list of device IP addresses.

--if-type **eth**

Specifies that the interface type is Ethernet, which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. For example: 0/50-52,0/45.

--speed { **100mbps** | **1gbps** | **10gbps** | **25gbps** | **40gbps** | **100gbps** | **1GbpsAN** | **AUTO** }

Specifies the speed for the port.

Usage Guidelines

In SLX-OS, you can use the **show interface ethernet** command to see the speed of the Ethernet interfaces on your device.

The configuration you set is persisted.

DRC and idempotency are supported.

Examples

This example sets the port speed on one device.

```
efa inventory device interface set-speed --ip 10.25.225.167
--if-name 0/20-22 --speed 10gbps
+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | Port Speed | Result | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.25.225.167 | 89 | 0/20 | ethernet | 10gbps | Success | |
+-----+-----+-----+-----+-----+-----+
| | 1 | 0/22 | ethernet | 10gbps | Success | |
+-----+-----+-----+-----+-----+-----+
| | 9 | 0/21 | ethernet | 10gbps | Success | |
+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 17.9586487s ---
```

This example attempts to set the port speed on two IP addresses, however one IP address is invalid.

```
efa inventory device interface set-speed --ip 10.25.225.167,10.10.10.10
--if-name 0/20-22 --speed 25gbps
```

DeviceIP	ID	Name	Interface Type	Port Speed	Result	Reason
10.10.10.10					Failed	Device does not exist with IP: 10.10.10.10
10.25.225.167	89	0/20	ethernet	25gbps	Success	
	1	0/22	ethernet	25gbps	Success	
	9	0/21	ethernet	25gbps	Success	

```
Interface Details
--- Time Elapsed: 18.3503896s ---
```

This example sets the port speed on multiple valid IP addresses.

```
efa inventory device interface set-speed --ip 10.25.225.167,10.24.48.131,10.24.51.135
--if-name 0/20-22 --speed 25gbps
```

DeviceIP	ID	Name	Interface Type	Port Speed	Result	Reason
10.25.225.167	9	0/21	ethernet	25gbps	Success	
	89	0/20	ethernet	25gbps	Success	
	1	0/22	ethernet	25gbps	Success	
10.24.51.135	16	0/21	ethernet	25gbps	Success	
	86	0/20	ethernet	25gbps	Success	
	48	0/22	ethernet	25gbps	Success	
10.24.48.131	142	0/20	ethernet	25gbps	Success	
	110	0/21	ethernet	25gbps	Success	
	148	0/22	ethernet	25gbps	Success	

```
Interface Details
--- Time Elapsed: 53.8631425s ---
```

This example sets the port speed for the specified fabric name.

```
efa inventory device interface set-speed --fabric nc_no_vni
--if-name 0/20-22 --speed 25gbps
```

DeviceIP	ID	Name	Interface Type	Port Speed	Result	Reason
10.24.51.135	86	0/20	ethernet	25gbps	Success	
	48	0/22	ethernet	25gbps	Success	
	16	0/21	ethernet	25gbps	Success	
10.24.48.131	142	0/20	ethernet	25gbps	Success	
	110	0/21	ethernet	25gbps	Success	

```
+-----+-----+-----+-----+-----+
|          | 148 | 0/22 | ethernet      | 25gbps      | Success |
+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 36.9974805s ---
```

efa inventory device interface unset-breakout

Reverts the breakout of multiple ports back to the original configuration.

Syntax

```
efa inventory device interface unset-breakout [--ip device-ip | --if-type
eth | --if-name if-list ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated list of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--if-type *eth*

Specifies that the interface type is Ethernet, which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. Example: 0/50-52,0/45.

Usage Guidelines

In SLX-OS, you can use the **show running-config hardware** command to see whether breakout mode is configured for a device.

When you run this command, the breakout interfaces are deconfigured and deleted. The original Ethernet interface in the default configuration is created automatically.

The configuration you set is persisted.

DRC and idempotency are supported.

Examples

This example removes breakout mode on one device.

```
efa inventory device interface unset-breakout --ip 10.24.80.158
--if-name 0/1-3
+-----+-----+-----+-----+
| DeviceIP | Interface ID | Interface Name | Interface Type | Result |
+-----+-----+-----+-----+
| 10.24.80.158 | 171 | 0/2 | ethernet | Success |
+-----+-----+-----+-----+
| | 175 | 0/3 | ethernet | Success |
+-----+-----+-----+-----+
| | 182 | 0/1 | ethernet | Success |
+-----+-----+-----+-----+
Interface Details
```

This example attempts to remove breakout mode on two devices, however one IP address is invalid.

```
efa inventory device interface unset-breakout --ip 10.24.80.158,10.10.10.10
--if-name 0/6-9
+-----+-----+-----+-----+-----+
| DeviceIP | Interface ID | Interface Name | Interface Type | Result | Reason |
+-----+-----+-----+-----+-----+
| 10.24.80.158 | 171 | 0/6 | ethernet | Success | |
+-----+-----+-----+-----+-----+
| 10.10.10.10 | 172 | 0/7 | ethernet | Error | Invalid IP |
+-----+-----+-----+-----+-----+
```



```

+-----+-----+-----+-----+-----+-----+
| 10.10.10.10 | | | | | Failed | Device does
| | | | | | | not exist
| | | | | | | with IP:
| | | | | | | 10.10.10.10
+-----+-----+-----+-----+-----+-----+
| 10.24.80.158 | 229 | 0/6 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 227 | 0/7 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 228 | 0/8 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 230 | 0/9 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 1m6.5402944s ---

```

This example removes breakout mode on multiple devices.

```

efa inventory device interface unset-breakout
--ip 10.24.80.158,10.24.80.159 --if-name 0/9-12
+-----+-----+-----+-----+-----+-----+
| DeviceIP | Interface ID | Interface Name | Interface Type | Result |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.158 | 248 | 0/10 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 250 | 0/11 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 249 | 0/12 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 247 | 0/9 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.159 | 252 | 0/10 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 254 | 0/11 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 253 | 0/12 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 251 | 0/9 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 1m52.8562333s ---

```

This example removes breakout mode for the specified fabric name.

```

efa inventory device interface unset-breakout --if-name 0/19-20
--fabric fabric1
+-----+-----+-----+-----+-----+-----+
| DeviceIP | Interface ID | Interface Name | Interface Type | Result |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.159 | 279 | 0/19 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 280 | 0/20 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.158 | 281 | 0/19 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
| | 282 | 0/20 | ethernet | Success |
+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 1m19.226463s ---

```

efa inventory device interface unset-fec

Unsets port forwarding error correction (FEC) on the SLX interface.

```
efa inventory device interface unset-fec [--ip device-ip | --if-type eth
| --if-name if-list ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--if-type **eth**

Specifies that the interface type is `eth` (Ethernet), which is the default setting.

--if-name *if-list*

Specifies a comma-separated list of interface names. Example: 0/50-52,0/45,0/5:4,0/5:1-2.

Usage Guidelines

The configuration you set is persisted in the EFA database.

The default value of FEC configured by SLX is `auto`, that is auto-negotiation.

DRC and idempotency are supported.

```
efa inventory device interface unset-fec --ip 1.1.1.1 --if-name 0/10
+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | Result | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.1.1.1 | 3 | 0/10 | ethernet | Success | |
+-----+-----+-----+-----+-----+-----+
Interface Details
--- Time Elapsed: 5.263239407s ---
```

efa inventory device interface unset-link-error-disable

Unsets port dampening (link-error-disable) on the SLX interface.

Syntax

```
efa inventory device interface unset-link-error-disable [--ip device-ip  
| --if-name if-list ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--if-name *if-list*

Specifies a comma-separated list of interface names. Example: 0/50-52,0/45,0/5:4,0/5:1-2.

Usage Guidelines

The configuration you set is persisted in the EFA database.

DRC and idempotency are supported.

Examples

```
efa inventory device interface unset-link-error-disable --ip 10.20.246.10 --if-type eth --  
if-name 0/20
```

efa inventory device list

Lists all devices known to the inventory.

Syntax

```
efa inventory device list [ --orphan | --fabric name | --role { leaf |
border-leaf | spine | super-spine } | --ips device-ips ]
```

Parameters

--orphan

Fetches all devices that are not associated to a fabric.

--fabric *name*

Specifies fabric name. If this modifier is used, the command shows only those devices in the specified fabric.

--role { leaf | border-leaf | spine | super-spine }

Specifies device role.

--ips *device-ips*

Specifies comma separated device IPs.

Examples

These examples show a portion of all possible output.

```
$ efa inventory device list
```

```
+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model | Chassis Name | Firmware |
+-----+-----+-----+-----+-----+
| 10.139.44.156 | SLX      | 4000 | BR-SLX9540 | 20.2.2slxos20.2.2_200914_1800 |
+-----+-----+-----+-----+-----+
| 10.139.44.155 | SLX      | 4000 | BR-SLX9540 | 20.2.2slxos20.2.2_200914_1800 |
+-----+-----+-----+-----+-----+
```

```
$ efa inventory device register --ip 10.139.44.154
--username=admin --password=password
```

```
+-----+-----+-----+-----+-----+
| ID | IP Address | Host Name | Model | Chassis Name | Firmware |
+-----+-----+-----+-----+-----+
| 5 | 10.139.44.154 | SLX      | 4000 | BR-SLX9540 | 20.2.2slxos20.2.2_200914_1800 |
+-----+-----+-----+-----+-----+
```

```
$ efa inventory device list
```

```
+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model | Chassis Name | Firmware |
+-----+-----+-----+-----+-----+
| 10.139.44.156 | SLX      | 4000 | BR-SLX9540 | 20.2.2slxos20.2.2_200914_1800 |
+-----+-----+-----+-----+-----+
| 10.139.44.155 | SLX      | 4000 | BR-SLX9540 | 20.2.2slxos20.2.2_200914_1800 |
+-----+-----+-----+-----+-----+
| 10.139.44.154 | SLX      | 4000 | BR-SLX9540 | 20.2.2slxos20.2.2_200914_1800 |
+-----+-----+-----+-----+-----+
```

efa inventory device lldp list

Lists the LLDP (Link Layer Discovery Protocol) neighbors for a device.

```
efa inventory device lldp list [--ip device-ip | --type { all | edge | fabric } ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--type { **all** | **edge** | **fabric** }

Specifies the type of neighbors to show in the list. The default is **all**.

This example returns LLDP neighbors of type **fabric**. This example shows a portion of all possible output.

```
efa inventory device lldp list --ip 10.20.246.1 --type fabric
+-----+-----+-----+-----+-----+
| DeviceIP | Local Interface | Local Int MAC | Remote Interface | Remote Interface MAC |
+-----+-----+-----+-----+-----+
| 10.20.246.1 | Ethernet 0/13 | f46e.95a2.b813 | Ethernet 0/13 | f46e.95a0.c813 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/14 | f46e.95a2.b814 | Ethernet 0/14 | f46e.95a0.c815 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/15 | f46e.95a2.b815 | Ethernet 0/15 | f46e.95a0.c815 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/31 | f46e.95a2.b825 | Ethernet 0/3 | 489b.d57f.f009 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/32 | f46e.95a2.b826 | Ethernet 0/3 | 489b.d57f.6809 |
+-----+-----+-----+-----+-----+
```

This example returns LLDP neighbors of type **edge**. This example shows a portion of all possible output.

```
efa inventory device lldp list --ip 10.20.246.1 --type edge
+-----+-----+-----+-----+-----+
| DeviceIP | Local Interface | Local Int MAC | Remote Interface | Remote Interface MAC |
+-----+-----+-----+-----+-----+
| 10.20.246.1 | Ethernet 0/1 | f46e.95a2.b807 | Ethernet 0/1 | 609c.9f87.0a05 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/2 | f46e.95a2.b808 | Ethernet 0/2 | 609c.9f87.0a06 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/4 | f46e.95a2.b80a | Ethernet 0/4 | 609c.9f87.0c08 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/5 | f46e.95a2.b80b | Ethernet 0/5 | 609c.9fcd.f309 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/24 | f46e.95a2.b81e | Ethernet 0/53 | f46e.959f.1744 |
+-----+-----+-----+-----+-----+
```

This example returns all LLDP neighbors. This example shows a portion of all possible output.

```
efa inventory device lldp list --ip 10.20.246.1 --type all
+-----+-----+-----+-----+-----+
| DeviceIP | Local Interface | Local Int MAC | Remote Interface | Remote Interface MAC |
+-----+-----+-----+-----+-----+
| 10.20.246.1 | Ethernet 0/1 | f46e.95a2.b807 | Ethernet 0/1 | 609c.9f87.0a05 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/2 | f46e.95a2.b808 | Ethernet 0/2 | 609c.9f87.0a06 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/4 | f46e.95a2.b80a | Ethernet 0/4 | 609c.9f87.0c08 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/5 | f46e.95a2.b80b | Ethernet 0/5 | 609c.9fcd.f309 |
+-----+-----+-----+-----+-----+
| | Ethernet 0/6 | f46e.95a2.b80c | Ethernet 0/6 | 609c.9fcd.f30a |
+-----+-----+-----+-----+-----+
```

efa inventory device ntp disable-server

Lists NTP servers configured using EFA.

```
efa inventory device ntp disable-server [ --ip device-ip | --enable { yes | no } | --list | --fabric name ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--enable { **yes** | **no** }

Turns the NTP server on or off.

--list

Lists disable-server on devices.

--fabric *name*

Fabric name.

To turn off an NTP server on a given device:

```
efa inventory device ntp disable-server --enable yes --ip 10.20.246.10
```

To turn on an NTP server on a given device:

```
efa inventory device ntp disable-server --enable no --ip 10.20.246.10
```

To turn off an NTP server at the fabric level:

```
efa inventory device ntp disable-server --enable yes --fabric clos_fabric
```

To turn on an NTP server at the fabric level:

```
efa inventory device ntp disable-server --enable no --fabric clos_fabric
```

efa inventory device ntp server create

Creates an NTP server.

Syntax

```
efa inventory device ntp server create [ --ip device-ip | --ntp-ip ntp-ip  
  | --auth-key key | --auth-key-name key-name | --encryption-type { md5  
  | sha1 } | --trusted-key | --fabric name ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example: 10.1.1.13,10.1.1.50,10.1.1.101. Required if you do not specify **--fabric**.

--ntp-ip *ntp-ip*

Specifies the IP address of the NTP server.

--auth-key *key*

Specifies the authentication key ID. Values can be from 1 to 65535.

--auth-key-name *key-name*

Specifies the name of the key.

--encryption-type { **md5** | **sha1** }

Specifies the encryption type. .

--trusted-key *key*

Specifies the trusted key.

--fabric *name*

Specifies the name of the fabric. Required if you do not specify **--ip**.

Usage Guidelines

The command creates an NTP server. By default, EFA sets the key name to be encrypted 7. SLX supports encryption levels of 0 (clear text) and 7 (encrypted).

Here, one **auth-key** ID can be used by different NTP servers in SLX. To simplify, EFA enforces a unique key id per NTP server. Here, device IP addresses and fabric options are mutually exclusive.

Examples

```
efa inventory device ntp server create -ntp-ip 3.3.3.3  
--auth-key 1 --auth-key-name ntpsecret --encryption-type md5 -trusted-key  
--ip 10.20.246.10
```

```
efa inventory device ntp server create -ntp-ip 3.3.3.3 --auth-key 1  
--auth-key-name ntpsecret --encryption-type md5 -trusted-key --fabric clos_fabric
```

efa inventory device ntp server delete

Deletes an NTP server.

```
efa inventory device ntp server delete [ --ip device-ip | --ntp-ip ntp-ip  
  | --fabric name ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example: 1.1.1.1-3,1.1.1.2,2.2.2.2.

--ntp-ip *ntp-ip*

Specifies the IP address of the NTP server.

--fabric *name*

Specifies the name of the fabric.

```
efa inventory device ntp server delete --ntp-ip 3.3.3.3  
--ip 10.20.246.10
```

```
efa inventory device ntp server delete  
--ntp-ip 3.3.3.3 --fabric clos_fabric
```


efa inventory device ntp server list

Lists NTP servers configured using EFA.

Syntax

```
efa inventory device ntp server list [ --ip device-ip | --fabric name ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *name*

Specifies the name of the fabric.

Examples

```
efa inventory device ntp server list --ip 10.20.246.10
```

```
efa inventory device ntp server list --fabric clos_fabric
```

efa inventory device running-config persist

Saves the running-config to the startup-config on SLX devices.

Syntax

```
efa inventory device running-config persist [ --ip ip-addr | --fabric fabric-name ]
```

Parameters

--ip *ip-addr*

Specifies a comma-separated list of IP address and IP address ranges of the devices for which you want the configuration to persist.

--fabric *fabric-name*

Specifies the fabric that contains the devices for which you want the configuration to persist.

Usage Guidelines

Device configurations configured by the Fabric Service and Tenant Service do not automatically persist on SLX devices.

Use this command to specify the IP addresses or fabric for which configurations should persist.

Examples

This example configures persistence for two IP addresses.

```
# efa inventory device running-config persist --ip 10.20.50.212,10.20.50.213

Persist Devices Running-Config (success)
IP Address      Device Name      Fabric           Status
10.20.50.212   Leaf-1-3         stage5          Success
10.20.50.213   Leaf-1-3         stage5          Success

Persist Running-Config Details
--- Time Elapsed: 12.2902836s ---
```

This example configures persistence for the same devices, but by fabric instead of IP address.

```
# efa inventory device running-config persist --fabric stage5

Persist Devices Running-Config (success)
IP Address      Device Name      Fabric           Status
10.20.50.212   Leaf-1-3         stage5          Success
10.20.50.213   Leaf-1-3         stage5          Success

Persist Running-Config Details
--- Time Elapsed: 11.4899986s ---
```

efa inventory device register

Registers a device.

Syntax

```
efa inventory device register [ --ip ip-list | --username username | --password password | --maintmode-enable-on-reboot ]
```

Parameters

--ip *ip-ist*

Specifies a comma separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--username *username*

Specifies the user name to connect to the device.

--password *password*

Specifies password to connect to the device.

--maintmode-enable-on-reboot

Configures a device to reboot in maintenance mode.

Usage Guidelines

If, during device registration, the device is found to be in maintenance mode, then an error is returned and registration fails. To take the device out of maintenance mode before reattempting device registration, you must perform a manual procedure.

efa inventory device reload

Reloads a running SLX device.

Syntax

```
efa inventory device reload [ --ip device-ips | --fabric fabric-name | -h, --help ]
```

Parameters

--ip *device-ips*

Specifies comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Specifies the name of the fabric.

-h, **--help**

Help for device reload.

Usage Guidelines

This operation causes the specified devices to reboot. When **--fabric** is specified, all devices in the fabric are reloaded. All existing sessions on the devices must be restarted.

Examples

This example reloads a device with IP address 10.20.30.40.

```
efa inventory device reload --ip 10.20.30.40

Warning: This operation will cause the device(s) to reboot and requires all existing
telnet, secure telnets and SSH sessions to be restarted on the device(s).

Warning: EFA access may be lost temporary if EFA is running on reloading device.

Are you sure you want to reload device(s) 10.20.30.40 (Y/N):
```

efa inventory device setting show

Displays device settings.

Syntax

```
efa inventory device setting show [ --ip device-ip ]
```

Parameters

--ip *device-ip*

Specifies IP address of the device.

This example shows the current inventory settings for the specified device.

```
efa inventory device setting show --ip 10.20.24.10
+-----+-----+
|           NAME           | VALUE |
+-----+-----+
| Maintenance Mode Enable | No    |
| Reboot                   |      |
+-----+-----+
| Maintenance Mode Enable | No    |
+-----+-----+
| Maintenance Convergence |      |
| Time                     |      |
+-----+-----+
| Maintenance Convergence |      |
| Recommended Time       |      |
+-----+-----+
| Health Check Enabled    | No    |
+-----+-----+
| Health Check Interval   | 6m    |
+-----+-----+
| Health Check Heartbeat  | 2     |
| Miss Threshold         |      |
+-----+-----+
| Periodic Backup Enabled | Yes   |
+-----+-----+
| Config Backup Interval  | 24h   |
+-----+-----+
| Config Backup Count     | 4     |
+-----+-----+
| Prefix Independent      | No    |
| Convergence             |      |
+-----+-----+
| Maximum Load Sharing   | 64    |
| Paths                   |      |
+-----+-----+
```

efa inventory device setting update

Configures device settings such as maintenance mode.

Syntax

```
efa inventory device setting update [ --ip device-ips | --fabric fabric-name | --maint-mode-enable-on-reboot { Yes | No } --maint-mode-enable { Yes | No } --health-check-enable { Yes | No } --health-check-interval minutes | --health-check-heartbeat-miss-threshold { 2 | 3 | 4 | 5 } --config-backup-periodic-enable { Yes | No } --config-backup-interval interval | --number-of-config-backups count | --mct-bring-up-delay delay | --maint-mode-convergence-time time | --prefix-independent-convergence { Yes | No } | --maximum-load-sharing-paths { 8 | 16 | 32 | 64 } | -h, --help ]
```

Parameters

--ip *device-ips*

Specifies a comma-separated range of device IP addresses. For example:

```
10.1.1.13,10.1.1.50,10.1.1.101.
```

--fabric *fabric-name*

Specify the name of the fabric.

--maint-mode-enable-on-reboot { **Yes** | **No** }

Specify Yes to activate maintenance mode on reboot and No to deactivate it.

--maint-mode-enable { **Yes** | **No** }

Specify Yes to activate maintenance mode and No to deactivate it.

--health-check-enable { **Yes** | **No** }

Specify Yes to enable health check and No to disable it.

--health-check-interval *minutes*

Specifies the health check interval in minutes. Valid values are 6m through 24h.

--health-check-heartbeat-miss-threshold { **2** | **3** | **4** | **5** }

Specifies the health check miss threshold.

--config-backup-periodic-enable { **Yes** | **No** }

Specify Yes to enable periodic configuration backup and No to disable it.

--config-backup-interval *interval*

Specifies configuration backup interval, in minutes. Valid values are 3m through 30h. The default is 24h.

--number-of-config-backups *count*

Specifies configuration backup count. Valid values are 2-20. The default is 4.

--mct-bring-up-delay *delay*

Specifies delay, in seconds, before an MCT cluster bring up. Valid values are 10-600. Default is 90. Set a value of 0 to deconfigure.

--maint-mode-convergence-time *time*

Specifies the maximum time, in seconds, that maintenance mode is allowed to complete operations. Valid values are 100-500. Default is 300. Set a value of 0 to deconfigure.

--prefix-independent-convergence { **Yes** | **No** }

Specify Yes to enable BGP PIC and No to de-configure it.

--maximum-load-sharing-paths *string*

Specifies the route load-sharing maximum paths. Valid values are 8,16,32,64. The default is 64 paths.

-h, --help

Help for update.

Examples

This example specifies a delay of 200 seconds before an MCT cluster bring up.

```
$ efa inventory device setting update --ip 10.20.24.10
--mct-bring-up-delay 200
```

This example specifies a maximum time of 100 seconds that maintenance mode is allowed to complete operations.

```
$ efa inventory device setting update --ip 10.20.24.10
--maint-mode-convergence-time 100
```

This example enables prefix independent convergence on the device.

```
efa inventory device setting update --ip 10.20.24.10 --prefix-independent-convergence
Yes
Warning: It is recommended to clear all routes after configuring PIC, reloading
device may done instead.
Execute the CLI to clear : efa inventory device clear route-all --ip 10.20.24.1
```

efa inventory device snmp community create

Creates an SNMP community and SNMP group.

Syntax

```
efa inventory device snmp community create [--ip device-ip | --name community | --group group | --enable-read-access | --enable-write-access | --enable-notify-access ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--name *community*

Specifies an SNMP community name.

--group *group*

Specifies an SNMP group name.

--enable-read-access

Sets read access for the view.

--enable-write-access

Sets write access for the view.

--enable-notify-access

Sets notify access for the view.

Usage Guidelines

The command creates the SNMP community along with the SNMP group. This command is valid for SNMP v2c version only.

EFA internally creates the v2 group and manages the community to group mapping. The group is created on the first community that is associated with the group.

Examples

```
efa inventory device snmp community create --name tempv2community  
--group tempv2group --enable-read-access --enable-write-access  
--enable-notify-access --ip 10.20.246.10
```


efa inventory device snmp community delete

Deletes an SNMP community.

Syntax

```
efa inventory device snmp community delete [--ip device-ip | --name community]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--name *community*

Specifies an SNMP community name.

Usage Guidelines

The command deletes the SNMP community. This command is valid for SNMP v2c version only.

EFA internally manages the community to group mapping and deletes the group if this is the last community associated with the group.

The SNMP group is automatically deleted when the last SNMP community associated with the group is removed. If you want to edit SNMP group attributes, all the corresponding communities must be deleted and created with modified group settings.

Examples

```
efa inventory device snmp community delete --name tempv2community  
--ip 10.20.246.10
```

efa inventory device snmp community list

Lists SNMP communities.

Syntax

```
efa inventory device snmp community list [--ip device-ip]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

The command lists SNMP communities. This command is valid for SNMP v2c version only.

Examples

```
efa inventory device snmp community list --ip 10.20.246.1-2
```

IP Address	AppState	Community Name	Group	Read view	Write view	Notify
10.20.246.2	cfg-not-managed	\$9\$jcpLTpu9FlqFFcWN7uhHig==				
10.20.246.2	cfg-in-sync	\$9\$smklvisSghOZEQvXJKBDeA==	cgrp1	efav3View	efav3View	
10.20.246.1	cfg-not-managed	\$9\$Y4Cy6IcflH9B0KM932JUNw==	v2group	efav3View	efav3View	
10.20.246.1	cfg-not-managed	\$9\$yS/jSs6dBJyG5o0yy+coKw==				
10.20.246.1	cfg-not-managed	\$9\$kr+plantb3TRIgkiBkV83Q==	v2group	efav3View	efav3View	
10.20.246.1	cfg-not-managed	\$9\$mdDwWNPg2Om04Fqnodl+Bw==	gtest1	efav3View	efav3View	
10.20.246.1	cfg-in-sync	\$9\$smklvisSghOZEQvXJKBDeA==	cgrp1	efav3View	efav3View	

Snm community details

efa inventory device snmp host create

Creates an SNMP v2c or v3 host.

Syntax

```
efa inventory device snmp host create [--ip device-ip | --host-ip { IPv4 | IPv6 | FQDN } | --community community | --user user | --version { v2c | v3 } | --notify-type { traps | informs } | --engine-id remote_id ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--host-ip { IPv4 | IPv6 | FQDN }

Specifies a host IP address.

--community *community*

Specifies a community name. Applicable for v2c only.

--user *user*

Specifies an SNMP v3 user.

--version { v2c | v3 }

Specifies the SNMP version.

--notify-type { traps | informs }

Specifies the notification type. Informs are valid for v3 only.

--engine-id *remote_id*

Specifies the remote engine ID of manager.

Hex string format with a colon as a separator. The string size must be between 10-65. Example: 02:03:04:05:06.

Usage Guidelines

The command creates an SNMP v2c or v3 host.

Engine-id and informs are applicable only for SNMP v3 hosts.

Examples

```
efa inventory device snmp host create --host-ip 20.20.20.2
--community tempv2community --version v2c --notify-type traps
--ip 10.20.246.10
```

```
efa inventory device snmp host create -host-ip 30.30.30.2
--user v3user --version v3 --notify-type informs -engine-id 02:03:04:05:06
--ip 10.20.246.10
```

efa inventory device snmp host delete

Deletes an SNMP v2c or v3 host.

Syntax

```
efa inventory device snmp host delete [--ip device-ip | --host-ip { IPv4  
| FQDN } | --community community | --user username ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example:
10.1.1.13,10.1.1.50,10.1.1.101.

--host-ip { **IPv4** | **FQDN** }

Specifies an SNMP host address.

--community *community*

Specifies an SNMP community. Applicable for v2c only.

--user *username*

Specifies an SNMP v3 user.

Examples

```
efa inventory device snmp host delete  
--host-ip 20.20.20.2 --community tempv2community --ip 10.20.246.10
```

```
efa inventory device snmp host delete  
--host-ip 30.30.30.2 --user v3user --ip 10.20.246.10
```

efa inventory device snmp host list

Lists SNMP v2c or v3 hosts.

Syntax

```
efa inventory device snmp host list [--ip device-ip]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

Examples

```
efa inventory device snmp host list --ip 10.20.246.2
```

```
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| IP Address | Host IP | User | Community | Notify Type |
| EngineID | Source Interface | Vrf | UDP port | Severity |
| AppState |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.255.241 | efav3User | mgmt-vrf | traps |
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | 162 | None |
| cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.10 | efav3User | mgmt-vrf | traps |
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | 162 | None |
| cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.24 | efav3User | mgmt-vrf | traps |
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | 162 | None |
| cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.7 | efav3User | mgmt-vrf | traps |
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | 162 | None |
| cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.85 | efav3User | mgmt-vrf | traps |
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | 162 | None |
| cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.255.115 | efav3User | mgmt-vrf | traps |
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | 162 | None |
| cfg-not-managed |
```

```
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 3.3.3.3      | user7      |          | informs    |
02:03:04:05:06      | management chassis-ip | mgmt-vrf   | 162      | None      |
cfg-in-sync      |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
```

efa inventory device snmp user create

Creates an SNMP user and SNMP group.

Syntax

```
efa inventory device snmp user create [--ip device-ip | --name community  
 | --group group | --enable-read-access | --enable-write-access | --  
enable-notify-access | --auth-protocol { md5 | sha } | --auth-pass  
authphrase | --priv-protocol { AES128 | DES } | --priv-pass  
privphrase ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--name *community*

Specifies an SNMP community name.

--group *group*

Specifies an SNMP group name.

--enable-read-access

Sets read access for the view.

--enable-write-access

Sets write access for the view.

--enable-notify-access

Sets notify access for the view.

--auth-protocol { **md5** | **sha** }

Sets notify access for the view. This parameter is set to off, by default.

--auth-pass *passphrase*

Authentication password.

--priv-protocol { **AES128** | **DES** }

Privacy protocol.

--priv-pass *privphrase*

Privacy password.

Usage Guidelines

The command creates an SNMP user along with an SNMP group. This command is valid for SNMP v3 version only.

SLX stores the `auth-pass` and `priv-pass` in encrypted format. By default, EFA encrypts all passwords.

EFA internally creates the v3 group and manages the user to group mapping. The group is created with the first user associated with the group.

The SNMP group is automatically deleted when the last SNMP community associated with the group is removed. If you want to edit SNMP group attributes, all the corresponding users must be deleted and created with modified group settings.

Examples

```
efa inventory device snmp user create --ip 10.20.246.10 --name tempv3user --group  
tempv3group  
--enable-read-access --enable-write-access --enable-notify-access --auth-protocol md5  
--auth-pass tempv3passwd --priv-protocol AES128 --priv-pass tempv3priv
```

efa inventory device snmp user delete

Deletes an SNMP user.

Syntax

```
efa inventory device snmp user delete [--ip device-ip | --name user]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--name *user*

Specifies an SNMP user name.

Usage Guidelines

The command deletes the SNMP user. This command is valid for SNMP v3 version only.

Examples

```
efa inventory device snmp user delete --name tempv3user --ip 10.20.246.10
```

efa inventory device snmp user list

Lists SNMP users.

Syntax

```
efa inventory device snmp user list [--ip device-ip]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

Usage Guidelines

The command lists SNMP user information, including IP address, group, view permissions, and pass phrases. This command is valid for SNMP v3 version only.

efa inventory device timezone debug-show

Displays the list of valid timezones supported on an SLX device.

Syntax

```
efa inventory device timezone debug-show
```

efa inventory device timezone list

Lists timezones configured on an SLX device.

Syntax

```
efa inventory device timezone list [ --ip device-ip | --fabric fabric-name ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Specifies the name of the fabric.

-h, --help

Help for list

efa inventory device timezone set

Sets the timezone on an SLX device.

Syntax

```
efa inventory device timezone set [ --ip device-ip | --fabric fabric-name  
  | --timezone region ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Specifies the name of the fabric.

--timezone *region*

Specifies the timezone region or city. For example: America /Los_Angeles.

-h, --help

Help for set.

Examples

```
efa inventory device timezone set --timezone America/Los_Angeles  
--ip 10.20.246.10
```

```
efa inventory device timezone set --timezone Europe/Paris --fabric fabric1
```

efa inventory device timezone unset

Unsets the timezone on an SLX device.

Syntax

```
efa inventory device timezone unset [ --ip device-ip | --fabric fabric-name ]
```

Parameters

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--fabric *fabric-name*

Specifies the name of the fabric.

-h, --help

Help for unset

Examples

```
efa inventory device timezone unset --ip 10.20.246.13
```

efa inventory device tpvm-upgrade execute

Asynchronously launches the TPVM upgrade process.

Syntax

```
efa inventory device tpvm-upgrade execute [--ip device-ip | --firmware-host firmware-host-ip-address | --tpvm-image tpvm-image-path-on-firmware-host | --trusted-peer-sudo-user sudo-user-id | --trusted-peer-password sudo-user-password]
```

Parameters

--ip *device-ip*

Specifies the IP address of the device.

--firmware-host *firmware-host-ip-address*

Specifies the IP address of the firmware host. The host must be registered beforehand.

--tpvm-image *tpvm-image-path-on-firmware-host*

Specifies the path and file name of the TPVM deb file located on the `firmware-host`.

--trusted-peer-sudo-user *sudo-user-id*

Specifies the user ID of the trusted peer sudo user to be reconfigured after TPVM is upgraded. If you do not specify a user, the default sudo user is used: `extreme`.

--trusted-peer-password *sudo-user-password*

Specifies the password of the trusted peer sudo user to be reconfigured after TPVM is upgraded.

Examples

This example successfully upgrades TPVM.

```
$ efa inventory device tpvm-upgrade execute --ip 10.20.245.9
--firmware-host 10.20.241.101 --tpvm-image /build/slx/tpvm/tpvm-4.2.2-1.amd64.deb
--trusted-peer-sudo-user <sudo-user-id> --trusted-peer-password <sudo-user-password>

TPVM Upgrade Execute [success]

Monitor TPVM upgrade execution progress using:

efa inventory device tpvm-upgrade show --ip 10.20.245.9

efa inventory device tpvm-upgrade show --uuid b32eee26-8b59-4bc5-8ab3-1f9884e89929
```

efa inventory device tpvm-upgrade show

Polls and monitors the progress of the TPVM upgrade workflow.

Syntax

```
efa inventory device tpvm-upgrade show [ --execution-id id | --ip ip-address ]
```

Parameters

--execution-id *id*

Specifies the ID provided from the **tpvm-upgrade execute** command.

--ip *ip-address*

Specifies the IP address of the device.

Usage Guidelines

Specify either the **--execution-id** or the **--ip** option.

Examples

This example shows only a portion of the possible output.

```
$ efa inventory device tpvm-upgrade show --ip 10.20.62.114
+-----+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model | Chassis Name | ASN | Role |
+-----+-----+-----+-----+-----+-----+
| 10.20.62.114 | POD-A-leaf4 | 4004 | SLX9740-40C | 66000 | border-leaf |
+-----+-----+-----+-----+-----+-----+
```


efa inventory device update

Updates the Asset service with the latest information from the switch.

Syntax

```
efa inventory device update [ --ip device-ip | --fabric fabric-name | --username username | --password password ]
```

Parameters

--ip *device-ip*

Specifies the IP address of the device that you want to synchronize.

--fabric *fabric-name*

Specifies the fabric of the devices that you want to synchronize.

--username *username*

Specifies the user name to connect to the device. If the user name must be updated, use this option.

--password *password*

Specifies the password to connect to the device. If the password must be updated, use this option.

Examples

This example (which shows only a portion of all possible output) synchronizes IP address 10.24.80.158.

```
$ efa inventory device update --ip 10.24.80.158
-----
ID      IP Address      Host Name      Model      Chassis Name      Firmware
-----
2       10.24.80.158   SLX            3012      SLX9250-32C      20.1.2slxos20.1.2x
-----
```

efa inventory drift-reconcile

Identifies drift in device configuration and performs reconciliation.

Syntax

```
efa inventory drift-reconcile history [ --ip ip-address | --reverse ]  
efa inventory drift-reconcile execute [ --ip ip-address | --reconcile ]  
efa inventory drift-reconcile delete [ --key{ ip-address | dr-uuid} ]  
efa inventory drift-reconcile detail [ --uuid dr-uuid ]
```

Parameters

--ip *ip-address*
Specifies the IP address of the device.

--key { **ip-address** | *dr-uuid* }
Specifies IP address of the device or the drift and reconcile UUID to be deleted.

--reverse
Sort display in reverse start time.

--reconcile
Identifies drift and reconcile with device.

--uuid *dr-uuid*
Specifies the UUID for the drift reconcile.

Usage Guidelines

Use the **efa inventory drift-reconcile history** command to display the history of drift and reconcile operations.

Use the **efa inventory drift-reconcile execute** command to perform a drift and reconcile operation.

Use the **efa inventory drift-reconcile delete** command to delete a drift and reconcile record for the specified IP address or UUID.

Use the **efa inventory drift-reconcile detail** command to display drift and reconcile details.

Examples

The following examples show **efa inventory drift-reconcile** commands.

```
# efa inventory drift-reconcile execute --ip 10.24.14.133 --reconcile  
# efa inventory drift-reconcile history --device-ip 10.24.14.133  
# efa inventory drift-reconcile detail --uuid 1111-1111-1111
```

```
# efa inventory drift-reconcile delete --key 10.24.14.133  
# efa inventory drift-reconcile delete --key 1111-1111-1111
```

efa inventory execution

Displays the list of inventory execution operations.

Syntax

```
efa inventory execution show [ --id id | --limit value | --status  
  { failed | succeeded | all } ]
```

Parameters

--id *id*

Filters the operations by ID. If **--id** is available, **--limit** and **--status** are ignored.

--limit *value*

Limits the number of operations to be listed. Value 0 lists all operations. The default is 10.

--status { **failed** | **succeeded** | **all** }

Filters the operations based on the status. The default is **all**.

efa inventory firmware-host delete

Removes a firmware host.

Syntax

```
efa inventory firmware-host delete [--ip host-ip ]
```

Parameters

--ip *host-ip*

Specifies the IP address of the firmware host.

efa inventory firmware-host list

Displays all registered firmware hosts.

Syntax

```
efa inventory firmware-host list [ --ip host-ips ]
```

Parameters

--ip *host-ips*

Specifies a comma-separated list of firmware host IP addresses.

efa inventory firmware-host register

Registers the firmware host that will be used to download firmware builds to the devices.

Syntax

```
efa inventory firmware-host register [--ip host-ip | --protocol { scp | ftp | sftp | http } | --username username | --password password ]
```

Parameters

--ip *host-ip*

Specifies the IP address of the firmware host.

--protocol { **scp** | **ftp** | **sftp** | **http** }

Specifies the protocol to use for downloading the firmware.

--username *user name*

Specifies the user name to use to log in to the firmware download host.

--password *password*

Specifies the password for the user name.

Usage Guidelines

Simple connectivity test to the firmware-host by given IP. The complete firmware-host sanity check is performed later when a device is prepared and again when firmware-download is run.

efa inventory firmware-host update

Updates the login credentials or file transfer protocol to be used by a device when downloading firmware from the firmware host.

Syntax

```
efa inventory firmware-host update [--ip host-ip | --protocol { scp | ftp | sftp | http } | --username username | --password password ]
```

Parameters

--ip *host-ip*

Specifies the IP address of the firmware host.

--protocol { **scp** | **ftp** | **sftp** | **http** }

Specifies the protocol to use for downloading the firmware.

--username *user name*

Specifies the user name to use to log in to the firmware download host.

--password *password*

Specifies the password for the user name.

efa inventory kvstore

Configures a key-value pair.

Syntax

```
efa inventory kvstore create [ --key name | --value value | --encrypt ]
```

```
efa inventory kvstore delete [ --key name ]
```

```
efa inventory kvstore list [ --decrypt | --key name | --prefix prefix ]
```

Parameters

--decrypt

Decrypts the secret fields.

--encrypt

Encrypts the secret fields.

--key *name*

Specifies a key-value pair name.

--prefix *prefix*

Retrieve list of a key-value pairs matching the prefix.

--value *value*

Specifies a value for the key.

efa inventory rma

Initiates Return Material Authorization (RMA).

Syntax

```
efa inventory rma history [ --ip ip-address ]  
efa inventory rma detail [ --uuid uuid ]  
efa inventory rma delete [ --key ip-address ]  
efa inventory rma execute [ --ip ip-address | --config-backup-id uuid ]
```

Parameters

--ip *ip-address*
Specifies the IP address of the device.

--uuid *uuid*
Specifies the UUID of the drift reconcile.

--key *ip-address*
Specifies IP address of the device or RMA UUID to be deleted.

--config-backup-id *uuid*
Specifies the UUID of the configuration to be replayed.

Usage Guidelines

Use the **efa inventory rma history** command to display the RMA history for the specified IP address.

Use the **efa inventory rma detail** command to display device replacement details.

Use the **efa inventory rma delete** command to delete an RMA record for the specified IP address.

Use the **efa inventory rma execute** command to perform an RMA operation.

Examples

The following example shows **efa inventory rma** commands.

```
# efa inventory rma execute --ip 10.24.14.133  
--config-backup-id 1111-1111-1111  
# efa inventory rma history -ip 10.24.14.133  
# efa inventory rma detail -uuid 123e4567-e89b-12d3-a456-426614174000
```

efa login

Logs into the EFA application.

Syntax

```
efa login [ --username | --password ]
```

Parameters

--username

Specifies name of the user to login.

--password

Specifies password of the user.

efa logout

Logs out of the EFA application.

Syntax

```
efa logout
```

efa mgmt route create

Creates a virtual management route in a multi-node deployment for the Multiple Management IP Networks feature.

Syntax

```
efa mgmt route create [ --src mmip-vip | --to dest-cidr | --via next-hop-ip ]
```

Parameters

--src *mmip-vip*

Identifies the source IP address for the route.

--to *dest-cidr*

Identifies the destination network for outbound traffic, in CIDR format: x.x.x/xx.

--via *next-hop-ip*

Identifies the next-hop or gateway IP address through which access to the destination network is provided.

Examples

The following example creates a virtual route.

```
$ efa mgmt route create --src 10.21.30.40 --to 192.168.100.0/24 --via 10.21.30.41
Virtual Route created successfully
```

efa mgmt route delete

Deletes a virtual management route in a multi-node deployment for the Multiple Management IP Networks feature.

Syntax

```
efa mgmt route delete [ --src mmip-vip | --to dest-cidr | --via next-hop-ip ]
```

Parameters

--src *mmip-vip*

Identifies the source IP address for the route.

--to *dest-cidr*

Identifies the destination network for outbound traffic, in CIDR format: x.x.x/xx.

--via *next-hop-ip*

Identifies the next-hop or gateway IP address through which access to the destination network is provided.

Examples

The following example deletes a virtual route.

```
$ efa mgmt route delete --src 10.21.30.40 --to 192.168.100.0/24 --via 10.21.30.41
Virtual Route deleted successfully
```

efa mgmt route show

Displays a list of all virtual management routes in a multi-node deployment for the Multiple Management IP Networks feature.

Syntax

```
efa mgmt route show
```

Examples

The following example shows typical output for the command.

```
$ efa mgmt route show

+-----+-----+-----+
| Route-Src | Route-To | Route-Via |
+-----+-----+-----+
| 10.21.30.40 | 192.168.100.0/24 | 10.21.30.41 |
+-----+-----+-----+
```

efa mgmt subinterface create

Creates a subinterface in a Multiple Management IP network deployment.

Syntax

```
efa mgmt subinterface create [ --name sub | --vlan-id vlan-id | --ip-addr ip-addr ]
```

Parameters

--name *sub*

Specifies the subinterface name.

--vlan-id *vlan-id*

Specifies the ID of the VLAN for the subinterface.

--ip-addr *ip-addr*

Specifies the IP address of the subinterface, including the subnet mask. For example: 10.24.80.150/24.

Examples

```
$ efa mgmt subinterface create --name sub-200 --vlan-id 10  
--ip-address 192.168.150.8/24  
  
SubInterface sub-200 created successfully
```


efa mgmt subinterface delete

Deletes a specified subinterface in a Multiple Management IP network deployment.

Syntax

```
efa mgmt subinterface delete --name sub
```

Parameters

--name *sub*

Specifies the name of the subinterface that you want to delete.

Examples

The following example deletes the server1 subinterface.

```
$ efa mgmt subinterface delete --name server1  
Subinterface server1 deleted successfully
```

efa mgmt subinterface show

Displays the details of all subinterfaces in a Multiple Management IP network deployment or the details for a specified subinterface.

Syntax

```
efa mgmt subinterface show [ --name name ]
```

Parameters

--name *name*

Specifies the name of the subinterface for which you want to see details.

Examples

The following example shows the details for all configured subinterfaces.

```
$ efa mgmt subinterface show
+-----+-----+-----+-----+
| Name | Parent Interface | Vlan | IP Subnet |
+-----+-----+-----+-----+
| sub2 | eth0             | 20  | 20.20.20.2/24 |
+-----+-----+-----+-----+
| sub1 | eth0             | 50  | 50.50.50.2/24 |
+-----+-----+-----+-----+
```

The following example shows the details for the specified subinterface.

```
$ efa mgmt subinterface show --name sub2
+-----+-----+-----+-----+
| Name | Parent Interface | Vlan | IP Subnet |
+-----+-----+-----+-----+
| sub2 | eth0             | 20  | 20.20.20.2/24 |
+-----+-----+-----+-----+
```

efa mgmt subinterface staticip add

Sets static IP addresses for a subinterface.

Syntax

```
efa mgmt subinterface staticip add [ --subinterface sub | --ip1 ip1 | --  
ip2 ip2 ]
```

Parameters

--subinterface *sub*

Name of the subinterface for which you want to set static IP addresses.

--ip1 *ip1*

First static IP address. Example: 192.168.10.1/32

--ip2 *ip2*

Second static IP address. Example: 192.168.10.2/32

Usage Guidelines

You can assign a maximum of one pair of static IP addresses. Only one subinterface at a time can have static IP addresses.

Examples

```
$ efa mgmt subinterface staticip add --subinterface sub-200 --ip1 192.168.150.1/24  
--ip2 192.168.150.2/24
```

Static IP assigned to SubInterface sub-200 successfully

```
+-----+-----+  
| Node | IP Address |  
+-----+-----+  
| node1 | 192.168.150.1/24 |  
+-----+-----+  
| node2 | 192.168.150.2/24 |  
+-----+-----+  
Add SubInterface Static IP  
Details
```

efa mgmt subinterface staticip remove

Removes the IPs that are attached to the specified subinterface.

Syntax

```
efa mgmt subinterface staticip remove [ --subinterface sub ]
```

Parameters

--subinterface *sub*

Specifies the subinterface for which to remove IPs.

Examples

```
(efa:extreme)extreme@testuser:~$ efa mgmt subinterface staticip remove --subinterface
sub-200

Static IP deleted from SubInterface sub-200 successfully

+-----+-----+
| Node | IP Address |
+-----+-----+
| testuser | 192.168.150.1/24 |
+-----+-----+
| testuser2 | 192.168.150.2/24 |
+-----+-----+
Delete SubInterface Static IP
Details
```

efa mgmt subinterface staticip show

Shows all the subinterfaces and the IPs that are attached to them.

Syntax

```
efa mgmt subinterface staticip show
```

Examples

```
(efa:extreme)extreme@testuser:~$ efa mgmt subinterface staticip show
+-----+-----+-----+-----+-----+
| SubInterface | Node1   | Node1 IP       | Node2   | Node2 IP       |
+-----+-----+-----+-----+-----+
| sub-200      | testuser | 192.168.150.1/24 | testuser2 | 192.168.150.2/24 |
+-----+-----+-----+-----+-----+
Static IP Details
```

efa notification subscribers add-https

Registers a new subscriber to the Notification service with an HTTPS webhook.

Syntax

```
efa notification subscribers add-https [ --url address | --username name
| --password password | --insecure | --cacert string ]
```

Parameters

--url *address*

(Required) Specifies the URL of the subscriber.

--username *name*

(Required) Specifies the user name for access to the URL.

--password *password*

(Optional) Specifies the password for the user name.

--insecure

(Optional) Indicates that insecure SSL connection and transfers are used for sending notifications. By default, the SSL connection and transfers are secure.

--cacert *string*

(Optional) Local path to the cacert pem file for SSL verification.

The following example registers a new HTTPS webhook subscriber to the Notification Service.

```
efa notification subscribers add-https --url https://path/to/my/service:port
--username myusername --password mypassword --insecure
Successfully registered subscriber.
```

```
+-----+-----+
| attribute | value |
+-----+-----+
| id        | 1     |
+-----+-----+
| handler   | http  |
+-----+-----+
| endpoint  | https://path/to/my/service:port |
+-----+-----+
| config    | {"cacert":"","insecure":true,"password":"mypassword",
      "username":"myusername"} |
+-----+-----+
Notification Subscriber ID=1
--- Time Elapsed: 2.203878641s ---
```

efa notification subscribers add-syslog-relp

Registers a syslog subscriber to receive notifications over Reliable Event Logging Protocol (RELP) from the EFA Notification service.

Syntax

```
efa notification subscribers add-syslog-relp [ --address host:port ] [ --insecure ] [ --cacert local-path ] [ --conn-timeout seconds ]
```

Command Default

By default, subscribers do not receive syslog notifications over RELP.

Parameters

--address *host:port*

(Required) Specifies the address of the syslog server in *host:port* format. The default port is 514.

--insecure

(Optional) Indicates that insecure SSL connection and transfers are used for sending notifications. By default, the SSL connection and transfers are secure.

--cacert *local-path*

(Optional) Specifies the local path to the cacert.pem file for SSL verification. Required *only* when the **--insecure** parameter is not specified.

--conn-timeout *seconds*

(Optional) Specifies the maximum amount of time allowed to open a connection to the syslog server before the request times out. The default is 10 seconds.

Usage Guidelines

Any external server that is configured with RELP can be registered as a subscriber to EFA notifications. For more information, see "Notification Service" in the *Extreme Fabric Automation Administration Guide*.

Examples

The following example registers 10.x.x.x:20514 as a subscriber for insecure notifications.

```
$ efa notification subscribers add-syslog-relp --address
10.x.x.x:20514 --insecure
Successfully registered subscriber.
```

```
+-----+-----+
| attribute | value |
+-----+-----+
| id        | 1     |
+-----+-----+
| handler   | relp  |
+-----+-----+
```

```
+-----+-----+
| endpoint | 10.x.x.x:20514 |
+-----+-----+
| config   | {"cacert":"","conn-timeout":10,"insecure":true}|
+-----+-----+
Notification Subscriber ID=1
--- Time Elapsed: 2.399195253s ---
```


efa notification subscribers delete

Deletes the specified subscriber from the Notification service.

Syntax

```
efa notification subscribers delete id
```

Parameters

id

Specifies the ID of the subscriber that you want to delete.

Examples

The following example deletes the specified subscriber from the Notification service.

```
efa notification subscribers delete 1
Successfully unregistered subscriber.
--- Time Elapsed: 186.568274ms ---
```

efa notification subscribers get

Retrieves subscription details about the Notification service for the specified subscriber.

Syntax

```
efa notification subscribers get id
```

Parameters

id

Specifies the ID of the subscriber that you want to review.

Examples

The following is sample output for the subscriber with the ID of '1'.

```
efa notification subscribers get 1

+-----+-----+
| attribute | value |
+-----+-----+
| id        | 1     |
+-----+-----+
| handler   | http  |
+-----+-----+
| endpoint  | https://path/to/my/service:port |
+-----+-----+
| config    | {"cacert":"","insecure":true,"password":"mypassword",
|           | "username":"myusername"} |
+-----+-----+
Notification Subscriber ID=1
--- Time Elapsed: 2.203878641s ---
```

efa notification subscribers list

Provides a list of subscribers to the Notification service.

Syntax

```
efa notification subscribers list
```

Examples

The following example lists the subscribers to the Notification service.

```
efa notification subscribers list
+-----+-----+-----+
| id | handler | endpoint |
+-----+-----+-----+
| 1 | http | https://path/to/my/service:port |
+-----+-----+-----+
Notification Subscribers Count=1
--- Time Elapsed: 121.443493ms ---
```

efa openstack debug

Displays OpenStack debug information.

Syntax

```
efa openstack debug [ network | network-interface | tenant | router |  
router-interface | router-route ]
```

```
efa openstack debug network delete [--neutron-id id]
```

Deletes the network and its summary information.

```
efa openstack debug network-interface delete [--neutron-id id]
```

Deletes the network interface and its summary information.

```
efa openstack debug tenant cleanup [--name name ]
```

Deletes a network.

```
efa openstack debug router delete [id]
```

Deletes the router and its summary information.

```
efa openstack debug router-interface delete [--router-id id --subnet-id  
id]
```

Deletes the router interface and its summary information.

```
efa openstack debug router-route delete [--router-id id | --destination  
cidr | --nexthop ip]
```

Deletes the router route.

Parameters

cleanup

Cleans up all OpenStack assets associated with a tenant.

delete

Deletes the selected network element.

network

Specifies network commands.

network-interface

Specifies network interface commands.

tenant

Specifies tenant commands.

router

Specifies router commands.

--router-interface

Specifies router interface commands.

--neutron-id *id*

Specifies the Neutron ID of the network.

--router-id *id*

Specifies a comma-separated range of router IDs.

--subnet-id *id*

Specifies the subnet ID.

--name *name*

Specifies the name of the tenant.

--destination *cidr*

Specifies the destination CIDR (classless inter-domain routing).

--nexthop *ip*

Specifies the IP address of the next hop.

efa openstack execution

Provides OpenStack execution commands.

Syntax

```
efa openstack execution delete [ --days days ]
```

```
efa openstack execution show [ --id id | --limit num | --status {failed |  
succeeded | all} ]
```

Parameters

--days *days*

Deletes execution entries older than the specified days. The default is 30.

--id *id*

Filters the executions based on execution id. `limit` and `status` flags are ignored when the `id` flag is specified.

--limit *num*

Limits the number of executions to be listed. 0 will list all the executions. The default is 10.

--status {**failed** | **succeeded** | **all**}

Filters the executions based on the status. The default is all.

efa openstack network show

Displays OpenStack network information.

Syntax

```
efa openstack network show [ --id id ]
```

Parameters

--id *id*

(Optional) Specifies the network ID.

Examples

This example shows typical results.

```
$ efa openstack network show
+-----+-----+-----+-----+
|           Neutron ID           | Tenant | MTU | CTAG |
+-----+-----+-----+-----+
| 9e44630a-5f4c-4f08-98a1-c9d03147116d | RegionTwo | 2100 | 400 |
+-----+-----+-----+-----+
| 9769f154-4f5c-4350-88de-88c0b0d9f00c | RegionTwo | 2100 | 1400 |
+-----+-----+-----+-----+
| 37a496fc-0ad1-49c2-a497-28e2902b9f1b | RegionTwo | 2100 | 401 |
+-----+-----+-----+-----+
| 70905215-c566-4a75-bce8-fef8d8008fe4 | RegionTwo | 2100 | 1401 |
+-----+-----+-----+-----+
```

efa openstack network-interface show

Displays OpenStack network interface information.

Syntax

```
efa openstack network-interface show [ --id id ]
```

Parameters

--id *id*

Specifies the network ID.

Examples

This example shows typical, but truncated, results.

```
$ efa openstack network-interface show

+-----+-----+
|          Neutron Port ID          |          Neutron Network ID          |
+-----+-----+
| e8a0227b-9e59-487d-8a7d-0322ce2f362e | 9e44630a-5f4c-4f08-98a1-c9d03147116d |
+-----+-----+
| 70680aad-caa0-40a0-993b-550d58fd9906 | 9e44630a-5f4c-4f08-98a1-c9d03147116d |
+-----+-----+
| 77b0adea-6db4-4389-8cd0-a1bbb9e06a18 | 9769f154-4f5c-4350-88de-88c0b0d9f00c |
+-----+-----+
| fa41f4e2-aadd-4b47-9171-4dbc07591fe2 | 9769f154-4f5c-4350-88de-88c0b0d9f00c |
+-----+-----+
```


efa openstack router show

Displays OpenStack router information.

Syntax

```
efa openstack router show [ --id id ]
```

Parameters

--id *id*

Specifies the network ID.

Examples

This example shows typical, but truncated, results.

```
efa openstack router show
+-----+-----+
| Router ID | VRF Name |
+-----+-----+
| bc482e46-8896-442c-a29d-2dd76abca2ae | bc482e468896442ca29d2dd76abca2ae |
+-----+-----+
| 99ec4fd8-fbc4-4117-a529-5e9a40049906 | 99ec4fd8fbc44117a5295e9a40049906 |
+-----+-----+
```

efa openstack router-interface show

Displays OpenStack router interface information.

Syntax

```
efa openstack router-interface show[ --id id ]
```

Parameters

--id *id*

Specifies the network ID.

Examples

This example shows typical results.

```
# efa openstack router-interface show
+-----+-----+
|          Subnet ID          |          Router ID          |
+-----+-----+
| 323e4567-e89b-12d3-a456-426655440001 | 523e4567-e89b-12d3-a456-426655440001 |
+-----+-----+
| 323e4567-e89b-12d3-a456-426655440002 | 523e4567-e89b-12d3-a456-426655440001 |
+-----+-----+
```

efa openstack router-route show

Displays destination, next hop, and device IP addresses on the specified router.

Syntax

```
efa openstack router-route show [--id router-id]
```

Parameters

--id *router-id*
Specifies the router ID.

efa openstack subnet show

Displays OpenStack subnet information, including network ID, CIDR, and gateway IP address.

Syntax

```
efa openstack subnet show [ --id id ]
```

Parameters

--id *id*

Specifies the subnet ID.

efa openstack sync start

Syncs entries such as networks, network interfaces, routers, and router interfaces from OpenStack to a tenant if they are out of sync.

Syntax

```
efa openstack sync start [ --tenant tenant name ]
```



Note

There is no output for this command.

Parameters

--tenant

The name of the tenant

efa policy prefix-list create

Creates an IPv4 prefix list.

Syntax

```
efa policy prefix-list create [ --type string | --name string | --rule
stringArray ]
```

Parameters

--type *string*

Specifies the type of prefix list. The only valid value is `ipv4`.

--name *string*

Specifies the name of the prefix list.

--rule *stringArray*

Specifies the rule of the prefix list, in the following format: `seq[seq-num],action[permit/deny],prefix[IPv4 prefix],ge[prefix-len],le[prefix-len]`. For example: `seq[5],action[permit],prefix[10.0.0.0/8],ge[10],le[24]`.

Usage Guidelines

Execute **efa policy prefix-list update --operation add-device** to push the configuration on to the device.

Examples

This example creates an IPv4 prefix list.

```
$ efa policy prefix-list create --name prefix_v4
--type ipv4 --rule seq[5],action[permit],prefix[10.0.0.0/8],ge[16]
```

This example creates an IPv4 prefix list with two rules.

```
efa policy prefix-list create --type ipv4
--name test1 --rule seq[5],action[permit],prefix[10.0.0.0/8],ge[16],le[24]
--rule seq[10],action[deny],prefix[20.0.0.0/8],ge[14],le[24]
Name: test1
+-----+-----+-----+-----+-----+-----+
| Type | Seq num | Action | Prefix | Ge | Le | Status |
+-----+-----+-----+-----+-----+-----+
| ipv4 | 5       | permit | 10.0.0.0/8 | 16 | 24 | Success |
+-----+-----+-----+-----+-----+-----+
| ipv4 | 10      | deny  | 20.0.0.0/8 | 14 | 24 | Success |
+-----+-----+-----+-----+-----+-----+
Prefix-list details
--- Time Elapsed: 597.965605ms ---
```

efa policy prefix-list update

Configures prefix list rules on specified devices.

Syntax

```
efa policy prefix-list update [ --type string | --name string | --  
  operation { add-rule | remove-rule | add-device | remove-device } | --  
  rule stringArray | --ip string ]
```

Parameters

--type *string*

Specifies the type of prefix list. The only valid value is `ipv4`.

--name *string*

Specifies the name of the prefix list.

--operation { **add-rule** | **remove-rule** | **add-device** | **remove-device** }

Specifies the operation to perform.

--rule *stringArray*

Specifies the rule of the prefix list, in the following format: `seq[seq-num],action[permit/deny],prefix[IPv4 prefix],ge[prefix-len],le[prefix-len]`. For example: `seq[5],action[permit],prefix[10.0.0.0/8],ge[10],le[24]`.

--ip *string*

Specifies a comma-separated range of device IP addresses. For example: `10.1.1.13,10.1.1.50,10.1.1.101`.

Usage Guidelines

- The **add-device** and **remove-device** operations configure the prefix list rules on the specified devices.
- The **add-rule** and **remove-rule** operations add or remove a prefix list rule on the specified devices. If the prefix list is configured on the device, the rule is added or removed from the device.
- This command is also used to add or delete a prefix-list rule.

Examples

This example configures a prefix list on the specified devices.

```
efa policy prefix-list update --name prefix_v4  
--type ipv4 --operation add-device --ip 10.20.246.10-11
```

This example removes the prefix list from the specified devices.

```
efa policy prefix-list update --name prefix_v4 --type ipv4  
--operation remove-device --ip 10.20.246.10-11
```

This example adds a rule to an existing prefix list.

```
efa policy prefix-list update --name prefix_v4
--type ipv4 --operation add-rule
--rule seq[5],action[permit],prefix[10.0.0.0/8],ge[16]
```

This example removes a rule from an existing prefix list.

```
efa policy prefix-list update --name prefix_v4
--type ipv4 --operation remove-rule
--rule seq[5],action[permit],prefix[10.0.0.0/8],ge[16]
```


efa policy prefix-list delete

Deletes the IPv4 prefix list on all devices and from EFA.

Syntax

```
efa policy prefix-list delete [ --type string | --name string ]
```

Parameters

--type *string*

Specifies the type of prefix list. The only valid value is `ipv4`.

--name *string*

Specifies the name of the prefix list.

Examples

This example deletes an IPv4 prefix list named `prefix_v4`.

```
$ efa policy prefix-list delete --type ipv4 --name prefix_v4
```

efa policy prefix-list list

Lists the IPv4 prefix list on the specified devices.

Syntax

```
efa policy prefix-list list [ --type string | --ip string ]
```

Parameters

--type *string*

Specifies the type of prefix list. The only valid value is `ipv4`.

--ip *string*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

Examples

This example displays the IPv4 prefix list on two devices.

```
$ efa policy prefix-list list --type ipv4 --ip 10.20.246.10-11
Prefix-list details:

Name: prefix_v4_1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type | Seq num | Action | Prefix | Ge | Le | DeviceIP | AppState |
+-----+-----+-----+-----+-----+-----+-----+-----+
| ipv4 | 10      | permit | 10.0.0.0/8 | 16 |   | 10.20.246.10 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+-----+
| ipv4 | 10      | permit | 10.0.0.0/8 | 16 |   | 10.20.246.11 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

efa policy route-map create

Creates a route map with one or more rules.

Syntax

```
efa policy route-map create [ --name string | --rule stringArray]
```

Parameters

--name *string*

Specifies the name of the route map.

--rule *stringArray*

Specifies the rule in the following format: seq[seq-num],action[permit/deny].

Usage Guidelines

Execute **efa policy route-map update --operation add-device** to push the configuration on to the device.

Examples

This example creates a route map, rmap_1, with two rules.

```
efa policy route-map create --name rmap_1 --rule seq[5],action[permit]
--rule seq[10],action[permit]
```

efa policy route-map update

Adds or deletes the route map configuration on a list of devices. Can be used to update the action for a given route map rule.

Syntax

```
efa policy route-map update [ --name string | --rule string | --operation  
  { add-device | remove-device | update-action } | --ip string]
```

Parameters

--name *string*

Specifies the name of the route map that you want to update.

--rule *string*

Specifies the rule in the following format: seq[seq-num],action[permit/deny].

--operation { **add-device** | **remove-device** | **update-action** }

Specifies the type of operation that you want to perform.

--ip *string*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

Usage Guidelines

A route map can be associated with multiple rules.

For an **add-device** operation, specify **only** the device IP address and route map name. The operation adds all the rules of the route map on the specified devices.

For a **remove-device** operation, specify **only** the device IP address and route map name. The operation deletes all the rules of the route map on the specified devices.



Note

The **remove-device** operation fails if the route map is bound to a BGP neighbor.

For an **update-action** operation, specify the route map name and the rule. You can modify the action to permit or deny for a specific rule. You can provide only one rule at a time.

Examples

This example configures a route map rule on devices 10.20.246.10 and 10.20.246.11. Assume there are two route map rules for a map named rmap_1 that already exists in EFA:

- rmap_1 seq 5 action permit
- rmap_1 seq 10 action permit

```
efa policy route-map update --name rmap_1 --operation add-device  
--ip 10.20.246.10-11
```

This example removes a route map from the specified devices.

```
efa policy route-map update --name rmap_1 --operation remove-device  
--ip 10.20.246.10-11
```

This example changes the action from permit to deny for the specified rule.

```
efa policy route-map update --name rmap_1  
--rule seq[5],action[deny] --operation update-action
```

efa policy route-map delete

Deletes a route map rule, based on the specified sequence number.

Syntax

```
efa policy route-map delete [ --name string | --seq sequence-numbers ]
```

Parameters

--name *string*

Specifies the name of the route map.

--seq *sequence-numbers*

Specifies a comma-separated list of the sequence numbers of the rules you want to delete, or specifies **all** to delete all rules.

Usage Guidelines

The command removes the route map rule from the EFA database and from the associated devices.

You can delete a specific rule of a route map by specifying the route map name and the sequence number of the rule.

You can delete all rules for the specified route map by entering **--seq all**.

The result of this command depends on whether the route map is bound with a BGP neighbor.

- You cannot delete the last route map rule when the route map is bound to a BGP peer.
- If the route map has no bindings, the command deletes the configuration on all devices associated with the route map.

Examples

This example deletes two rules with sequence numbers 5 and 10 from a route map (rmap_1) that has three rules:

- rmap_1 seq 5 action permit
- rmap_1 seq 10 action permit
- rmap_1 seq 20 action permit

```
efa policy route-map delete --name rmap_1 --seq 5,10
```

This example deletes all rules for the route map.

```
efa policy route-map delete --name rmap_1 --seq all
```

efa policy route-map list

Displays the route maps for the specified devices.

Syntax

```
efa policy route-map list [ --ip string]
```

Parameters

--ip *string*

Specifies a comma-separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50,10.1.1.101.

Usage Guidelines

In the command output, the App State column reflects the state of configuration on the specified device. When there is drift in a rule, the App State is shown as `cfg-refreshed`.

Examples

This example displays the route maps for the devices at 10.20.246.10 and 10.20.246.11.

```
efa policy route-map list --ip 10.20.246.10-11

Route-map details:
Name: rmap_1
Seq: 5
Action: permit
Match-ipv4-prefixlist:
  Prefix-list: prefix_1
Name: rmap_2
Seq: 5
Action: permit
Match-ipv4-prefixlist:
  Prefix-list: prefix_1
IP Addresses:
+-----+-----+-----+-----+
|      Name      | Seq | IP Address | App State |
+-----+-----+-----+-----+
| rmap_1         | 5   | 10.20.246.10 | cfg-in-sync |
+-----+-----+-----+-----+
| rmap_1         | 5   | 10.20.246.11 | cfg-in-sync |
+-----+-----+-----+-----+
| rmap_2         | 5   | 10.20.246.10 | cfg-in-sync |
+-----+-----+-----+-----+
| rmap_2         | 5   | 10.20.246.11 | cfg-in-sync |
+-----+-----+-----+-----+
```

efa policy route-map-match create

Creates the match criterion for a route map rule.

Syntax

```
efa policy route-map-match create [ --name string | --rule string | --match-ipv4-prefix string ]
```

Parameters

--name *string*

Specifies the name of the route map.

--rule *string*

Specifies the rule in the following format: seq[seq-num],action[permit/deny].

--match-ipv4-prefix *string*

Specifies the name of the IPv4 prefix list that contains the rule that you want to match.

Usage Guidelines

- The IPv4 prefix list is the only match supported.
- There is no separate CLI to display the route-map match criteria. The information is shown in **efa policy route-map list**.

Examples

```
efa policy route-map-match create --name rmap_1  
--rule seq[5],action[permit] --match-ipv4-prefix prefix_1
```


efa policy route-map-match delete

Deletes the match criterion for a route map rule.

Syntax

```
efa policy route-map-match delete [ --name string | --rule string | --match-ipv4-prefix string ]
```

Parameters

--name *string*

Specifies the name of the route map.

--rule *string*

Specifies the rule in the following format: seq[seq-num],action[permit/deny].

--match-ipv4-prefix *string*

Specifies the name of the IPv4 prefix list that contains the rule that you want to delete.

Usage Guidelines

The IPv4 prefix list is the only match supported.

Examples

This example deletes the rule with sequence number 5.

```
efa policy route-map-match delete --name rmap_1  
--rule seq[5],action[permit]
```

efa rbac execution

Displays or deletes the execution logs of the Auth and RBAC services.

Syntax

```
efa rbac execution show [ --id id | --limit exec-limit | --status
  { failed | succeeded | all } ]
efa rbac execution delete [ --day days ]
```

Parameters

--id *id*

Filters the log entries based on execution ID. The `limit` and `status` flags are ignored when the `id` flag is used.

--limit *exec-limit*

Limits the number of log entries to be listed. Enter 0 to show all of the executions. The default is 10.

--status { **failed** | **succeeded** | **all** }

Filters the log entries based on the status. The default is `all`.

--day *days*

Deletes log entries older than the specified number of days. The default is 30 days.

efa rbac role show

Displays the static and dynamic roles available in EFA.

Syntax

```
efa rbac role show [ --name role-name ]
```

Parameters

--name *role-name*

Specifies name of the role.

Usage Guidelines

Output of the command is as follows.

Role	Description
FabricAdmin	<ul style="list-style-type: none"> Registers devices to the fabric Configures fabric parameters Validates all devices in the fabric Configures switches for IP fabric with overlay and without overlay Creates tenants Creates networks inside tenants, such as VRF, EPG, and PO Performs fabric debug activities Has privileges for OpenStack, Hyper-V, and vCenter operations
SecurityAdmin	Performs user management, PKI, and key management operations
NetworkOperator	<ul style="list-style-type: none"> Has view-only privileges for fabric configurations, information for tenants and inventory, and all ecosystem information Cannot make changes in the system
SystemDebugger	<ul style="list-style-type: none"> Has privileges to perform supportsave and system backup, and to view the running system configurations Has privileges to perform fabric debug operations Sets debug levels for services Has privileges to collect execution logs from services

Role	Description
SystemAdmin	Has complete privileges to all operations in the system
<Tenant>Admin * Created dynamically per tenant	Performs tenant administration within the assigned tenant, such as the following: <ul data-bbox="789 369 1430 474" style="list-style-type: none">• Adding networks to the tenant• Configuring network parameters• Configuring switches with tenant-specific information Cannot perform actions for any other tenant

efa scvmm delete

Deletes the SCVMM server.

Syntax

```
efa scvmm delete [--host ip | --cleanup-epgs ]
```

Parameters

--host *ip*

Specifies the IP address or host name of the SCVMM server to be deleted.

--cleanup-epgs

Cleans up EPGs from the tenant service when the SCVMM server is deleted.

efa scvmm links physical

Configures SCVMM physical links.

Syntax

```
efa scvmm links physical [ --host ip-host | --hyperv hostname ]
```

Parameters

--host *ip-host*

Specifies the IP address or host name of the SCVMM server.

--hyperv *hostname*

Specifies the name of the Hyper-V host.

efa scvmm links virtual

Configures SCVMM virtual links.

Syntax

```
efa scvmm links virtual [ --hyperv hostname ]
```

Parameters

--hyperv *hostname*

Specifies the name of the Hyper-V host.

efa scvmm list

Lists all SCVMM servers and their summary information.

Syntax

```
efa scvmm list [ --host string]
```

Parameters

--host *string*

Specifies the name or IP address of the SCVMM server.

efa scvmm register

Registers an SCVMM server.

Syntax

```
efa scvmm register [ --host ip | --username user name | --password password | --tenant tenant ]
```

Parameters

--host *ip*

Specifies the IP address or name of the SCVMM server to connect to.

--username *user name*

Specifies the user name with which to connect to the SCVMM server.

--password *password*

Specifies the password with which to connect to the SCVMM server.

--tenant *tenant*

Specifies the tenant name for the SCVMM server.

efa scvmm settings show

Lists SCVMM Service settings.

Syntax

```
efa scvmm settings show
```

efa scvmm settings update

Updates SCVMM service settings.

Syntax

```
efa scvmm settings update [ --poll-frequency frequency ]
```

Parameters

--poll-frequency *frequency*

Specifies the polling frequency in hours. Valid values range from 1 to 24.

efa scvmm update

Updates the SCVMM server.

Syntax

```
efa scvmm update [ --host ip-host | --username user name | --password  
password ]
```

Parameters

--host *ip-host*

Specifies the IP address or host name of the SCVMM server you want to update.

--username *user name*

Specifies the user name with which to connect to the SCVMM server.

--password *password*

Specifies the password with which to connect to the SCVMM server.

efa snmp subscriber

Lists, adds, or unregisters SNMP trap subscribers.

Syntax

```
efa snmp subscriber list
```

```
efa snmp subscriber add [ --host ip | --type type | --community community
  | --v3user user | --auth auth_protocol | --authpassword password | --
  priv priv_protocol | --privpassword password]
```

```
efa snmp subscriber unregister [ --host ip ]
```

Parameters

--host *ip*

(Required) Specifies the IP address or host name of the SCVMM server you want to update.

--type *type*

(Required) v2c or v3.

--community *community*

(Optional) Community string for v2c trap type.

--v3user *user*

(Required for v3 type) v3 user when trap type is v3.

--auth *auth_protocol*

(Optional) v3 authentication protocol (MD5 or SHA1) when trap type is v3.

--authpassword *user name*

(Optional) v3 authentication password when trap type is v3. Default is "**<prompt>**".

--priv *priv_protocol*

(Optional) v3 privacy protocol (DES or AES) when trap type is v3.

--privpassword *password*

(Optional) v3 privacy password when trap type is v3. Default is "**<prompt>**".

Examples

This example shows how to add a version 2c trap receiver.

```
efa snmp subscriber add --host 10.20.241.86 --type v2c --community
  testcommunity
Register SNMP Trap forwarding
Successfully registered trap subscriber.
+-----+-----+
| attribute | value |
+-----+-----+
| id        | 1     |
+-----+-----+
| host      | 10.20.241.86 |
+-----+-----+
```

```

| type          | v2c          |
+-----+-----+
| community     | testcommunity |
+-----+-----+
| v3 user       | efaUser      |
+-----+-----+
| v3 Auth Type  |              |
+-----+-----+
| v3 Priv Type  |              |
+-----+-----+
Trap Subscriber Host=10.20.241.86

```

This example shows how to add an SMTP version 3 trap receiver.

```

efa snmp subscriber add --host 2.1.1.1 --type v3 --v3user efaUser
--auth SHA1 --authpassword testing --priv AES --privpassword testingPrivacy
Register SNMP Trap forwarding
Successfully registered trap subscriber.
+-----+-----+
| attribute    | value        |
+-----+-----+
| id           | 2            |
+-----+-----+
| host         | 2.1.1.1      |
+-----+-----+
| type         | v3           |
+-----+-----+
| community    |              |
+-----+-----+
| v3 user      | efaUser      |
+-----+-----+
| v3 Auth Type | SHA1         |
+-----+-----+
| v3 Priv Type | AES          |
+-----+-----+
Trap Subscriber
Host=2.1.1.1

```

This example lists currently registered trap subscribers.

```

efa snmp subscriber list
+-----+-----+-----+-----+-----+-----+-----+
| id | type | host   | community | user   | Auth Protocol | Privacy Protocol |
+-----+-----+-----+-----+-----+-----+-----+
| 64 | v3   | 2.1.1.1 |           | efaUser | SHA1          | AES              |
+-----+-----+-----+-----+-----+-----+
SNMP Trap Subscribers Count=2
--- Time Elapsed: 263.997115ms ---

```

This example deletes a trap subscriber.

```

efa snmp subscriber unregister --host 2.1.1.1
Delete SNMP Trap forwarding subscriber
Successfully unregistered trap subscriber

```

efa show-running-config

Displays the running-config of all current EFA configurations.

Syntax

```
efa show-running-config
```

Usage Guidelines

Run this command to display the running-config of current configurations for core services.

The output is displayed in the following order: Asset, Fabric, Tenant commands.

The output contains the default values for each configuration line item.

You can use the command output for CLI playback on an empty EFA deployment. The output is a useful tool for recovery.

This example shows a partial list of typical output.

```
$ efa show-running-config

efa inventory device register --ip "10.24.80.191" --username admin --password password

efa inventory device setting update --ip "10.24.80.191" --maint-mode-enable-on-reboot No
--maint-mode-enable No --health-check-enable No --health-check-interval 6m
--health-check-heartbeat-miss-threshold 2 --config-backup-periodic-enable Yes
--config-backup-interval 24h --number-of-config-backups 4

efa inventory device register --ip "10.24.80.192" --username admin --password password

efa inventory device setting update --ip "10.24.80.192" --maint-mode-enable-on-reboot No
--maint-mode-enable No --health-check-enable No --health-check-interval 6m
--health-check-heartbeat-miss-threshold 2 --config-backup-periodic-enable Yes
--config-backup-interval 24h --number-of-config-backups 4

efa fabric create --name "default" --type clos --stage 3 --description "Default Fabric"
```

efa status

Displays the role and the status of the nodes in a single-node or multi-node deployment.

Syntax

```
efa status
```

Usage Guidelines

Roles are either Active or Standby.

Examples

The following example shows output for a single node installation.

```
xmcddev@xmcddev-virtual-machine:~$ efa status
+-----+-----+-----+-----+
| Node Name          | Role  | Status | IP           |
+-----+-----+-----+-----+
| xmcddev-virtual-machine | active | up      | 10.20.255.114 |
+-----+-----+-----+-----+
--- Time Elapsed: 2.224282775s ---
```

The following example shows output for a multi-node installation.

```
(efa:extreme)extreme@tpvm:~$ efa status
+-----+-----+-----+-----+
| Node Name | Role   | Status | IP           |
+-----+-----+-----+-----+
| tpvm      | active | up     | 10.20.255.114 |
+-----+-----+-----+-----+
| tpvm2     | standby | up    | 10.20.255.115 |
+-----+-----+-----+-----+
--- Time Elapsed: 2.318350499s ---
```


efa system backup

Backs up the EFA system, including the database and certificates.

Syntax

```
efa system backup [ --device-ip stringArray | --remote | --fabric-all | --fabric-name string ]
```

Parameters

--device-ip *stringArray*

Specifies a comma-separated range of device IPs for which the configuration backup will be taken. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

--remote *boolean*

Specifies if backup should be stored remotely. Default is false (only on local). When set to true, backup is stored both locally and remotely.

--fabric-all

Indicates to back up all devices that are part of all fabrics running the configuration backup.

--fabric-name *string*

Indicates to back up all devices that are part of this fabric running the configuration backup.

Usage Guidelines

You can restore a backed-up database for various reasons, such as if the database becomes corrupted or you want to revert to a previous configuration. The backup process creates a backup tar file, which you specify for the restore process. The backup tar file is saved to one of the following locations:

- Server: `/var/log/efa/backup`
- TPVM: `/apps/efa_logs/backup`

Run this command as a sudo user for the base operating system.

Examples

This example performs system backup without device configuration backup.

```
efa system backup -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.735.tar
--- Time Elapsed: 5.741750131s ---
```

This example performs system backup by taking configuration backup of all devices that are part of the fabric specified.

```
efa system backup --fabric default -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.736.tar
--- Time Elapsed: 5.741750131s ---
```

This example performs system backup by taking configuration backup of all fabrics and its devices.

```
efa system backup --fabric-all -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.734.tar
--- Time Elapsed: 5.741750131s ---
```

This example performs system backup by taking configuration backup of all devices that are specified.

```
efa system backup --device-ip 10.20.1.2,10.20.1.3,10.20.1.4 -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.737.tar
--- Time Elapsed: 5.741750131s ---
```

This example shows the error message: Fabric does not exist.

```
efa system backup --fabric default -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.733.tar
Backup is partially success:
    Fabric does not exist
--- Time Elapsed: 5.741750131s ---
```

This example shows the error message: Device not found.

```
efa system backup --device-ip 10.20.1.5,10.20.1.6 -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.738.tar
Backup is partially success:
    Device 10.20.1.5 not found
--- Time Elapsed: 5.741750131s ---
```

This example shows the error message: Operation not allowed.

```
efa system backup --device-ip 10.20.1.2,10.20.1.3 -remote
Generating backup of EFA...
Backup Location: root@10.10.10.10:/tmp/efa/EFA-2021-04-07T03-04-48.739.tar
Backup is partially success:
    Devices [10.20.1.2] failed to get config backup as its locked for configuration
    change by process [Firmware download].
--- Time Elapsed: 5.741750131s ---
```

efa system backup-list

Displays a list of all available EFA backup files.

Syntax

```
efa system backup-list
```

Examples

This example shows two backup .tar files that are available to use in a restore operation.

```
efa system backup-list
+---+-----+-----+-----+
+-----+
| ID |           File           | Version  | Generated By |
Location      |
+---+-----+-----+-----+
+-----+
| 1  | EFA-2021-04-07T00-00-00.000.tar | 2.4.1-8799 | System
|           |
+---+-----+-----+-----+
+-----+
| 2  | EFA-2021-04-07T03-04-48.736.tar | 2.5.0-1   | User
10.10.10.10:/tmp/efa |
+---+-----+-----+-----+
+-----+
--- Time Elapsed: 16.263551ms ---
```

efa system cleanup

Cleans up archives.

Syntax

```
efa system cleanup [ --archive-type { backup | support } | --name  
  string ]
```

Parameters

--archive-type { **backup** | **support** }

Type of archive to delete. The default is backup.

--name *string*

Name of archive to delete.

efa system feature show

Displays the status of the automatic in-progress (in-flight) transaction recovery feature.

Syntax

```
efa system feature show
```

Usage Guidelines

Use the **efa system feature update** command to enable the feature.

For more information, see the "In-flight Transaction Recovery" topic in the [Extreme Fabric Automation Administration Guide, 2.7.0](#).

Examples

The following example shows that the feature is turned off.

```
# efa system feature show
+-----+-----+
|           FEATURE NAME           | STATUS |
+-----+-----+
| Inflight Transaction Auto Recovery | Disabled |
+-----+-----+
```

efa system feature update

Enables the automatic recovery of in-flight (in-progress) transactions after a service restart or high-availability failover.

Syntax

```
efa system feature update [--inflight-transaction-auto-recovery { enable | disable } | --tenant-api-concurrency { enable | disable }]
```

Parameters

```
--inflight-transaction-auto-recovery { enable | disable }
```

Turns on or turns off the automatic recovery of in-flight transactions after a service restart or high-availability failover.

```
--tenant-api-concurrency { enable | disable }
```

Turns on or turns off the tenant API concurrency feature.

Usage Guidelines

After a service restart or high-availability failover, EFA can recover in-progress transactions by rolling them back or rolling them forward. In-progress transactions are those that are outstanding in the execution log after a restart or a failover.

- When transactions are rolled back, the requested action is undone. This category includes transactions launched by create commands and by update commands with add operations.
- When transactions are rolled forward, the requested action is completed. This category includes transactions launched by delete commands and by update commands with delete operations.

When this feature is enabled, all Day-1 to Day-N tenant operations support automatic in-progress transaction recovery. This feature is enabled by default.

When this feature is turned off, Day-1 to Day-N tenant operations do not support automatic in-progress transaction recovery.

For more information about the feature, see the "In-flight Transaction Recovery" topic in the *Extreme Fabric Automation Administration Guide*.

Examples

This example enables automatic in-flight transaction recovery.

```
efa system feature update --inflight-transaction-auto-recovery enable
Feature Setting Updated Successful
--- Time Elapsed: 634.557118ms ---
```

This example turns off automatic in-flight transaction recovery.

```
efa system feature update --inflight-transaction-auto-recovery disable
Feature Setting Updated Successful
--- Time Elapsed: 634.557125ms ---
```

efa system restore

Restores the EFA system from a backup.

Syntax

```
efa system restore [ --backup-tar filename]
```

Parameters

--backup-tar *filename*

Specifies the file name for the backup tar file.

Usage Guidelines

You can restore a backed-up database for various reasons, such as if the database becomes corrupted or you want to revert to a previous configuration. The backup process creates a backup tar file, which you specify for the restore process.

If you do not specify the **--backup-tar** parameter, the restore process displays the list of available backup files for you to choose from.

After the database is restored, a list of all pods is displayed, showing their status, number of restarts, and age.

Run this command as a sudo user for the base operating system.

Examples

This syntax restores the database.

```
efa system restore --backup-tar EFA-2020.06.08-15.52.09.tar
Backup file name: EFA-2020.06.08-15.52.09.tar
Restore result :: Restore operation successful
--- Time Elapsed: 3m11.269648985s ---
```

efa system settings reset

Resets system settings to default values.

Syntax

```
efa system settings reset [ --backup-schedule time | --max-backup-files |  
  --max-supportsave-files | --remote-server-settings settings | --  
  periodic-device-config-backup backup ]
```

Parameters

--backup-schedule *time*

The default time is scheduled for Sunday 12 AM.

--max-backup-files

The default is five backup files.

--max-supportsave-files

The default is five supportsave files.

--remote-server-settings *settings*

By specifying this flag, you set the following remote server settings to **null**: **remote-server-ip**, **remote-server-username**, **remote-server-password**, **remote-server-directory** and the **remote-server-protocol** setting to **scp**.

--periodic-device-config-backup *backup*

By specifying this flag, you turn off the backup.

Usage Guidelines

You must issue the **efa system settings reset** command with at least one flag. For the flags that you specify, the command resets values to their defaults.

Examples

The following command sets all the flags that you specify to their default values: --max-backup-files to 5, --max-supportsave-files to 5, --backup-schedule to 0 * * *, --remote-server-settings to null --periodic-device-config-backup to off.

```
efa system settings reset --max-backup-files --max-supportsave-files  
--backup-schedule --remote-server-settings --periodic-device-config-backup
```


efa system settings show

Displays the settings for the maximum number of backup and supportsave files to save, and the schedule for backing up the database.

Syntax

```
efa system settings show
```

Examples

This command prints the current (default) system settings.

```
$ efa system settings show
+-----+
|          SETTING          | VALUE |
+-----+-----+
| Max Backup File Limit    | 5     |
+-----+-----+
| Max Supportsave File Limit | 5     |
+-----+-----+
| Backup Schedule          | 0:*:*:* |
+-----+-----+
| Remote Server Ip         |        |
+-----+-----+
| Remote Server Username   |        |
+-----+-----+
| Remote Server Password   | ***** |
+-----+-----+
| Remote Server Directory  |        |
+-----+-----+
| Remote Transfer Protocol | scp    |
+-----+-----+
| Periodic Device Config Backup | Disabled |
+-----+-----+
--- Time Elapsed: 221.412325ms ---
```

efa system settings update

Configures the schedule for automatic database backups and determines the maximum number of backup and supportsave files to save before they are deleted.

Syntax

```
efa system settings update [--backup-schedule time | --max-backup-files num | --max-supportsave-files num | --remote-server-ip ip | --remote-transfer-protocol { ftp | scp } | --remote-server-username name | --remote-server-password password | --remote-server-directory directory | --periodic-device-config-backup { Enabled | Disabled } ]
```

Parameters

--backup-schedule *time*

Specifies the schedule for periodic, automatic backups, in the format HOUR:DOM:MON:DOW (hour, day of month, month, day of week). The default time is scheduled for Sunday 12 AM.

--max-backup-files *num*

Specifies the maximum number of backup files to keep. As new backup files are created, older backup files in excess of this value are deleted. The range of values is 2-20. The default is 5.

--max-supportsave-files *num*

Specifies the maximum number of supportsave files to keep. As new supportsave files are created, older files in excess of this value are deleted. The range of values is 2-20. The default is 5.

--remote-server-ip *ip*

Specifies the remote server IP where backup is copied.

--remote-transfer-protocol { **ftp** | **scp** }

Specifies a valid transfer protocol.

--remote-server-username *name*

Specifies the remote server user name.

--remote-server-password *password*

Specifies the remote server password. The default is "<prompt>".

--remote-server-directory *directory*

Specifies the remote server directory where backup is copied.

--periodic-device-config-backup { **Enabled** | **Disabled** }

Specifies whether to include device configuration backup in periodic backup. Default is No.

Usage Guidelines

The settings described here enable you to configure details of the remote server where the backup is copied.



Note

The details of the remote server are verified before configuring the remote parameters. If you specify a value for any **one** of the following, you must specify values for **all the others**:

- `remote-server-ip`
- `remote-server-username`
- `remote-server-password`
- `remote-server-directory`

If the four values have already been set, they are retained in the database. If you then modify any one of them, the other values are used.

The transfer of backup on remote servers is performed through the Secure Copy Protocol (SCP).

Examples

```
efa system settings update --remote-server-ip 10.20.255.177 --remote-server-username
extreme
--remote-server-password password --remote-server-directory efa/backup
Setting Update Successful
```

```
efa system settings update --remote-server-ip 10.20.255.177
Error: Remote server settings missing or incomplete
```

```
efa system settings update --remote-server-username extreme
Error: Remote server settings missing or incomplete
```

```
efa system settings update --remote-server-ip 10.20.241.7
--remote-server-username root --remote-server-password pass
--remote-server-directory /root/vinod/ --remote-transfer-protocol scp

Setting Update Successful

--- Time Elapsed: 148.800033ms ---
```

efa system supportsave

Collects the system supportsave of the Inventory, Tenant, and Fabric service logs, and their associated database dumps.

Syntax

```
efa system supportsave [--device-ip stringArray | --fabric-all | --  
fabric-name string ]
```

Parameters

--device-ip *stringArray*

Specifies a comma-separated range of device IPs to collect supportsave from. Example:
10.1.1.13,10.1.1.50,10.1.1.101.

--fabric-all

Indicates to collect supportsave from all devices.

--fabric-name *string*

Indicates to collect supportsave from all devices that are part of the specified fabric.

Usage Guidelines

Location of supportsave files:

- For TPVM deployments, the file is saved to `/apps/efa_logs/efa_<log_id>.logs.zip`.
- For non-TPVM deployments, the file is saved to `/var/log/efa/efa_<log_id>.logs.zip`.



Note

If SLX devices are part of the supportsave collection, the command triggers the **copy support** command (supportsave) on SLX devices. You must check the SLX devices to confirm that the supportsave copy is complete before running the command again. SLX supportsave can take up to 15 minutes to complete on each switch.

Examples

This example collects the system supportsave logs.

```
(efa)extreme@tpvm:~$ efa system supportsave  
SupportSave File Location: /apps/efa_logs/efa_2022-04-28T19-23-59.706.logs.zip  
--- Time Elapsed: 51.441277667s ---
```

efa tenant create

Creates tenant reserved resources such as the total number of Layer 3 VNIs, VLANs, VRFs, and bridge domains for fabrics with non-auto VNI settings, which can later be applied to an endpoint group.

Syntax

```
efa tenant create [ --name name | --description description | --type
  { private | shared } | --l2-vni-range range | --l3-vni-range range |
  --vlan-range range | --vrf-count vrfs | --enable-bd { true | false }
  | --port list of ports | --mirror-destination-port list of ports ]
```

Parameters

--name *name*

Specifies the name of the tenant.

--description *description*

Describes the tenant.

--type { **private** | **shared** }

Specifies whether the endpoints and VRFs of a shared tenant are available to other tenants. The default is private.

--l2-vni-range *range*

Specifies the contiguous range of Layer 2 VNIs in ascending order to be reserved for the tenant within the scope of a fabric. Valid VNI values are 1 through 16777215. Example: 14201-14300.

--l3-vni-range *range*

Specifies the contiguous range of Layer 3 VNIs in ascending order starting from l3-vni-range to be reserved for the tenant within the scope of a fabric. Valid VNI values are 1 through 16777215. Example: 14201-14300.

--vlan-range *range*

Specifies the contiguous range of VLANs to be reserved for the tenant. Valid values are 2 through 4090. Example: 2-100.

--vrf-count *vrf*s

Specifies the number of VRFs reserved for the tenant.

--enable-bd { **true** | **false** }

Enables bridge domain capability for networks created under this tenant. Use the following format: `--enable-bd=true` or `--enable-bd=false`.

--port *list of ports*

Lists physical ports of devices that are reserved for the tenant. Example:

```
SW1_IP[0/1,0/2,0/12-15,0/5:4],SW2_IP[0/1,0/4,0/5:1-2,0/5:3,0/9-20].
```

--mirror-destination-port *list of ports*

xxx

Examples

This example creates a VLAN-based tenant with manual VNI mapping.

```
$ efa tenant create --name tenant11 --l2-vni-range 10002-14190
--l3-vni-range 14191-14200 --vlan-range 2-4090 --vrf-count 10
--port 10.20.216.15[0/11-20],10.20.216.16[0/11-20] --description Subscriber1
```

Tenant created successfully.

--- Time Elapsed: 455.141597ms ---

This example creates a BD-based tenant.

```
$ efa tenant create --name tenant21 --l2-vni-range 30002-34190
--l3-vni-range 34191-34200 --vlan-range 2-4090 --vrf-count 10 --enable-bd
--port 10.20.216.15[0/21-28],10.20.216.16[0/21-28]
```

Tenant created successfully.

--- Time Elapsed: 501.176996ms ---

This example creates a tenant with auto-VNI with breakout ports.

```
$ efa tenant create --name tenant12 --vlan-range 2-100 --vrf-count 10
--port 10.20.216.103[0/1-10],10.20.216.104[0/1-5,0/6:1-4]
```

Tenant created successfully.

--- Time Elapsed: 427.73527ms ---

This example creates a shared tenant.

```
$ efa tenant create --name ST --type shared
--port 10.20.216.15[0/1-10],10.20.216.16[0/1-10]
```

Tenant created successfully.

--- Time Elapsed: 381.182892ms ---

efa tenant debug

Identifies drift in device configuration.

Syntax

```
efa tenant debug device drift [ --device-ip | --filter { po | evpn | ovg
    | lif | bd | vlan | intf | ve | vrf | mct | routerbgp | pwprofile |
    all } | --reconcile ]

efa tenant debug set [ --level { debug | info | debugdb | nodebugdb } ]

efa tenant debug device config sync [ --device-ip ip-addr ]
```

Parameters

--device-ip *ip-addr*

Specifies a comma-separated list of device IP addresses.

--filter { **po** | **evpn** | **ovg** | **lif** | **bd** | **vlan** | **intf** | **ve** | **vrf** | **mct** | **routerbgp** | **pwprofile** | **all** }

Filters the displayed draft data by the type of data you select. The default is all.

--level { **debug** | **info** | **debugdb** | **nodebugdb** }

Specifies the debug level.

--reconcile

Reconciles configuration on the device.

Usage Guidelines

You can also use the API to set up debug level and debug for the module.

Examples

This example identifies drift.

```
$ efa tenant debug device drift --device-ip 10.20.216.15
=====
Device           : 10.20.216.15

Operation succeeded.

--- Time Elapsed: 626.727571ms ---
```

This example reconciles the drifted configuration.

```
$ efa tenant debug device drift --device-ip 10.20.216.15 --reconcile
=====
Device           : 10.20.216.15
===== VLAN Drift =====
VLAN             : 10
App-state        : cfg-refreshed

Drifted Interface
```

```

+-----+-----+-----+
| Interface-Type | Interface-Name | App-State |
+-----+-----+-----+
| ethernet      | 0/11           | cfg-refreshed |
+-----+-----+-----+
===== Interface Drift =====

Drifted Interface
+-----+-----+-----+
| Interface-Type | Interface-Name | App-State |
+-----+-----+-----+
| ethernet      | 0/11           | cfg-in-sync |
+-----+-----+-----+

===== Reconciliation Status =====
+-----+-----+-----+
| CONFIG TYPE | STATUS | ERROR |
+-----+-----+-----+
| Port-Property | Success |      |
| Portchannel   | Success |      |
| Vlan          | Success |      |
+-----+-----+-----+

Operation succeeded.

--- Time Elapsed: 32.202412126s ---

```

This example synchronizes the tenant database with the inventory database.

```

$ efa tenant debug device config sync --device-ip 10.20.216.15

+-----+-----+-----+-----+
| IP Address | Role | State | Result |
+-----+-----+-----+-----+
| 10.20.216.15 | Leaf | READY | Device 10.20.216.15 is already present in Tenant |
|              |      |      | DB with READY state |
+-----+-----+-----+-----+

Device Details

--- Time Elapsed: 451.271557ms ---

```

This example sets the debug level.

```

$ efa tenant debug set --level debugdb

Level debugdb set successfully.

--- Time Elapsed: 116.43581ms ---

```


efa tenant delete

Deletes a tenant.

Syntax

```
efa tenant delete [ --name tenant-name | --force ]
```

Parameters

--name *tenant-name*

Specifies a tenant and a tenant name.

--force

Forces the deletion on the tenant if the option is provided.

Examples

The following example deletes a tenant.

```
(efa:extreme)extreme@node-1:~$ efa tenant delete --name tenant12
Tenant deleted successfully.
--- Time Elapsed: 1.972258324s ---
```

The following example deletes the tenant and all the related entities forcibly.

```
(efa:extreme)extreme@node-1:~$ efa tenant delete --name tenant12 --force
Tenant delete with force will delete associated Vrfs, EndpointGroups and PortChannels.
Do you want to proceed (Y/N): y
Tenant deleted successfully.
--- Time Elapsed: 3.258645739s ---
```

efa tenant epg configure

Pushes or removes a pending endpoint group configuration.

Syntax

```
efa tenant epg configure [ --name name | --tenant name ]
```

Parameters

--name *name*

Specifies the name of the endpoint group.

--tenant *name*

Specifies the name of the tenant.

Examples

Pushes or removes a pending endpoint group configuration.

```
$ efa tenant epg configure --name epg1 --tenant tenant11  
  
EndpointGroup configured successfully.  
  
--- Time Elapsed: 444.898897ms ---
```

efa tenant epg create

Creates a Layer 3 endpoint group.

Syntax

```
efa tenant epg create [--name epg-name | --tenant tenant-name | --
  description desc | --port ip-ethport | --po po-name | --switchport-mode
  { access | trunk | trunk-no-default-native } | --type { l3-hand-off |
  extension } | --switchport-native-vlan-tagging | --switchport-native-
  vlan value | --ctag-range range | --ctag description desc | --vrf vrf-
  name | --l3-vni vni | --l2-vni vni | --anycast-ip ipv4 | --anycast-
  ipv6 ipv6 | --local-ip ipv4 | --local-ipv6 ipv6 | --bridge-domain bd-
  name | --ipv6-nd-mtu mtu-value | --ipv6-nd-managed-config flag | --
  ipv6-nd-other-config other-flag | --ipv6-nd-prefix ipv6-prefix | --
  ipv6-nd-prefix-valid-lifetime lifetime | --ipv6-nd-prefix-preferred-
  lifetime pref-lifetime | --ipv6-nd-prefix-no-advertise | --ipv6-nd-
  prefix-config-type { no-autoconfig | no-onlink | off-link } | --
  single-homed-bfd-session-type { auto | software | hardware } | --ip-
  mtu mtu-value | --suppress-arp array | --suppress-nd array | --pp-
  mac-acl-in ext-mac-permit-any-mirror-acl | --pp-mac-acl-out ext-mac-
  permit-any-mirror-acl | --pp-ip-acl-in ext-ip-permit-any-mirror-acl |
  --pp-ip-acl-out ext-ip-permit-any-mirror-acl | --pp-ipv6-acl-in ext-
  ipv6-permit-any-mirror-acl | --np-mac-acl-in ctag:ext-mac-permit-any-
  mirror-acl | --np-mac-acl-out ctag:ext-mac-permit-any-mirror-acl | --
  np-ip-acl-in ctag:ext-ip-permit-any-mirror-acl | --np-ip-acl-out
  ctag:ext-ip-permit-any-mirror-acl | --np-ipv6-acl-in ctag:ext-ipv6-
  permit-any-mirror-acl | --help ]
```

Parameters

--name *epg-name*

Specifies the name of the endpoint group.

--tenant *tenant-name*

Specifies the name of the associated tenant.

--description *desc*

Describes the endpoint group.

--port *ip-ethport*

Specifies the device IP address and Ethernet port details. Example: SW1_IP[0/1], SW2_IP[0/5,0/6], SW3_IP[0/7-10]

--po *po-name*

Lists port channels. Example: po1, po2

--switchport-mode { **access** | **trunk** | **trunk-no-default-native** }

Configures switch port mode on the interfaces. The default is `trunk`.

--type { **l3-hand-off** | **extension** }

Configures the BGP service type. Valid values are l3-hand-off or extension. The default is extension.

--switchport-native-vlan-tagging

Enables the native VLAN characteristics on the ports of this endpoint group. Valid only if the switchport-mode parameter is set to trunk.

--switchport-native-vlan *value*

Configures native VLAN on the interfaces. Valid values are 2 through 4090, corresponding to the value of the ctag-range parameter.

--ctag-range *range*

Specifies the customer VLAN range in comma and hyphen separated format. Example: 2-20,30,40,50-55.

--ctag-description *desc*

Specifies a unique description of the ctag in the following format: ctag:12-vni.

--vrf *vrf-name*

Specifies the VRF to which these networks are attached.

--l3-vni *vni*

Specifies the Layer 3 VNI to be used for this VRF.

--l2-vni *vni*

Specifies the Layer 2 VNI to be used for this network in the following format: ctag:12-vni.

--anycast-ip *ipv4*

Specifies the IPv4 anycast address in the following format: ctag:anycast-ip.

--anycast-ipv6 *ipv6*

Specifies the IPv6 anycast address in the following format: ctag:anycast-ipv6.

--local-ip *ipv4*

Specifies the IPv4 local address in the following format: ctag,device-ip:local-ip.

--local-ipv6 *ipv6*

Specifies the IPv6 local address in the following format: ctag,device-ip:local-ipv6.

--bridge-domain *bd-name*

Specifies the bridge domain name in the following format: ctag:bridge-domain.

--ipv6-nd-mtu *mtu-value*

Sets the maximum transmission unit (MTU) for IPv6 neighbor discovery. Valid values range from 1280 through 65535. The format is ctag:mtu.

--ipv6-nd-managed-config *flag*

Sets the managed configuration flag for IPv6 router advertisement. The format is ctag:managedflag.

--ipv6-nd-other-config *other-flag*

Sets the other configuration flag for IPv6 router advertisement. The format is ctag:otherflag.

--ipv6-nd-prefix *ipv6-prefix*

Configures the IPv6 prefix address in the following format: `ctag:prefix1,prefix2`.

--ipv6-nd-prefix-valid-lifetime *lifetime*

Sets IPv6 prefix valid lifetime from 0 through 4294967295 in seconds. The format is `ctag,prefix:validTime`.

--ipv6-nd-prefix-preferred-lifetime *pref-lifetime*

Sets the IPv6 prefix preferred lifetime from 0 through 4294967295 in seconds. The format is `ctag,prefix:preferredTime`.

--ipv6-nd-prefix-no-advertise

Enables the prevention of prefix advertisement. The format is `ctag,prefix:noadvertiseflag`.

--ipv6-nd-prefix-config-type { **no-autoconfig** | **no-onlink** | **off-link** }

Sets the configuration type for the IPv6 prefix. The format is `ctag,prefix:configType`.

--single-homed-bfd-session-type { **auto** | **software** | **hardware** }

Specifies the BFD session type for the endpoint group. The default is `auto`, which means that the BFD session type is automatically determined based on the value of the `type` parameter: `extension` or `L3 hand-off`.

--ip-mtu *mtu-value*

Sets the IP maximum transmission unit (MTU) for the tenant network. Valid values range from 1280 through 9194. The format is `ctag:ip-mtu`.

--suppress-arp *value*

Sets `suppress-arp` flag to this network. The format is `ctag:suppress-arp`. Example: `1002:true`.

--suppress-nd *value*

Sets `suppress-nd` flag to this network. The format is `ctag:suppress-nd`. Example: `1002:true`.

--pp-mac-acl-in *ext-mac-permit-any-mirror-acl*

xxx

--pp-mac-acl-out *ext-mac-permit-any-mirror-acl*

xxx

--pp-ip-acl-in *ext-ip-permit-any-mirror-acl*

xxx

--pp-ip-acl-out *ext-ip-permit-any-mirror-acl*

xxx

--pp-ipv6-acl-in *ext-ipv6-permit-any-mirror-acl*

xxx

--np-mac-acl-in *ctag:ext-mac-permit-any-mirror-acl*

xxx

--np-mac-acl-out *ctag:ext-mac-permit-any-mirror-acl*

xxx

```

--np-ip-acl-in ctag:ext-ip-permit-any-mirror-acl
XXX
--np-ip-acl-out ctag:ext-ip-permit-any-mirror-acl
XXX
--np-ipv6-acl-in ctag:ext-ipv6-permit-any-mirror-acl
XXX

```

Usage Guidelines

An empty endpoint group has no network-policy, network-property, or port-property.

An endpoint group can be created with a port-property but without a port-group. However, an endpoint group cannot be created with a port-group but without a port-property.

ARP suppression is enabled for all the possible broadcast domains, VLAN or BD, on the device.

CEP is handled by replicating all the tenant configuration on the MCT neighbor except for the endpoint configuration, because the endpoint does not exist on the MCT neighbor.

Event handling sets the corresponding tenant networks to the `cfg-refreshed` state. However, there is no way to re-push the refreshed configuration onto the devices.

The value of `--single-homed-bfd-session-type` is configured for one endpoint group and then propagated to all Ethernet and single-homed port channel interfaces defined for that endpoint group.

EFA does not distinguish between SRIOV (single-root input/output virtualization) and non-SRIOV connections. Therefore, it treats both connections the same way. If you want to use hardware-based BFD sessions for CEP non-SRIOV connections, then create an endpoint group that contains all the CEP non-SRIOV connections and set the `--single-homed-bfd-session-type` to `hardware`.

You use the `--ip-mtu` parameter to configure the Maximum Transmission Unit (MTU) for the tenant network. This value is then configured on the interface VE on the SLX device. The output of the `efa tenant epg show --detail` command includes the configured `--ip-mtu <mtu-value>`.

Examples

This example creates a VLAN-based Layer 3 endpoint group.

```

$ efa tenant epg create --name epg1 --tenant tenant11 --switchport-mode trunk
--switchport-native-vlan 10 --switchport-native-vlan-tagging --port
10.20.216.15[0/11],10.20.216.16[0/11]
--po po1 --vrf blue11 --ctag-range 10 --l2-vni 10:10010 --l3-vni 14191 --anycast-ip
10:10.10.10.1/24
--anycast-ipv6 10:10::1/125 --local-ip 10,10.20.216.15:1.1.10.3/28 --local-ip
10,10.20.216.16:1.1.10.4/28
--local-ipv6 10,10.20.216.15:10a:10::3/125 --local-ipv6 10,10.20.216.16:10a:10::4/125 --
ipv6-nd-mtu 10:9000
--ipv6-nd-prefix 10:1002::/125,1003::/125,1004::/125 --ipv6-nd-prefix-valid-lifetime
10,1002::/125:infinite
--ipv6-nd-prefix-preferred-lifetime 10,1002::/125:1020304 --ipv6-nd-prefix-valid-lifetime
10,1003::/125:1020304
--ipv6-nd-prefix-preferred-lifetime 10,1003::/125:1020304 --ipv6-nd-prefix-valid-lifetime
10,1004::/125:1020304

```

```
--ipv6-nd-prefix-preferred-lifetime 10,1004::/125:infinite --ipv6-nd-prefix-config-type
10,1004::/125:no-onlink
--ipv6-nd-prefix-config-type 10,1003::/125:off-link --ipv6-nd-prefix-config-type
10,1002::/125:no-autoconfig
--ipv6-nd-managed-config 10:true --ipv6-nd-other-config 10:true --ctag-description
10:Network-10
```

EndpointGroup created successfully.

--- Time Elapsed: 16.922083265s ---

This example creates a VLAN-based L3-hand-off endpoint group.

```
$ efa tenant epg create --tenant tenant11 --name epg2
--type l3-hand-off --switchport-mode trunk --port 10.20.216.15[0/18],10.20.216.16[0/18] --
po po2
--vrf blue11 --ctag-range 12 --l2-vni 12:10012 --l3-vni 14191 --local-ipv6
12,10.20.216.16:10:12a::1/127
--local-ipv6 12,10.20.216.15:10:12a::2/127 --local-ip 12,10.20.216.16:1.1.12.1/29
--local-ip 12,10.20.216.15:1.1.12.2/29
```

EndpointGroup created successfully.

--- Time Elapsed: 8.605943783s ---

This example creates a bridge-domain-based Layer 3 endpoint group.

```
$ efa tenant epg create --tenant tenant21 --name epg3 --type extension
--switchport-mode trunk --po po11 --ctag-range 1002 --bridge-domain 1002:Net-30002 --l2-
vni 1002:30002
--vrf red11 --anycast-ip 1002:10.20.30.1/24
```

EndpointGroup created successfully.

--- Time Elapsed: 13.469697138s ---

This example creates a VLAN-based Layer 2 endpoint group.

```
$ efa tenant epg create --name epg4 --tenant tenant11
--ctag-range 101-103 --switchport-mode trunk-no-default-native --port 10.20.216.15[0/17]
```

EndpointGroup created successfully.

--- Time Elapsed: 19.83265s ---

This example creates an endpoint group for which the BFD session type is automatically determined.

```
$ efa tenant epg create --name epg5 --tenant tenant11 --port 10.20.216.15[0/11]
,10.20.216.16[0/11] --po po1 --switchport-mode trunk --single-homed-bfd-session-type auto
```

This example creates an endpoint group with MTU values for Ctag 11 and Ctag 12.

```
$ efa tenant epg create --name tenlepg1 --tenant ten1 --port 10.20.246.17[0/1],
10.20.246.18[0/1] --switchport-mode trunk --ctag-range 11-12 --anycast-ip11:10.0.11.1/24
--anycast-ip12:10.0.12.1/24 --anycast-ipv6 11:11::1/127 --anycast-ipv6 12:12::1/127
--vrf tenlvrf1 --ip-mtu 11:7900 --ip-mtu 12:8900
```

efa tenant epg delete

Deletes an endpoint group.

Syntax

```
efa tenant epg delete [ --force | --name epg-name | --tenant tenant-name ]
```

Parameters

--force

Forces the deletion of any endpoint group that is tied to the tenant.

--name *epg-name*

Specifies the name of the endpoint group or a comma-separated list of names to be deleted.

Example: `EPG-1` or `EPG-1,EPG2,EPG3`.

--tenant *tenant-name*

Specifies the tenant name.

Usage Guidelines

Before deleting a tenant, delete any endpoint group that is tied to the tenant to remove any port level configurations on the device that are defined by the endpoint group. Failure to do this causes EFA to return an error. Using the `-force` option overrides this error and deletes the underlying endpoint group configuration on the device.

Examples

The following example deletes a specific EPG.

```
(efa:extreme)extreme@node-1:~$ efa tenant epg delete --name epg2 --tenant tenant11
EndpointGroup: epg2 deleted successfully.
--- Time Elapsed: 5.853148185s ---
```

The following example deletes a specific EPG forcefully and ignores any error.

```
(efa:extreme)extreme@node-1:~$ efa tenant delete --name tenant12 --force
Tenant delete with force will delete associated Vrfs, EndpointGroups and PortChannels.
Do you want to proceed (Y/N): y
Tenant deleted successfully.
--- Time Elapsed: 3.258645739s ---
```


efa tenant epg detach

Splits an endpoint group into multiple endpoint groups.

Syntax

```
efa tenant epg detach [ --tenant name | --source-epg epg-name | --destination-epg epg-name | --destination-epg-description | --destination-epg-ctag-range ctag-range ]
```

Parameters

- tenant** *name*
Specifies the tenant name.
- source-epg** *epg-name*
Specifies the name of the source endpoint group.
- destination-epg** *epg-name*
Specifies the name of the destination endpoint group.
- destination-epg-description**
Displays destination endpoint group description.
- destination-epg-ctag-range** *ctag-range*
Specifies the destination endpoint group CTAG range.

Examples

The following example shows an existing endpoint group, splits it into multiple endpoint groups, and then shows the individual endpoint groups.

```
$ efa tenant epg show --name epg4 --tenant tenant11
+-----+-----+-----+-----+-----+-----+
+-----+
| Name | Tenant | Type | Ports | PO | SwitchPort |
Native Vlan |
| | | | | | Mode |
Tagging |
+-----+-----+-----+-----+-----+-----+
+-----+
| epg4 | tenant11 | extension | 10.20.216.15[0/17] | | trunk-no-default-native |
false |
| | | | | | |
| | | | | | |
+-----+-----+-----+-----+-----+-----+
+-----+
+-----+-----+-----+-----+
| Ctag Range | Vrf | L3Vni | State |
+-----+-----+-----+-----+
| 101-103 | | | | |
| | | | | |
+-----+-----+-----+-----+
EndpointGroup Details
```

```

PO/Vrf [Flags : * - Unstable]

For 'unstable' entities, run 'efa tenant po/vrf show' for details

--- Time Elapsed: 156.85406ms ---

$ efa tenant epg detach --tenant tenant11 --source-epg epg4
--destination-epg epg5 --destination-epg-ctag-range 103

EndpointGroup: epg4 split successfully.

--- Time Elapsed: 263.570996ms ---

$ efa tenant epg show --name epg4 --tenant tenant11
+-----+-----+-----+-----+-----+-----+
+-----+
| Name | Tenant | Type | Ports | PO | SwitchPort | Native |
| Vlan | | | | | | Vlan |
| Tagging | | | | | | Mode |
+-----+-----+-----+-----+-----+-----+
+-----+
| epg4 | tenant11 | extension | 10.20.216.15[0/17] | | trunk-no-default-native |
| false | | | | | |
| | | | | | |
+-----+-----+-----+-----+-----+-----+
+-----+
+-----+-----+-----+-----+
| Ctag Range | Vrf | L3Vni | State |
| | | | |
+-----+-----+-----+-----+
| 101-102 | | | |
| | | | |
+-----+-----+-----+-----+
EndpointGroup Details

PO/Vrf [Flags : * - Unstable]

For 'unstable' entities, run 'efa tenant po/vrf show' for details

--- Time Elapsed: 155.775375ms ---
$ efa tenant epg show --name epg5 --tenant tenant11
+-----+-----+-----+-----+-----+-----+
+-----+
| Name | Tenant | Type | Ports | PO | SwitchPort | Native |
| Vlan | | | | | | Vlan |
| Tagging | | | | | | Mode |
+-----+-----+-----+-----+-----+-----+
+-----+
| epg5 | tenant11 | | 10.20.216.15[0/17] | | trunk-no-default-native |
| false | | | | | |
| | | | | | |
+-----+-----+-----+-----+-----+-----+
+-----+
+-----+-----+-----+-----+
| Ctag Range | Vrf | L3Vni | State |
| | | | |
+-----+-----+-----+-----+
| 103 | | | |
| | | | |

```

```
|          |      |          |
+-----+-----+-----+
EndpointGroup Details

PO/Vrf [Flags : * - Unstable]

For 'unstable' entities, run 'efa tenant po/vrf show' for details

--- Time Elapsed: 154.028719ms ---
```

efa tenant epg error show

Reports errors related to the creation or update of endpoint groups.

Syntax

```
efa tenant epg error show [ --name epg-name | --tenant tenant-name | --summary ]
```

Parameters

--name *epg-name*

Specifies the name of the endpoint group.

--tenant *tenant-name*

Specifies the tenant name.

--summary

Displays a summary of endpoint groups in the error state.

Examples

```
efa tenant epg error show --name epg_nonclos --tenant tenant1
=====
Name : epg_nonclos
Tenant : tenant1
Errors
-----
MgmtIp Ctag NetworkPolicyErrorList
-----
10.25.225.173 2001 Adding/Removing VLANs under MCT Cluster failed due to
netconf rpc [error] '"2001:2002:2003" is an invalid value.'
-----
10.25.225.173 2002 Adding/Removing VLANs under MCT Cluster failed due to
netconf rpc [error] '"2001:2002:2003" is an invalid value.'
-----
10.25.225.173 2003 Adding/Removing VLANs under MCT Cluster failed due to
netconf rpc [error] '"2001:2002:2003" is an invalid value.'
-----
10.25.225.174 2001 Adding/Removing VLANs under MCT Cluster failed due to
netconf rpc [error] '"2001:2002:2003" is an invalid value.'
-----
10.25.225.174 2002 Adding/Removing VLANs under MCT Cluster failed due to
netconf rpc [error] '"2001:2002:2003" is an invalid value.'
-----
10.25.225.174 2003 Adding/Removing VLANs under MCT Cluster failed due to
netconf rpc [error] '"2001:2002:2003" is an invalid value.'
-----
```

efa tenant epg show

Shows brief or detailed information for all endpoint groups or a specific endpoint group.

Syntax

```
efa tenant epg show [ --name epg-name | --detail | --tenant tenant-name |
  --vrf vrf-name ]
```

Parameters

--name *epg-name*

Specifies the name of the endpoint group.

--detail

Displays all attributes of the specified endpoint group or all endpoint groups. When this parameter is not used, only brief tabular output is returned.

--tenant *tenant-name*

Specifies the name of the tenant.

--vrf *vrf-name*

Specifies the name of the VRF.

Usage Guidelines

The following table lists the details that this command can return. Your details may vary, depending on the deployment.

Information Type	Brief Output	Detailed Output
Name	X	X
Tenant	X	X
Type	X	X
Ports	X	X
Switchport mode	X	X
Native VLAN tagging	X	X
CTAG range	X	X
VRF	X	X
L3 VNI	X	X
State	X	X
Dev state		X
App state		X
Port property states		X

Information Type	Brief Output	Detailed Output
Network property flags		X
IPv6 ND prefix flags		X

Examples

Use this syntax to return brief details for all configured endpoint groups.

```
$ efa tenant epg show
```

This syntax returns full details for all configured endpoint groups.

```
$ efa tenant epg show --detail
```

This syntax returns brief details for the specified endpoint group.

```
$ efa tenant epg show --name epg2
```

This syntax returns full details for the specified endpoint group and tenant.

```
$ efa tenant epg show --name epg3 --tenant tenant21 --detail
```

This syntax returns full details for all endpoint groups for the specified tenant.

```
$ efa tenant epg show --tenant tenant21 --detail
```

efa tenant epg update

Updates an endpoint group.

Syntax

```
efa tenant epg update [--name epg-name | --tenant tenant-name | --
  operation { port-group-add | port-group-delete | ctag-range-add |
  ctag-range-delete | vrf-add | vrf-delete | local-ip-add | local-ip-
  delete | anycast-ip-add | anycast-ip-delete } | --port ip-ethport | --
  po po-name | --switchport-mode { access | trunk | trunk-no-default-
  native } | --switchport-native-vlan-tagging | --switchport-native-vlan
  value | --ctag-range range | --ctag-description desc | --vrf vrf-name
  | --l3-vni vni | --l2-vni vni | --anycast-ip ipv4 | --anycast-ipv6
  ipv6 | --local-ip ipv4 | --bridge-domain bd-name | --single-homed-
  bfd-session-type { auto | software | hardware } | --ip-mtu mtu-value
  | --suppress-arp array | --suppress-nd array | --pp-mac-acl-in ext-
  mac-permit-any-mirror-acl | --pp-mac-acl-out ext-mac-permit-any-
  mirror-acl | --pp-ip-acl-in ext-ip-permit-any-mirror-acl | --pp-ip-
  acl-out ext-ip-permit-any-mirror-acl | --pp-ipv6-acl-in ext-ipv6-
  permit-any-mirror-acl | --np-mac-acl-in ctag:ext-mac-permit-any-
  mirror-acl | --np-mac-acl-out ctag:ext-mac-permit-any-mirror-acl | --
  np-ip-acl-in ctag:ext-ip-permit-any-mirror-acl | --np-ip-acl-out
  ctag:ext-ip-permit-any-mirror-acl | --np-ipv6-acl-in ctag:ext-ipv6-
  permit-any-mirror-acl | --help ]
```

Parameters

--name *epg-name*

Specifies the name of the endpoint group.

--tenant *tenant-name*

Specifies the name of the associated tenant.

--operation { **port-group-add** | **port-group-delete** | **ctag-range-add** | **ctag-range-delete** | **vrf-add**, **vrf-delete** | **local-ip-add** | **local-ip-delete** | **anycast-ip-add** | **anycast-ip-delete** }

Specifies the operation to be performed.

--port *ip-ethport*

Specifies the device IP address and Ethernet port details. Example: SW1_IP[0/1], SW2_IP[0/5,0/6], SW3_IP[0/7-10]

--po *po-name*

Lists port channels. Example: po1, po2.

--switchport-mode { **access** | **trunk** | **trunk-no-default-native** }

Configures switch port mode on the interfaces. The default is `trunk`.

--switchport-native-vlan-tagging

Enables the native VLAN characteristics on the ports of this endpoint group. Valid only if the **switchport-mode** flag is set to `trunk`.

--switchport-native-vlan *value*

Configures native VLAN on the interfaces. Valid values are 2 through 4090 corresponding to the value of the **ctag-range** parameter.

--ctag-range *range*

Specifies the customer VLAN range in comma and hyphen separated format. Example: 2-20,30,40,50-55.

--ctag-description *desc*

Specifies a unique description of the ctag in the following format: `ctag:12-vni`.

--vrf *vrf-name*

Specifies the VRF to which these networks are attached.

--l3-vni *vni*

Specifies the Layer 3 VNI to be used for this VRF.

--l2-vni *vni*

Specifies the Layer 2 VNI to be used for this network in the following format: `ctag:12-vni`.

--anycast-ip *ipv4*

Specifies the IPv4 anycast address in the following format: `ctag:anycast-ip`.

--anycast-ipv6 *ipv4*

Specifies the IPv6 local address in the following format: `ctag,device-ip:local-ipv6`.

--local-ip *ipv4*

Specifies the IPv4 local address in the following format: `ctag,device-ip:local-ip`.

--bridge-domain *bd-name*

Specifies the bridge domain name in the following format; `ctag:bridge-domain`.

--single-homed-bfd-session-type { **auto** | **software** | **hardware** }

Specifies the BFD session type for the endpoint group. The default is `auto`, which means that the BFD session type is automatically determined based on the value of the **--type** parameter: `extension` or `L3 hand-off`.

--ip-mtu *mtu-value*

Sets the IP maximum transmission unit (MTU) for the tenant network. Valid values range from 1280 through 9194. The format is `ctag:ip-mtu`.

--suppress-arp *array*

Sets suppress-arp flag to this network. Format **ctag:suppress-arp**. Example: **1002:true**.

--suppress-nd *array*

Sets suppress-nd flag to this network. Format **ctag:suppress-arp**. Example: **1002:true**.

--pp-mac-acl-in *ext-mac-permit-any-mirror-acl*

xxx

--pp-mac-acl-out *ext-mac-permit-any-mirror-acl*

xxx


```
--pp-ip-acl-in ext-ip-permit-any-mirror-acl
XXX
--pp-ip-acl-out ext-ip-permit-any-mirror-acl
XXX
--pp-ipv6-acl-in ext-ipv6-permit-any-mirror-acl
XXX
--np-mac-acl-in ctag:ext-mac-permit-any-mirror-acl
XXX
--np-mac-acl-out ctag:ext-mac-permit-any-mirror-acl
XXX
--np-ip-acl-in ctag:ext-ip-permit-any-mirror-acl
XXX
--np-ip-acl-out ctag:ext-ip-permit-any-mirror-acl
XXX
--np-ipv6-acl-in ctag:ext-ipv6-permit-any-mirror-acl
XXX
```

Usage Guidelines

An empty endpoint group has no network-policy, network-property, or port-property.

An endpoint group can be created with a port-property and without a port-group. But an endpoint group cannot be created with a port-group and without a port-property.

ARP suppression is enabled for all the possible broadcast domains VLAN or BD on the device.

CEP is handled by replicating all the tenant configuration on the MCT neighbor except for the endpoint configuration, since the endpoint does not exist on the MCT neighbor.

The update operation for a bridge domain-based endpoint group is similar to that of a VLAN-based endpoint group. During a port-group add or delete operation, the logical interface configurations will be created or deleted for the existing ctags, and the corresponding bridge-domains.

During a ctag-range-add or delete operation, the logical interface and bridge-domain configurations are updated on the endpoint group.

During vrf-add or delete operation, the corresponding Layer 3 configurations are added to or deleted from the endpoint group.

Event handling sets the corresponding tenant networks to the cfg-refreshed state. However, there is no way to re-push the refreshed configuration onto the devices.

The value of --single-homed-bfd-session-type is configured for one endpoint group and then propagated to all Ethernet and single-homed port channel interfaces defined for that endpoint group.

EFA does not distinguish between SRIOV (single-root input/output virtualization) and non-SRIOV connections. Therefore, it treats both connections the same way. If you want to use hardware-based BFD sessions for CEP non-SRIOV connections, then create an endpoint group that contains all the CEP non-SRIOV connections and set the `--single-homed-bfd-session-type` to `hardware`.

During `vrf-add` and `ctag-range-add` operations, you can use the `--ip-mtu` parameter to configure the MTU for the tenant network. This value is then configured on the interface VE on the SLX device. The output of the `efa tenant epg show --detail` command includes the configured `--ip-mtu <mtu-value>`.

Examples

The following example adds a port to the endpoint group.

```
$ efa tenant epg update --name epg1
--tenant tenant11 --operation port-group-add --port 10.20.216.15[0/20]

EndpointGroup updated successfully.

--- Time Elapsed: 32.208253521s ---
```

The following example adds a Ctag with network properties to endpoint group.

```
$ efa tenant epg update --name epg1 --tenant tenant11
--operation ctag-range-add --ctag-range 100 --anycast-ip 100:1.1.100.1/24
--local-ip 100,10.20.216.15:100.100.1.1/28

EndpointGroup updated successfully.

--- Time Elapsed: 37.428381252s ---
```

The following example adds a automatic BFD session type to an endpoint group.

```
$ efa tenant epg update --name epg5 --tenant tenant11 --operation port-group-add
--port 10.20.216.15[0/11],10.20.216.16[0/11] --po po1 --switchport-mode trunk
--single-homed-bfd-session-type auto
```

The following example configures the MTU during a `vrf-add` operation.

```
$ efa tenant epg update --name tenlepg1 --tenant ten1 --operation vrf-add
--anycast-ip11:10.0.11.1/24 --anycast-ipv6 11:11::1/127 --vrf ten1vrf1 --ip-mtu 11:5990
```

The following example configures the MTU during a `ctag-range-add` operation.

```
$ efa tenant epg update --name tenlepg1 --tenant ten1 --operation ctag-range-add
--ctag-range 12 --anycast-ip12:10.0.12.1/24 --anycast-ipv6 12:12::1/127 --ip-mtu 12:6990
```

efa tenant execution

Displays the list of runs and event histories and deletes entries older than the specified number of days.

Syntax

```
efa tenant execution delete [--days days ]
efa tenant execution show [ --limit runs | --status { failed | succeeded
  | all | incomplete | reconcile-failed | reconcile-succeeded } | --id
  id ]
efa tenant execution show-event [--device ip-addr | --execution-id uuid ]
efa tenant execution delete [ --days value ]
```

Parameters

--days *days*

Deletes run entries older than the specified days (default 30).

--id *id*

Filter the executions based on the ID. The **limit** and **status** flags are ignored when the **id** flag is specified.

--device *ip-addr*

Filters run entries on IP address.

--limit *runs*

Limits the number of runs to be listed. Value "0" will list all the runs. Default is 10.

--status { **failed** | **succeeded** | **all** | **incomplete** | **reconcile-failed** | **reconcile-succeeded** }

Filters the runs based on the status. Default is **all**.

--execution-id *uuid*

Filters run entries on run UUID.

The following example returns the following information for all statuses: ID, entity, command, status, start time, end time, and user name.

```
$ efa tenant execution show
```

The following example shows the list of all event histories.

```
$ efa tenant execution show-event
```

The following example deletes run entries older than one day.

```
$ efa tenant execution delete --days 1
```

efa tenant po configure

Pushes or removes pending port channel configuration.

Syntax

```
efa tenant po configure [ --name po-name | --tenant name ]
```

Parameters

--name *po-name*

Specifies the name of the port channel.

--tenant *name*

Specifies the name of the tenant.

This example pushes or removes pending port channel configuration.

```
$ efa tenant po configure --name po1 --tenant tenant11
```

```
PortChannel: po1 configured successfully.
```

efa tenant po create

Creates a port channel for a tenant.

Syntax

```
efa tenant po create [ --name po-name | --tenant tenant-name | --description description | --speed speed | --negotiation { active | passive | static } | --port port-info | --min-link-count num-links | --number interface_number | --lACP-timeout { long | short } | --mtu int32 ]
```

Parameters

--name *po-name*

Specifies the port channel name.

--tenant *tenant-name*

Specifies the tenant name.

--description *description*

Describes the port channel.

--speed *speed*

Configures the speed for the port channel and its member ports. Valid values are 100Mbps, 1Gbps, 10Gbps, 25Gbps, 40Gbps, and 100Gbps.

--negotiation { **active** | **passive** | **static** }

Configures LACP Negotiation mode for a port channel.

--port *port-info*

Specifies the device IP address and Ethernet port details. Example: SW1_IP[0/1],SW2_IP[0/5]

--min-link-count *num-links*

Specifies the least number of links that need to be operationally up to declare the port channel as operationally up. Valid values are 1 through 64. The default value is 1.

--number *interface_number*

Specifies the port channel interface number generated by the service.

--lACP-timeout { **long** | **short** }

Specifies the length of the timeout.

--mtu **int32**

MTU configuration for the port channel. Valid values are 1500 through 9216.

Examples

This example creates a dual-homed PO.

```
(efa:extreme)extreme@node-1:~$ efa tenant po create --tenant tenant11 --name po1
--speed 100Gbps --negotiation active --lACP-timeout short --port 10.20.216.15[0/12-13],
10.20.216.16[0/12-13]
--min-link-count 2
```

```
PortChannel created successfully.
```

This example creates a single-homed PO.

```
((efa:extreme)extreme@node-1:~$ efa tenant po create --tenant "tenant11" --name "po2"  
--speed 10Gbps --negotiation static --port 10.20.216.15[0/15] --min-link-count 1  
--description po2
```

```
PortChannel created successfully.
```

```
--- Time Elapsed: 3.894884521s ---
```

efa tenant po delete

Deletes a port channel.

Syntax

```
efa tenant po delete [ --name name | --force | --tenant name ]
```

Parameters

--name *name*

Specifies port channel name or comma-separated port channel names. For example: po1 or po1,po2,po3.

--force

Forces the port channel deletion if the option is provided.

--tenant *name*

Specifies the tenant name.

Examples

This example deletes the specified port channels.

```
$ efa tenant po delete --name po1,po2 --tenant tenant11
```

```
PortChannel: po1 deleted successfully.
```

```
PortChannel: po2 deleted successfully.
```

This example deletes a port channel even when it is associated with an EPG.

```
$ efa tenant po delete --name po1 --tenant tenant11 --force
```

```
PortChannel Delete with force will update associated EndpointGroups and  
Networks and deletes them if there are no other ports associated to them (N/Y): y
```

```
PortChannel: po1 deleted successfully.
```

efa tenant po show

Shows brief or detailed output of the port channel of all tenants, a given tenant, or a given port channel.

Syntax

```
efa tenant po show [--name po-name | --tenant tenant-name | --detail ]
```

Parameters

--name *po-name*

Specifies the port channel name.

--tenant *tenant-name*

Specifies the tenant name.

--detail

Displays detailed output of the port channel. When this parameter is not used, only brief tabular output is returned.

Examples

This example shows brief output of all port channels.

```
$ efa tenant po show
+-----+-----+-----+-----+-----+-----+-----+
| Name | Tenant | ID | Description | Speed | Negotiation | MinLinkCount |
+-----+-----+-----+-----+-----+-----+-----+
| po1 | tenant11 | 1 | EFA Port-channel po1 | 100Gbps | active | 2 |
| | | | | | | |
+-----+-----+-----+-----+-----+-----+-----+
| po2 | tenant11 | 2 | po2 | 10Gbps | static | 1 |
| | | | | | | |
+-----+-----+-----+-----+-----+-----+-----+
| po11 | tenant21 | 3 | EFA Port-channel po11 | 25Gbps | active | 1 |
| | | | | | | |
+-----+-----+-----+-----+-----+-----+-----+
| Ports | LacpTimeout | State | Dev-State | App-State |
+-----+-----+-----+-----+-----+-----+
| 10.20.216.15[0/12-13] | short | po-created | provisioned | cfg-in-sync |
| 10.20.216.16[0/12-13] | | | | |
+-----+-----+-----+-----+-----+-----+
| 10.20.216.15[0/15] | | po-created | provisioned | cfg-in-sync |
| 10.20.216.16[0/15] | | | | |
+-----+-----+-----+-----+-----+-----+
| 10.20.216.15[0/22] | short | po-created | provisioned | cfg-in-sync |
| 10.20.216.16[0/22] | | | | |
+-----+-----+-----+-----+-----+-----+

PortChannel Details

--- Time Elapsed: 832.496716ms ---
```

This example shows detailed output of all port channels.

```
$ efa tenant po show --detail
=====
Name : po1
```



```
Tenant      : tenant11
ID          : 1
Description : EFA Port-channel po1
Speed      : 100Gbps
Negotiation : active
Min Link Count : 2
Lacp Timeout : short
Ports      : 10.20.216.15[0/12-13]
           : 10.20.216.16[0/12-13]

State      : po-created
Dev State  : provisioned
App State  : cfg-in-sync

=====
Name       : po2
Tenant    : tenant11
ID        : 2
Description : EFA Port-channel po3
Speed     : 10Gbps
Negotiation : static
Min Link Count : 1
Lacp Timeout :
Ports     : 10.20.216.15[0/15]
           : 10.20.216.16[0/15]

State     : po-created
Dev State : provisioned
App State : cfg-in-sync

=====
Name       : po11
Tenant    : tenant21
ID        : 3
Description : EFA Port-channel po11
Speed     : 25Gbps
Negotiation : active
Min Link Count : 1
Lacp Timeout : short
Ports     : 10.20.216.15[0/22]
           : 10.20.216.16[0/22]

State     : po-created
Dev State : provisioned
App State : cfg-in-sync

=====

--- Time Elapsed: 506.117046ms ---
```

This example shows brief output of a specific port channel.

```
$ efa tenant po show --tenant tenant11 --name po1
+-----+-----+-----+-----+-----+-----+
| Name | Tenant | ID | Speed | Negotiation | Min Link | Lacp |
|      |        |   |       |              | Count   | Timeout |
+-----+-----+-----+-----+-----+-----+
| po1  | tenant11 | 1 | 100Gbps | active      | 1       | short  |
|      |          |   |         |             |         |        |
+-----+-----+-----+-----+-----+-----+
| Lacp |        |    |        | State      | Dev State | App State |
| Timeout |      |   |       |            |           |           |
+-----+-----+-----+-----+-----+-----+
| short | 10.20.216.15[0/12] | | po-created | provisioned | cfg-in-sync |
|      | 10.20.216.16[0/12] | |           |           |           |
```

```
+-----+-----+-----+-----+-----+
PortChannel Details
--- Time Elapsed: 150.30883ms ---
```

This example shows detailed output of all port channels belonging to a tenant.

```
$ efa tenant po show --tenant tenant21 --detail
=====
Name           : po11
Tenant         : tenant21
ID             : 3
Description    : EFA Port-channel po11
Speed         : 25Gbps
Negotiation    : active
Min Link Count : 1
Lacp Timeout  : short
Ports         : 10.20.216.15[0/22]
              : 10.20.216.16[0/22]
State         : po-created
Dev State     : provisioned
App State     : cfg-in-sync
=====
--- Time Elapsed: 223.892847ms ---
```

efa tenant po update

Updates port channel parameters such as name, operation, management IP, and port.

Syntax

```
efa tenant po update [ --name po-name | --tenant tenant-name | --operation { port-add | port-delete | lACP-timeout | description | min-link-count | mtu-add | mtu-delete } | --port port-list | --lACP-timeout { long | short } | --min-link-count num-links | --description description | --mtu int32 ]
```

Parameters

--name *po-name*

Specifies the name of the port channel.

--tenant *tenant-name*

Specifies tenant name.

--operation { **port-add** | **port-delete** | **lACP-timeout** | **description** | **min-link-count** | **mtu-add** | **mtu-delete** }

Adds or deletes operation on the ports.

--port *port-list*

Specifies device IP along with ethernet port details. Example: SW1_IP [0/1], SW2_IP[0/5]

--lACP-timeout { **long** | **short** }

Specifies LACP timeout configuration.

--min-link-count *num-links*

Specifies the least number of links that need to be operationally up to declare the port channel as operationally up. The default value is 1. Valid values are 1 through 64.

--description

Describes the port channel.

--mtu **int32**

MTU configuration for the port channel. Valid values are 1500 through 9216.

Examples

This example updates the min-link-count of an existing PO.

```
(efa:extreme)extreme@node-1:~$ efa tenant po update --name poll1 --tenant tenant21
--operation min-link-count --min-link-count 1
```

```
PortChannel: poll1 updated successfully.
```

```
--- Time Elapsed: 378.143836ms ---
```

This example updates the lACP-timeout of an existing PO.

```
(efa:extreme)extreme@node-1:~$ efa tenant po update --name poll1 --tenant tenant21
--operation lACP-timeout --lACP-timeout long
```

```
PortChannel: po11 updated successfully.  
--- Time Elapsed: 5.883514246s ---
```

efa tenant service bgp peer configure

Pushes or removes a pending configuration for a BGP peer instance.

Syntax

```
efa tenant service bgp peer configure [ --name peer-name | --tenant tenant-name ]
```

Parameters

--name *peer-name*

Specifies the name of the BGP peer instance.

--tenant *tenant-name*

Specifies the name of the tenant.

Examples

This example pushes or removes a pending configuration for a BGP peer instance.

```
$ efa tenant service bgp peer configure --name B2 --tenant tenant11  
BgpService configured successfully.
```

efa tenant service bgp peer create

Creates a BGP (Border Gateway Protocol) peer for a specified VRF.

Syntax

```
efa tenant service bgp peer create [ --name peer-name | --tenant tenant-name | --description description | --ipv4-uc-dyn-nbr string | --ipv4-uc-nbr string | --ipv4-uc-nbr-bfd string | --ipv4-uc-nbr-route-map stringArray | --ipv4-uc-nbr-next-hop-self string | --ipv4-uc-nbr-remove-private-as stringArray | --pg-remove-private-as string | --ipv4-uc-nbr-default-originate stringArray | --ipv4-uc-nbr-default-originate-route-map stringArray | --ipv4-uc-nbr-update-source-ip string | --ipv4-uc-nbr-md5-password string | --md5-password-prompt-enable { true | false } | --ipv6-uc-dyn-nbr string | --ipv6-uc-nbr string | --ipv6-uc-nbr-bfd string | --ipv6-uc-nbr-default-originate stringArray | --ipv6-uc-nbr-default-originate-route-map stringArray | --ipv6-uc-nbr-md5-password string | --ipv6-uc-nbr-next-hop-self string | --ipv6-uc-nbr-remove-private-as stringArray | --ipv6-uc-nbr-route-map stringArray | --ipv6-uc-nbr-update-source-ip string ]
```

Parameters

--description *description*

Describes the BGP service.

--ipv4-uc-dyn-nbr *string*

Identifies the IPv4 unicast dynamic neighbor with a string in the following format: *device-ip,vrf-name:ipv4-listen-range,peer-group-name,listen-limit*. For example: *10.x.x.x,red:11::22/127*. The *listen-limit* value is optional.

--ipv4-uc-nbr *string*

Identifies the IPv4 unicast neighbor with a string in the following format: *device-ip,vrf-name:ipv4-neighbor,remote-as*. For example: *10.x.x.x,red:10.20.30.40,5000*.

--ipv4-uc-nbr-bfd *string*

Identifies the IPv4 BFD unicast neighbor with a string in the following format: *device-ip,vrf-name:ipv4-neighbor,bfd-enable (true/false),bfd-interval,bfd-min-rx,bfd-multiplier*. For example:

10.x.x.x,red:10.20.30.40,true,100,200,5.

--ipv4-uc-nbr-route-map *stringArray*

Identifies the IPv4 unicast neighbor route map in the following format: *device-ip,vrf-name:ipv4-neighbor,route-map-name,direction (in/out)*.

--ipv4-uc-nbr-remove-private-as *stringArray*

Turns the *remove-private-as* setting on or off per BGP peer group, created for external connectivity, in the following format: *device-ip,vrf-name:neighbor-ip,true|false*.

For example: *10.20.246.15,ten1vrf1:10.20.30.40,true*.

By default, the setting is off.

--ipv4-uc-nbr-default-originate *stringArray*

Identifies the `default-originate` setting for the IPv4 unicast neighbor, with a string in the following format: `device-ip,vrf-name:neighbor-ip,true/false`. For example:
10.20.246.15,ten1vrf1:10.20.30.40,true.

By default, the setting is off.

--ipv4-uc-nbr-default-originate-route-map *stringArray*

Identifies the `default-originate` setting for the IPv4 unicast neighbor, with a string in the following format: `device-ip,vrf-name:neighbor-ip,route-map`. For example:
10.20.246.16,ten1vrf1:10.20.30.40,rmap1.

By default, the setting is off.

--ipv4-uc-nbr-next-hop-self *string*

Identifies the `next-hop-self` for the IPv4 unicast neighbor with a string in the following format: `device-ip,vrf-name:ipv4-neighbor,next-hop-self (true/false/always)`. For example: 10.x.x.x,red:10.y.y.y,true.

--ipv4-uc-nbr-update-source-ip *string*

Updates the source IP for the IPv4 unicast neighbor with a string in the following format: `device-ip,vrf-name:ipv4-neighbor,update-source-ip`. For example:
10.x.x.x,red:10.y.y.y,11.x.x.x.

--ipv4-uc-nbr-md5-password *string*

Identifies the IPv4 unicast neighbor MD5 password in the following format: `device-ip,vrf-name:ipv4-neighbor,ipv4-md5-password`.

Example for encrypted/password containing special characters: `--ipv4-uc-nbr-md5-password device-ip,vrf-name:ipv4-neighbor,'9MCGkGaNt6OASX68/7TC6Lw=='`

--name *peer-name*

Identifies the name of the BGP peer instance.

--tenant *tenant-name*

Identifies the name of the tenant.

--md5-password-prompt-enable { **true** | **false** }

Turns on secure input for the MD5 password. The default is `false`.

--ipv6-uc-dyn-nbr *string*

Identifies the IPv6 unicast dynamic neighbor with a string in the following format: `device-ip,vrf-name:ipv6-listen-range,peer-group-name,listen-limit`. The `listen-limit` value is optional.

--ipv6-uc-nbr *string*

Identifies the IPv6 unicast neighbor with a string in the following format: `device-ip,vrf-name:ipv6-neighbor,remote-as`. For example: 10.x.x.x,red:10::40,5000.

--ipv6-uc-nbr-bfd *string*

Identifies the IPv6 BFD unicast neighbor with a string in the following format: `device-ip,vrf-name:ipv6-neighbor,bfd-enable (true/false),bfd-interval,bfd-min-rx,bfd-multiplier`. For example: `10.x.x.x,red:10::40,true,100,200,5`.

--ipv6-uc-nbr-default-originate *stringArray*

Identifies the IPv6 unicast neighbor default originate in the following format: `device-ip,vrf-name:ipv6-neighbor,default-originate (true/false)`.

--ipv6-uc-nbr-default-originate-route-map *stringArray*

Identifies the IPv6 unicast neighbor default originate route map in the following format: `device-ip,vrf-name:ipv6-neighbor,route-map-name`.

--ipv6-uc-nbr-md5-password *string*

Identifies the IPv6 unicast neighbor MD5 password, in the following format: `device-ip,vrf-name:ipv6-neighbor,ipv6-md5-password`.

Example for encrypted/password containing special characters: `--ipv6-uc-nbr-md5-password device-ip,vrf-name:ipv6-neighbor,'9MCgKGaNT6OASX68/7TC6Lw=='`

--ipv6-uc-nbr-next-hop-self *string*

Identifies the next-hop-self for the IPv6 unicast neighbor with a string in the following format: `device-ip,vrf-name:ipv6-neighbor,next-hop-self (true/false/always)`. For example: `10.x.x.x,red:10::40,true`.

--ipv6-uc-nbr-remove-private-as *stringArray*

Identifies the IPv6 unicast neighbor remove private AS in the format `device-ip,vrf-name:ipv6-neighbor,remove-private-as (true/false)`.

--ipv6-uc-nbr-route-map *stringArray*

Identifies the IPv6 unicast neighbor route map in the format `device-ip,vrf-name:ipv6-neighbor,route-map-name,direction (in/out)`.

--ipv6-uc-nbr-update-source-ip *string*

Updates the source IP address for the IPv6 unicast neighbor with a string in the following format: `device-ip,vrf-name:ipv6-neighbor,update-source-ip`. For example: `10.x.x.x,red:10::40,11::22`.

Examples

This example creates a static BGP peer (IPv4 and IPv6).

```
$ efa tenant service bgp peer create --name B1 --tenant tenant11
--ipv4-uc-nbr 10.20.216.16,blue11:1.1.1.11,95001
--ipv4-uc-nbr-bfd 10.20.216.16,blue11:1.1.1.11,true,50,5000,50
--ipv4-uc-nbr-next-hop-self 10.20.216.16,blue11:1.1.1.11,always
--ipv4-uc-nbr-update-source-ip 10.20.216.16,blue11:1.1.1.11,10.11.12.13
--ipv6-uc-nbr 10.20.216.16,blue11:20a1:a::10,95001
--ipv6-uc-nbr-bfd 10.20.216.16,blue11:20a1:a::10,true,50,5000,50
--ipv6-uc-nbr-next-hop-self 10.20.216.16,blue11:20a1:a::10,always
--ipv6-uc-nbr-update-source-ip 10.20.216.16,blue11:20a1:a::10,20::10

BgpService created successfully.
```


This example creates a dynamic BGP peer (IPv4 and IPv6).

```
$ efa tenant service bgp peer create --name B2 --tenant tenant11
--ipv6-uc-dyn-nbr 10.20.216.16,blue11:15::/127,pg1,10
--ipv4-uc-dyn-nbr 10.20.216.16,blue11:15.15.15.0/28,pg1,10
```

BgpService created successfully.

This example creates a BGP peer MD5 password in a secure manner. For more information, see the "Configure BGP MD5 Authentication for Tenant BGP Peer and Peer-group Securely" topic in the [Extreme Fabric Automation Security Guide, 2.7.0](#).

```
$ efa tenant service bgp peer create --name bgp173-2501
--tenant tenant11
--ipv4-uc-nbr 10.20.246.6,v1:25.1.1.3,5901
--ipv4-uc-nbr-bfd 10.20.246.6,v1:25.1.1.3,true
--ipv6-uc-nbr 10.20.246.5,v1:25:1::3,5901
--ipv6-uc-nbr-bfd 10.20.246.5,v1:25:1::3,true
--md5-password-prompt-enable=true
```

```
$ efa tenant service bgp peer create --name ten1bgppeer1
--tenant ten1 --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-remove-private-as 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-remove-private-as 10.20.246.16,ten1vrf1:10.20.30.40,true
```

```
$ efa tenant service bgp peer create --name ten1bgppeer1
--tenant ten1 --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-default-originate 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-default-originate 10.20.246.16,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-default-originate-route-map 10.20.246.16,ten1vrf1:10.20.30.40,rmap1
```

```
$ efa tenant service bgp peer create --name ten1bgppeer1
--tenant ten1 --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-prefix-list-in 10.20.246.15,ten1vrf1:10.20.30.40,ipprefix1
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-prefix-list-out 10.20.246.16,ten1vrf1:10.20.30.40,ipprefix1
```

```
$ efa tenant service bgp peer create --name ten1bgppeer1
--tenant ten1 --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-route-map 10.20.246.15,ten1vrf1:10.20.30.40,rmap1,in
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-route-map 10.20.246.16,ten1vrf1:10.20.30.40,rmap1,out
--ipv6-uc-nbr 10.20.246.15,ten1vrf1:25:1::3,50000
--ipv6-uc-nbr-bfd 10.20.246.15,ten1vrf1:25:1::3,true
--ipv6-uc-nbr-route-map 10.20.246.15,ten1vrf1:25:1::3,rmap1,in
--ipv6-uc-nbr 10.20.246.16,ten1vrf1:25:1::3,50000
--ipv6-uc-nbr-bfd 10.20.246.16,ten1vrf1:25:1::3,true
--ipv6-uc-nbr-route-map 10.20.246.16,ten1vrf1:25:1::3,rmap1,out
```

```
$ efa tenant service bgp peer create --name ten1bgppeer1
--tenant ten1 --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr-default-originate 10.20.246.15,ten1vrf1:10.20.30.40,true
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.40,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.40,true
```

```
--ipv4-uc-nbr-default-originate 10.20.246.16,ten1vrf1:10.20.30.40,true  
--ipv4-uc-nbr-default-originate-route-map 10.20.246.16,ten1vrf1:10.20.30.40,rmap1
```

efa tenant service bgp peer delete

Deletes the BGP neighbors for a given VRF on the fabric device.

Syntax

```
efa tenant service bgp peer delete [ --force | --name service-name | --tenant tenant-name ]
```

Parameters

--force

Forces BGP service deletion when set to true.

--name

Specifies the name of the BGP service instance.

--tenant

Specifies the name of the tenant.

Examples

The following example deletes the BGP peer instance.

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer delete --name B1 --tenant
tenant11

BgpService deleted successfully.

--- Time Elapsed: 4.301665698s ---
```

The following example deletes the BGP peer instance forcefully avoiding any error.

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer delete --name B1 --tenant
tenant11 --force

Bgp Service delete with "force" option will delete the device configuration corresponding
to the bgp and also
deletes the Bgp record from the application. Do you want to proceed (Y/N): y

BgpService deleted successfully.

--- Time Elapsed: 4.301665698s ---
```

efa tenant service bgp peer operational show

Gets the operational state of the BGP peers belonging to both default-vrf and tenant-vrf.

Syntax

```
efa tenant service bgp peer operational show [--tenant name | --vrf
name ]
```

Parameters

--tenant *name*

Specifies the name of the tenant.

--vrf *name*

Specifies the name of the VRF.

Examples

```
efa tenant service bgp peer operational show --tenant tenant2 --vrf vrf101
+-----+-----+-----+-----+-----+-----+-----+-----+
| Tenant | VRF   | Source | Destination | Source Device | BGP Peer |
| BGP Peer | BGP Peer | Device IP | Device IP | Router ID | BGP Peer |
| | | | | | |
| Source ASN | Destination ASN | AFI | SAFI | Session State | IP |
+-----+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.24 | | 172.31.254.178 | 10.40.40.251 |
| 65000 | 65000 | ipv4 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.24 | | 172.31.254.178 | 121.10.1.1 |
| 65000 | 4200000000 | ipv4 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.24 | | 172.31.254.178 | 121:a::1 |
| 65000 | 4200000000 | ipv6 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.24 | | 172.31.254.178 | fd40:4040:4040:1::fd |
| 65000 | 65000 | ipv6 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.23 | | 172.31.254.114 | 10.40.40.250 |
| 65000 | 65000 | ipv4 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.23 | | 172.31.254.114 | 121.10.1.2 |
| 65000 | 4200000000 | ipv4 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
| tenant2 | vrf101 | 10.20.246.23 | | 172.31.254.114 | fd40:4040:4040:1::fc |
| 65000 | 65000 | ipv6 | unicast | CONN | |
+-----+-----+-----+-----+-----+-----+
```

```

+-----+-----+-----+-----+-----+
BGP Summary
+-----+-----+-----+-----+-----+
efa tenant service bgp peer operational show
+-----+-----+-----+-----+-----+
+-----+
| Tenant | VRF | Source | Destination | Source Device | BGP
Peer | BGP Peer | BGP Peer | BGP Peer | BGP Peer | BGP Peer |
| | | Device IP | Device IP | Router ID |
IP | Source ASN | Destination ASN | AFI | SAFI | Session State |
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.24 | | 172.31.254.178 |
10.40.40.251 | 65000 | 34566 | ipv4 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.24 | | 172.31.254.178 |
121.10.1.1 | 65000 | 4200000000 | ipv4 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.24 | | 172.31.254.178 |
121:a::1 | 65000 | 4200000000 | ipv6 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.24 | | 172.31.254.178 |
fd40:4040:4040:1::fd | 65000 | 34566 | ipv6 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.23 | | 172.31.254.114 |
10.40.40.250 | 65000 | 34566 | ipv4 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.23 | | 172.31.254.114 |
121.10.1.2 | 65000 | 4200000000 | ipv4 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf102 | 10.20.246.23 | | 172.31.254.114 |
fd40:4040:4040:1::fc | 65000 | 34566 | ipv6 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101 | 10.20.246.23 | | 172.31.254.114 |
10.40.40.250 | 65000 | 65000 | ipv4 | unicast | CONN
|
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101 | 10.20.246.23 | | 172.31.254.114 |
121.10.1.2 | 65000 | 4200000000 | ipv4 | unicast | CONN

```

```

|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101      | 10.20.246.23 |          | 172.31.254.114 |
fd40:4040:4040:1::fc | 65000      | 65000          | ipv6      | unicast  | CONN
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101      | 10.20.246.24 |          | 172.31.254.178 |
10.40.40.251      | 65000      | 65000          | ipv4      | unicast  | CONN
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101      | 10.20.246.24 |          | 172.31.254.178 |
121.10.1.1        | 65000      | 4200000000    | ipv4      | unicast  | CONN
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101      | 10.20.246.24 |          | 172.31.254.178 |
121:a::1          | 65000      | 4200000000    | ipv6      | unicast  | CONN
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| tenant2 | vrf101      | 10.20.246.24 |          | 172.31.254.178 |
fd40:4040:4040:1::fd | 65000      | 65000          | ipv6      | unicast  | CONN
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.21 |          | 172.31.254.237 |
10.10.10.13      | 66000      | 64512          | ipv4      | unicast  | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.21 |          | 172.31.254.237 |
10.10.10.15      | 66000      | 64512          | ipv4      | unicast  | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.21 |          | 172.31.254.237 |
10.10.10.13      | 66000      | 64512          | l2vpn     | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.21 |          | 172.31.254.237 |
10.10.10.15      | 66000      | 64512          | l2vpn     | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.21 |          | 172.31.254.237 |
10.20.20.9       | 66000      | 66000          | l2vpn     | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+

```

```

|          | default-vrf | 10.20.246.22 |          | 172.31.254.218 |
10.10.10.5      | 66000      | 64512        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.22 |          | 172.31.254.218 |
10.10.10.7      | 66000      | 64512        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.22 |          | 172.31.254.218 |
10.10.10.5      | 66000      | 64512        | l2vpn   | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.22 |          | 172.31.254.218 |
10.10.10.7      | 66000      | 64512        | l2vpn   | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.22 |          | 172.31.254.218 |
10.20.20.8      | 66000      | 66000        | l2vpn   | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.0      | 64512      | 65000        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.2      | 64512      | 65000        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.4      | 64512      | 66000        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.6      | 64512      | 66000        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.8      | 64512      | 65000        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.10     | 64512      | 65000        | ipv4    | unicast | ESTAB
|
+-----+-----+-----+-----+-----+

```

```

+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.12 | 64512      | 66000        | | ipv4      | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.14 | 64512      | 66000        | | ipv4      | unicast | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.0  | 64512      | 65000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.2  | 64512      | 65000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.4  | 64512      | 66000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.6  | 64512      | 66000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.8  | 64512      | 65000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.10 | 64512      | 65000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.12 | 64512      | 66000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.14 |          | 172.31.254.239 |
10.10.10.14 | 64512      | 66000        | | l2vpn     | evpn    | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.23 |          | 172.31.254.114 |
10.10.10.9  | 65000      | 64512        | | ipv4      | unicast | ESTAB

```



```

|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.23 |          | 172.31.254.114 |
10.10.10.11      | 65000      | 64512          | ipv4     | unicast  | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.23 |          | 172.31.254.114 |
10.10.10.9       | 65000      | 64512          | l2vpn    | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.23 |          | 172.31.254.114 |
10.10.10.11      | 65000      | 64512          | l2vpn    | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.23 |          | 172.31.254.114 |
10.20.20.6       | 65000      | 65000          | l2vpn    | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.24 |          | 172.31.254.178 |
10.10.10.1       | 65000      | 64512          | ipv4     | unicast  | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.24 |          | 172.31.254.178 |
10.10.10.3       | 65000      | 64512          | ipv4     | unicast  | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.24 |          | 172.31.254.178 |
10.10.10.1       | 65000      | 64512          | l2vpn    | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.24 |          | 172.31.254.178 |
10.10.10.3       | 65000      | 64512          | l2vpn    | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
|          | default-vrf | 10.20.246.24 |          | 172.31.254.178 |
10.20.20.7       | 65000      | 65000          | l2vpn    | evpn     | ESTAB
|
+-----+-----+-----+-----+-----+
+-----+
BGP Summary
--- Time Elapsed: 21.822450796s ---

```

efa tenant service bgp peer show

Shows brief or detailed output of the BGP neighbors for a given tenant.

Syntax

```
efa tenant service bgp peer show [ --name peer-name | --tenant tenant-name
  | --detail ]
```

Parameters

--name *peer-name*

Specifies the name of the BGP peer instance.

--tenant *tenant-name*

Specifies the name of the tenant.

--detail

Displays detailed output of the BGP peer instance. When this parameter is not used, only brief tabular output is returned.

Examples

The following example shows brief output of all BGP peer instances.

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer show
=====
Name       : B2
Tenant     : tenant11
State      : bs-state-created

+-----+-----+-----+-----+-----+-----+-----+-----+
| Device IP | VRF | AFI | SAFI | Remote IP | Remote ASN | Next Hop | Update |
|           |     |     |     |           |           | Self     | Source IP |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| BFD       |     | BFD       | Dev State | App State |
| Enabled  | [Interval,Rx,Multiplier] |           |           |
+-----+-----+-----+-----+-----+-----+-----+-----+
Static Peer Details

+-----+-----+-----+-----+-----+-----+-----+-----+
| Device-IP | VRF | AFI | SAFI | Listen Range | Listen | Peer Group |
|           |     |     |     |           | Limit |           |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv4 | unicast | 15.15.15.0/28 | 10 | pg1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv6 | unicast | 15::/127 | 10 | pg1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| Dev State | App State |
|           |           |
+-----+-----+-----+-----+-----+-----+-----+-----+
| provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+-----+
| provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

```

Dynamic Peer Details
=====
Name      : B1
Tenant    : tenant11
State     : bs-state-created

+-----+-----+-----+-----+-----+-----+-----+
| Device IP | VRF | AFI | SAFI | Remote IP | Remote ASN | Next Hop |
|           |     |     |     |           |           | Self     |
+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv4 | unicast | 1.1.1.11 | 95001 | always |
+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv6 | unicast | 20a1:a::10 | 95001 | always |
+-----+-----+-----+-----+-----+-----+-----+
| Update | BFD | BFD | Dev State | App State |
| Source IP | Enabled | [Interval,Rx,Multiplier] | | |
+-----+-----+-----+-----+-----+-----+-----+
| 10.11.12.13 | true | 50, 5000, 50 | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+
| 20::10 | true | 50, 5000, 50 | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+

Static Peer Details

+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Device-IP | VRF | AFI | SAFI | Listen Range | Listen | Peer Group | Dev State | App
State |
|           |     |     |     |               | Limit |           |           |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
Dynamic Peer Details

=====

--- Time Elapsed: 482.525744ms ---
    
```

The following example shows detailed output of a specific BGP peer instance.

```

efa tenant service bgp peer show --name B2 --tenant tenant11 --detail
=====
Name      : B2
Tenant    : tenant11
State     : bs-state-created
Description :

Static Peer
-----
          0 Records

Dynamic Peer
-----
Device IP      : 10.20.216.16
VRF            : blue11
AFI            : ipv4
SAFI           : unicast
Listen Range   : 15.15.15.0/28
Listen Limit   : 10
Peer Group     : pgl
Dev State      : provisioned
    
```

```

App State      : cfg-in-sync

Device IP     : 10.20.216.16
VRF          : blue11
AFI          : ipv6
SAFI         : unicast
Listen Range  : 15::/127
Listen Limit  : 10
Peer Group    : pg1
Dev State     : provisioned
App State     : cfg-in-sync

=====

--- Time Elapsed: 129.553981ms ---

```

The following example shows brief output of all BGP peer instances belonging to a tenant.

```

$ efa tenant service bgp peer show --tenant tenant11
=====
Name      : B1
Tenant    : tenant11
State     : bs-state-created

+-----+-----+-----+-----+-----+-----+-----+
| Device IP | VRF | AFI | SAFI | Remote IP | Remote ASN | Next Hop |
|           |     |     |     |           |           | Self     |
+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv4 | unicast | 1.1.1.11 | 95001 | always |
+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv6 | unicast | 20a1:a::10 | 95001 | always |
+-----+-----+-----+-----+-----+-----+-----+
| Update | BFD | BFD | Dev State | App State |
| Source IP | Enabled | [Interval,Rx,Multiplier] | | |
+-----+-----+-----+-----+-----+-----+-----+
| 10.11.12.13 | true | 50, 5000, 50 | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+
| 20::10 | true | 50, 5000, 50 | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+
Static Peer Details

+-----+-----+-----+-----+-----+-----+-----+
+-----+
| Device-IP | VRF | AFI | SAFI | Listen Range | Listen | Peer Group | Dev State | App
State |
|           |     |     |     |           |       |           |           |
|           |     |     |     |           |       |           |           |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
Dynamic Peer Details

=====
=====
Name      : B2
Tenant    : tenant11
State     : bs-state-created

+-----+-----+-----+-----+-----+-----+-----+
| Device IP | VRF | AFI | SAFI | Remote IP | Remote ASN | Next Hop |
|           |     |     |     |           |           | Self     |
+-----+-----+-----+-----+-----+-----+-----+
| Update | BFD | BFD | Dev State | App State |

```

```

| Source IP | Enabled | [Interval,Rx,Multiplier] |           |           |
+-----+-----+-----+-----+-----+-----+
Static Peer Details
+-----+-----+-----+-----+-----+-----+
| Device-IP | VRF    | AFI  | SAFI  | Listen Range | Listen |
|           |        |      |      |              | Limit  |
+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv6 | unicast | 15::/127 | 10 |
+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv4 | unicast | 15.15.15.0/28 | 10 |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+
| Peer Group | Dev State | App State |
|           |          |          |
+-----+-----+-----+-----+
| pg1       | provisioned | cfg-in-sync |
+-----+-----+-----+-----+
| pg1       | provisioned | cfg-in-sync |
+-----+-----+-----+-----+
Dynamic Peer Details
=====

--- Time Elapsed: 138.308272ms ---

```

efa tenant service bgp peer update

Adds, deletes, or updates a BGP peer instance.

Syntax

```
efa tenant service bgp peer update [ --description description | --ipv4-uc-dyn-nbr string | --ipv4-uc-nbr string | --ipv4-uc-nbr-bfd string | --ipv4-uc-nbr-remove-private-as stringArray | --ipv4-uc-nbr-default-originate stringArray | --ipv4-uc-nbr-default-originate-route-map stringArray | --ipv4-uc-nbr-next-hop-self string | --ipv4-uc-nbr-update-source-ip string | --ipv4-uc-nbr-route-map stringArray | --ipv4-uc-nbr-md5-password password | --name peer-name | --operation { peer-add | peer-delete | desc-update } | --tenant tenant-name --md5-password-prompt-enable { true | false } ]
```

Parameters

--description *description*

Describes the BGP service.

--ipv4-uc-dyn-nbr *string*

Identifies the IPv4 unicast dynamic neighbor with a string in the following format: *device-ip,vrf-name:ipv4-listen-range,peer-group-name,listen-limit*. For example: *10.x.x.x,red:11::22/127*. The *listen-limit* value is optional.

--ipv4-uc-nbr *string*

Identifies the IPv4 unicast neighbor. For *peer-add* operations, enter a string in the following format: *device-ip,vrf-name:ipv4-neighbor,remote-as*. For example: *10.x.x.x,red:10.20.30.40,5000*. The *remote-as* value is not required for *peer-delete* operations.

--ipv4-uc-nbr-bfd *string*

Identifies the IPv4 BFD unicast neighbor with a string in the following format: *device-ip,vrf-name:ipv4-neighbor,bfd-enable(true/false),bfd-interval,bfd-min-rx,bfd-multiplier*. For example: *10.x.x.x,red:10.20.30.40,true,100,200,5*. BFD parameters are optional for *peer-add* operations and not required for *peer-delete* operations.

--ipv4-uc-nbr-remove-private-as *stringArray*

Turns the *remove-private-as* setting on or off per BGP peer group, created for external connectivity, in the following format: *<device-ip,vrf-name:neighbor-ip,true|false*. For example: *10.20.246.15,ten1vrf1:10.20.30.40,true*.

By default, the setting is off.

--ipv4-uc-nbr-default-originate *stringArray*

Identifies the *default-originate* setting for the IPv4 unicast neighbor, with a string in the following format: *device-ip,vrf-name:neighbor-ip,true/false*. For example: *10.20.246.15,ten1vrf1:10.20.30.40,true*.

--ipv4-uc-nbr-default-originate-route-map *stringArray*

Identifies the default-originate setting for the IPv4 unicast neighbor, with a string in the following format: *device-ip,vrf-name:neighbor-ip,route-map*. For example:
10.20.246.16,ten1vrf1:10.20.30.40,rmap1.

This parameter is optional.

--ipv4-uc-nbr-next-hop-self *string*

Identifies the next-hop-self for the IPv4 unicast neighbor with a string in the following format: *device-ip,vrf-name:ipv4-neighbor,next-hop-self (true/false/always)*. For example: 10.x.x.x,red:10.y.y.y,true.

--ipv4-uc-nbr-update-source-ip *string*

Updates the source IP for the IPv4 unicast neighbor with a string in the following format: *device-ip,vrf-name:ipv4-neighbor,update-source-ip*. For example:
10.x.x.x,red:10.y.y.y,11.x.x.x.

--ipv4-uc-nbr-route-map *stringArray*

Identifies the IPv4 BFD route map, in the following format: *device-ip,vrf-name:neighbor-ip,route-map-name*. For example:
10.20.246.15,ten1vrf1:10.20.30.40,rmap1.

--ipv4-uc-nbr-md5-password *password*

IPv4 Unicast Neighbor md5-password in the format *device-ip,vrf-name:ipv4-neighbor,ipv4-md5-password*.

Example for encrypted/password containing special characters: `--ipv4-uc-nbr-md5-password device-ip,vrf-name:ipv4-neighbor,'9MCgKGaNT6OASX68/7TC6Lw=='`

--name *peer-name*

Identifies the name of the BGP peer instance.

--operation { **peer-add** | **peer-delete** | **desc-update** }

Identifies the type of operation you are performing: add a peer, delete a peer, or update a description.

--tenant *tenant-name*

Identifies the name of the tenant.

--md5-password-prompt-enable{ **true** | **false** }

Turns on secure input for the MD5 password. The default is `false`.

Examples

This example adds a peer to instance B2.

```
$ efa tenant service bgp peer update
--name B2 --tenant tenant11 --operation peer-add
--ipv6-uc-nbr 10.20.216.15,blue11:18::1,98100
--ipv4-uc-nbr 10.20.216.16,blue11:1.1.1.12,95001

BGP service updated successfully.
```

This example deletes a BGP peer for the specified VRF (red) in the bgpservice1 instance for tenant1.

```
$ efa tenant service bgp peer update --name bgpservice1 --tenant tenant1
--operation peer-delete --ipv4-uc-nbr 10.24.80.134,red:10.20.30.40
```

```
$ efa tenant service bgp peer update --name ten1bgppeer1 --tenant ten1
--operation peer-add --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-remove-private-as 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-remove-private-as 10.20.246.16,ten1vrf1:10.20.30.50,false
```

```
$ efa tenant service bgp peer update --name ten1bgppeer1 --tenant ten1
--operation peer-add --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-default-originate 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-default-originate-route-map 10.20.246.15,ten1vrf1:10.20.30.50,rmap2
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-default-originate 10.20.246.16,ten1vrf1:10.20.30.50,false
```

```
$ efa tenant service bgp peer update --name ten1bgppeer1 --tenant ten1
--operation peer-add --ipv4-uc-nbr 10.20.246.15,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-route-map 10.20.246.15,ten1vrf1:10.20.30.50,rmap2,in
--ipv4-uc-nbr 10.20.246.16,ten1vrf1:10.20.30.50,50000
--ipv4-uc-nbr-bfd 10.20.246.16,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-route-map 10.20.246.16,ten1vrf1:10.20.30.50,rmap2,out
```


efa tenant service bgp peer-group configure

Pushes or removes a partial configuration for a peer-group.

Syntax

```
efa tenant service bgp peer-group configure [ --name peer-group-name | --tenant tenant-name ]
```

Parameters

--name *peer-group-name*

Specifies the name of the BGP peer group instance.

--tenant *tenant-name*

Specifies the name of the tenant.

Examples

Pushes or removes a partial configuration for a peer-group.

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer-group configure --name PG1 --tenant tenant11
BgpService configured successfully.

--- Time Elapsed: 790.953838ms ---
```

efa tenant service bgp peer-group create

Creates a BGP peer group for the specified tenant.

Syntax

```
efa tenant service bgp peer-group create [ --name peer-group-name | --tenant tenant-name | --pg-asn string | --pg-bfd string | --pg-bfd-enable string | --pg-remove-private-as stringArray | --pg-name string | --pg-next-hop-self string | --pg-update-source-ip string | --pg-md5-password string | --pg-md5-password-prompt-enable { true | false } | --description string]
```

Parameters

--name *peer-group-name*

Specifies the name of the BGP peer group instance.

--tenant *tenant-name*

Specifies the name of the tenant.

--pg-asn *string*

Specifies the ASN of the BGP peer group with a string in the following format: *device-ip:peer-group-name,remote-asn*.

--pg-bfd *string*

Specifies the BFD properties of the BGP peer group with a string in the following format: *device-ip: peer-group-name,bfd-enable(true/false),interval,bfd-min-rx,bfd-multiplier*.

--pg-bfd-enable *string*

Turns on BGP peer group with a string in the following format: *device-ip,peer-group-name:md5-password*.

--pg-name *string*

Specifies the BGP peer group name with a string in the following format: *device-ip:peer-group-name*.

--pg-remove-private-as *stringArray*

Turns the *remove-private-as* setting on or off per BGP peer group, created for external connectivity, in the following format: *device-ip,pg-name:true|false*. For example: *10.20.246.16,pg1:true*.

By default, the setting is off.

--pg-next-hop-self *string*

Specifies the next-hop-self for the BGP peer group with a string in the following format: *device-ip:peer-group-name,next-hop-self(true/false/always)*.

--pg-update-source-ip *string*

Updates the source IP address of the BGP peer group with a string in the following format: *device-ip:peer-group-name,update-source-ip*.

--pg-md5-password *string*

Indicates the MD5 password.



Important

BGP MD5 authentication for tenant dynamic peers is not supported.

--pg-md5-password-prompt-enable { **true** | **false** }

Turns on secure input for the MD5 password. The default is false.



Important

BGP MD5 authentication for tenant dynamic peers is not supported.

--description *string*

Describes the BGP peer group.

Examples

This example creates a peer group instance.

```
$ efa tenant service bgp peer-group create --name ten1bgppg1
--tenant ten1 --pg-name 10.20.246.15:pg1 --pg-asn 10.20.246.15,pg1:55001
--pg-bfd-enable 10.20.246.15,pg1:true
--pg-remove-private-as 10.20.246.15,pg1:true --pg-name 10.20.246.16:pg1
--pg-asn 10.20.246.16,pg1:55001 --pg-bfd-enable 10.20.246.16,pg1:true
--pg-remove-private-as 10.20.246.16,pg1:true
```

This example creates a BGP peer-group MD5 password in a secure manner. For more information, see the "Configure BGP MD5 Authentication for Tenant BGP Peer and Peer-group Securely" topic in the [Extreme Fabric Automation Security Guide, 2.7.0](#).

```
$ efa tenant service bgp peer-group create --tenant "tenant11"
--name "v1-PeerGrp" --pg-name 10.20.246.5:v1-PeerGrp --pg-asn 10.20.246.5,v1-PeerGrp:5200
--pg-bfd-enable 10.20.246.5,v1-PeerGrp:true --pg-name 10.20.246.5:v3-PeerGrp
--pg-asn 10.20.246.5,v3-PeerGrp:5201 --pg-bfd-enable 10.20.246.5,v3-PeerGrp:true
--pg-name 10.20.246.6:v1-PeerGrp --pg-asn 10.20.246.6,v1-PeerGrp:5200
--pg-bfd-enable 10.20.246.6,v1-PeerGrp:true --pg-md5-password-prompt-enable=true
```

efa tenant service bgp peer-group delete

Deletes the specified BGP peer group.

Syntax

```
efa tenant service bgp peer-group delete [ --force | --name peer-group-name | --tenant tenant-name ]
```

Parameters

--force

Forces the deletion of the BGP service when set to true.

--name *peer-group-name*

Specifies the name of the BGP peer instance.

--tenant *tenant-name*

Specifies the name of the tenant.

Examples

This example removes a peer-group instance.

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer-group delete --name PG1 --tenant tenant11  
BgpService deleted successfully.  
--- Time Elapsed: 1.359719067s ---
```

This example removes an instance forcefully and ignores any error.

```
(efa:extreme)extreme@node-1:~$ efa tenant service bgp peer-group delete --name PG1 --tenant tenant11 --force  
Bgp service peer-group delete with "force" option will delete the device configuration corresponding to the bgp and also deletes the bgp record from the application. Do you want to proceed (Y/N): y  
BgpService deleted successfully.  
--- Time Elapsed: 1.359719067s ---
```

efa tenant service bgp peer-group show

Displays brief or detailed output of BGP peer groups for the specified tenant.

Syntax

```
efa tenant service bgp peer-group show [ --name peer-group-name | --
  tenant tenant-name | --detail ]
```

Parameters

--name *peer-group-name*

Specifies the name of the BGP peer group instance.

--tenant *tenant-name*

Specifies the name of the tenant.

--detail

Displays detailed output of the BGP peer group instance. When this parameter is not used, only brief tabular output is returned.

Examples

The following example shows brief output of all instances.

```
$ efa tenant service bgp peer-group show
=====
Name       : PG1
Tenant     : tenant11
State      : bgp-pg-state-created

+-----+-----+-----+-----+-----+-----+
| Device IP | Peer Group | Remote | Next Hop | Update | BFD |
|           |           | ASN   | Self    | Source IP | Enabled |
+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | pg1 | 95002 | true | 10.10.10.3 | true |
+-----+-----+-----+-----+-----+-----+
|           | BFD | Dev State | App State |
| [Interval,Rx,Multiplier] | | |
+-----+-----+-----+-----+
|           | 660, 506, 20 | provisioned | cfg-in-sync |
+-----+-----+-----+-----+
BGP PeerGroup Details

=====

--- Time Elapsed: 541.906975ms ---
```

The following example shows detailed output of a specific instance.

```
$ efa tenant service bgp peer-group show --name PG1 --tenant tenant11
--detail
=====
Name       : PG1
Tenant     : tenant11
State      : bgp-pg-state-created
```

```

Description      :
Peer Group
-----
Device IP       : 10.20.216.16
Peer Group      : pg1
Remote ASN      : 95002
Next Hop Self   : true
Update Source IP : 10.10.10.3
BFD Enabled     : true
BFD Interval    : 660
BFD Rx          : 506
BFD Multiplier  : 20
Dev State       : provisioned
App State       : cfg-in-sync

=====

--- Time Elapsed: 111.316727ms ---
    
```

The following example shows brief output of all instances created for a tenant.

```

$ efa tenant service bgp peer-group show --tenant tenant11
=====
Name       : PG1
Tenant     : tenant11
State      : bgp-pg-state-created

+-----+-----+-----+-----+-----+
| Device IP | Peer Group | Remote | Next Hop | Update |
|           |            | ASN    | Self     | Source IP |
+-----+-----+-----+-----+-----+
| 10.20.216.16 | pg1      | 95002  | true     | 10.10.10.3 |
+-----+-----+-----+-----+-----+

+-----+-----+-----+-----+-----+
| BFD      | BFD          | Dev State | App State |
| Enabled  | [Interval,Rx,Multiplier] |           |           |
+-----+-----+-----+-----+-----+
| true    | 660, 506, 20 | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+
BGP PeerGroup Details
=====

--- Time Elapsed: 99.370646ms ---
    
```

efa tenant service bgp peer-group update

Adds, deletes, or updates a BGP peer group for the specified tenant.

Syntax

```
efa tenant service bgp peer-group update [ --name peer-group-name | --
  operation { peer-group-add | peer-group-delete | desc-update } | --pg-
  asn string | --pg-bfd string | --pg-bfd-enable string | --pg-remove-
  private-as stringArray | --pg-md5-password string | --pg-md5-
  password-prompt-enable { true | false } | --pg-name string | --pg-
  next-hop-self string | --pg-update-source-ip string | --tenant
  tenant-name | --description description ]
```

Parameters

--name *peer-group-name*

Specifies the name of the BGP peer group instance.

--operation { **peer-group-add** | **peer-group-delete** | **desc-update** }

Specifies the type of operation you are performing: add a peer group, delete a peer group, or update a description.

--pg-asn *string*

Specifies the ASN of the BGP peer group with a string in the following format: *device-ip;peer-group-name,remote-asn*.

--pg-remove-private-as *stringArray*

Turns the *remove-private-as* setting on or off per BGP peer group, created for external connectivity, in the following format: *device-ip,pg-name:true|false*. For example: *10.20.246.16,pg1:true*.

By default, the setting is off.

--pg-bfd *string*

Specifies the BFD properties of the BGP peer group with a string in the following format: *device-ip:peer-group-name,bfd-enable(true/false),interval,min-rx,multiplier*.

--pg-bfd-enable *string*

Specifies the BGP peer-group BFD in the following format: *device-ip,peer-group-name:bfd-enable(true/false)*.

--pg-md5-password *string*

Indicates the MD5 password.



Important

BGP MD5 authentication for tenant dynamic peers is not supported.

--pg-md5-password-prompt-enable { **true** | **false** }

Turns on secure input for the MD5 password. The default is `false`.

**Important**

BGP MD5 authentication for tenant dynamic peers is not supported.

--pg-name *string*

Specifies the name of the BGP peer group with a string in the following format: `device-ip:peer-group-name`.

--pg-next-hop-self *string*

Specifies the next-hop-self for the BGP peer group with a string in the following format: `device-ip:peer-group-name,next-hop-self (true/false/always)`.

--pg-update-source-ip *string*

Updates the source IP address for the BGP peer group with a string in the following format: `device-ip:peer-group-name,update-source-ip`.

--tenant *tenant-name*

Specifies the name of the tenant.

--description *description*

Describes the BGP service.

Examples

This example updates a peer group instance by defining several of its attributes.

```
$ efa tenant service bgp peer-group update --name ten1bgppg1
--tenant ten1 --operation peer-group-add --pg-name 10.20.246.15:pg2
--pg-asn 10.20.246.15,pg2:55002 --pg-bfd-enable 10.20.246.15,pg2:true
--pg-remove-private-as 10.20.246.15,pg2:true --pg-name 10.20.246.16:pg2
--pg-asn 10.20.246.16,pg2:55002 --pg-bfd-enable 10.20.246.16,pg2:true
```

This example updates a peer group instance and deletes a peer group from an instance.

```
$ efa tenant service bgp peer-group update
--name PG1 --tenant tenant11 --operation peer-group-delete
--pg-name 10.20.216.15:pg1
```


efa tenant service mirror session create

Creates a mirror session for a specified tenant.

Syntax

```
efa tenant service mirror session create [-name session-name | --tenant
  tenant-name | --description string | --source stringArray | --type
  stringArray | --destination-type stringArray | --destination
  stringArray | --direction stringArray | -h, --help ]
```

Parameters

-name *session-name*

Specifies the name of the mirror service session.

--tenant *tenant-name*

Specifies the name of the mirror source tenant.

--description *stringArray*

Specifies the description of the mirror service session.

--source *stringArray*

Specifies the mirror source, in the following format: *deviceIP, IfType, IfName*. Valid values of **IfType** are **eth** | **po** | **vlan**. Global Vlan Source format: *vlan, vlan-range*. Vlan Range Source format: *vlan, a-b, c, d, e-f, g*.

Example: *10.20.20.1, eth, 0/1 10.20.20.1, po, po1 10.20.20.1, vlan, 100
vlan, 100 vlan, 100-110, 115, 120-130.*

A maximum of 253 characters are allowed for the Vlan Range.

--type *stringArray*

Specifies the mirror type in the following format: *sourceDeviceIP, sourceIfType, sourceIfName:Value*

Example: *10.20.20.1, eth, 0/1:port-based 10.20.20.1, po, po1:flow-based*.
Default is port-based. Valid values are **port-based** | **flow-based**.

--destination-type *stringArray*

Specifies the mirror destination type in the following format: *sourceDeviceIP, sourceIfType, sourceIfName:value*

Example: *10.20.20.1, eth, 0/1:span*. Valid Values are **span**.

--destination *string*

Specifies the mirror destination in the following format: *sourceDeviceIP, sourceIfType, sourceIfName:destinationDeviceIP, destinationIfType, destinationIfName*

Example: *10.20.20.1, eth, 0/1:10.20.20.1, eth, 0/5
10.20.20.1, vlan, 100:10.20.20.1, eth, 0/6 vlan, 100:10.20.20.1, eth, 0/6*

--direction *stringArray*

Specifies the mirror traffic direction in the following format: sourceDeviceIP, sourceIfType, sourceIfName:Value

Example: 10.20.20.1,eth,0/1:tx. Valid values are **tx** | **rx** | **both**.

Examples

This example demonstrates the configuration to enable traffic mirroring with the following features:

- Out-of-band mirroring
- Port-based mirroring
- Multi tenancy

(1) For the IP: 10.20.246.15.

```
efa tenant service mirror session create -name ten1mirrorsession1
--tenant ten1--source-if-type port-channel--source-if-name 10.20.246.15,ten1po1
--destination-if-type eth--destination-if-name 10.20.246.15,0/31
--direction both--type port-based
```

and

```
efa tenant service mirror session create -name ten2mirrorsession1
--tenant ten2--source-if-type port-channel--source-if-name 10.20.246.15,ten2po1
--destination-if-type eth--destination-if-name 10.20.246.15,0/31
--direction both--type port-based
```

(2) For the IP: 10.20.246.16

```
efa tenant service mirror session create -name ten1mirrorsession2
--tenant ten1--source-if-type port-channel--source-if-name 10.20.246.16,ten1po1
--destination-if-type eth--destination-if-name 10.20.246.16,0/31
--direction both--type port-based
```

and

```
efa tenant service mirror session create -name ten2mirrorsession2
--tenant ten2--source-if-type port-channel--source-if-name 10.20.246.16,ten2po1
--destination-if-type eth--destination-if-name 10.20.246.16,0/31
--direction both--type port-based
```

efa tenant service mirror session delete

Deletes the specified mirror session.

Syntax

```
efa tenant service mirror session delete [ --tenant tenant-name | --name  
mirror-session-name | --force ]
```

Parameters

--tenant *tenant-name*

Specifies the name of the tenant being mirrored.

--name *mirror-session-name*

Specifies the name of the mirror session to delete.

--force

Forces the deletion of the mirror session service.

efa tenant service mirror session configure

Pushes or removes a pending mirror session configuration.

Syntax

```
efa tenant service mirror session configure [ --tenant tenant-name | --name mirror-session-name ]
```

Parameters

--tenant *tenant-name*

Specifies the name of the source tenant.

--name *mirror-session-name*

Specifies the name of the mirror session.

efa tenant service mirror session show

Displays the mirror session details for the specified tenant or mirror session.

Syntax

```
efa tenant service mirror session show [ --tenant tenant-name | --name mirror-session-name ]
```

Parameters

- tenant** *tenant-name*
Specifies the name of the source tenant.
- name** *mirror-session-name*
Specifies the name of the mirror session.

efa tenant show

Displays brief or detailed tenant details.

Syntax

```
efa tenant show [ --name tenant-name | --detail ]
```

Parameters

--name *tenant-name*

Specifies the name of the tenant.

--detail

Displays detailed output of the tenant. When this parameter is not used, only brief tabular output is returned.

Examples

Use this syntax to return brief output of all configured tenants.

```
$ efa tenant show
```

Use this syntax to return brief output of a specific tenant.

```
$ efa tenant show --name tenant21
```

Use this syntax to return detailed output of a specific tenant.

```
$ efa tenant show --name tenant21 --detail
```

efa tenant update

Allows changes to be made to a tenant after it has been created.

Syntax

```
efa tenant update [ --name tenant-name | --description tenant-
description | --l2-vni-range range | --l3-vni-range range | --vlan-
range range | --vrf-count number | --enable-bd | --operation { desc-
update | vni-update | port-add | port-delete | vlan-add | vlan-delete
| vlan-update | num-vrf-update | enable-bd-update | mirror-
destination-port-add | mirror-destination-port-delete } | --port port-
list | --force | --port value | --mirror-destination-port value ]
```

Parameters

--name *tenant-name*

Specifies name of tenant.

--description *tenant-description*

Describes the tenant.

--l2-vni-range *range*

Specifies the contiguous range of Layer 2 VNIs in ascending order that are reserved for the tenant within the scope of a fabric.

--l3-vni-range *range*

Specifies the contiguous Range of L3 VNIs in ascending order will be reserved for the tenant within the scope of a fabric.

--vlan-range *range*

Specifies the range of VLANs to be reserved for the tenant.

--vrf-count **int** *number*

Specifies the number of VRFs to be reserved for the tenant.

--enable-bd

Enables BD capability for networks created under this tenant.

--operation { **desc-update** | **vni-update** | **port-add** | **port-delete** | **vlan-add** | **vlan-delete** | **vlan-update** | **num-vrf-update** | **enable-bd-update** | **mirror-destination-port-add** | **mirror-destination-port-delete** }

Specifies operation code.

--port *value*

Lists physical ports of devices which will be reserved for the asset. Example:
SW1_IP[0/1],SW2_IP[0/5]

--mirror-destination-port *value*

xxx

--force

Forces the deletion on the tenant if the option is provided.

Examples

This example removes a VLAN range from the tenant.

```
$ efa tenant update --name tenant21 --operation vlan-delete  
--vlan-range 201-300
```

```
Tenant updated successfully.
```

This example adds a port to an existing tenant.

```
$ efa tenant update --name tenant21 --operation port-add  
--port 10.20.216.15[0/29]
```

```
Tenant updated successfully.
```


efa tenant vrf create

Creates a tenant VRF.

Syntax

```
efa tenant vrf create [ --name vrf-name | --tenant tenant-name | --rt-
  type { both | import | export } | --rt value | --ipv4-static-route-
  next-hop route | --ipv6-static-route-next-hop route | --local-asn
  local-asn | --ipv4-static-route-bfd route | --ipv6-static-route-bfd
  route | --ipv4-static-network stringArray | --ipv4-static-network-
  distance stringArray | --ipv6-static-network-distance stringArray |
  --max-path unit | --redistribute { static | connected } | --rh-max-
  path { 8 | 16 | 64 } | --rh-ecmp-enable { true | false } | --
  graceful-restart-enable { true | false } | --routing-type
  { distributed | centralized } | --centralized-router string | --ipv4-
  aggregate-address stringArray | --ipv6-aggregate-address stringArray
  | --ipv4-aggregate-as-set stringArray | --ipv6-aggregate-as-set
  stringArray | --ipv4-aggregate-summary-only stringArray | --ipv6-
  aggregate-summary-only stringArray | --ipv4-aggregate-advertise-map
  stringArray | --ipv6-aggregate-advertise-map stringArray | --ipv4-
  aggregate-suppress-map stringArray | --ipv6-aggregate-suppress-map
  stringArray | --ipv4-network network | --ipv4-network-backdoor
  stringArray | --ipv4-network-weight weight | --ipv4-network-route-map
  route-map | --ipv6-network network | --ipv6-network-backdoor
  stringArray | --ipv6-network-weight weight | --ipv6-network-route-map
  route-map ]
```

Parameters

--name *vrf-name*

Specifies the name of the VRF.

--tenant *tenant-name*

Specifies the name of the tenant.

--rt-type { **both** | **import** | **export** }

Specifies the route's target VPN community. Valid values are both, import, or export. The default is an empty string array.

--rt *value*

Specifies the unique number for forming the Route Target and the Route Distinguisher.

--ipv4-static-route-next-hop *route*

Specifies the IPv4 static route next hop. Valid values for the route distance are 1 through 254.

Format: device IP, IPv4 static route network, next hop IP, and route distance separated by commas.

Example: 10.25.25.100,20.0.0.0/24,16.0.0.2,3.

--local-asn *local-asn*

Specifies the local ASN for the VRF.

--ipv4-static-route-bfd *route*

Specifies the IPv4 static route BFD.

Format: device IP, destination IPv4 address, source IPv4 address[interval,min-rx,multiplier]

Example: 10.25.25.100,1.1.1.1,2.2.2.2,123,456,3

--ipv6-static-route-bfd *route*

Specifies the IPv6 static route BFD.

Format: device-ip,dest-ipv6-addr,source-ipv6-addr[interval,min-rx,multiplier].

Example: 10.25.25.100,1:::1,2:::2,300,300,3

--ipv6-static-route-next-hop *route*

Specifies the IPv4 static route next hop in the following format: device IP, IPv6 static route network, next hop IP, and route distance separated by commas. For example:

10.25.25.100,2001:1:::/64,3001:::2,3. Valid values for the route distance are 1 through 254.

--max-path *unit*

Specifies the number of load-sharing paths for the VRF. Valid values are 1 through 64.

--redistribute { **static** | **connected** }

Specifies the redistribute type for routes. Valid values are static or connected. The default is an empty string array.

--rh-max-path { **8** | **16** | **64** }

Specifies the maximum number of resilient hashing paths allowed per tenant VRF. Valid values are 8, 16, or 64. The default is 0 and unless you specify it, it is not set.

--rh-ecmp-enable { **true** | **false** }

Turns on or turns off resilient hashing for a tenant VRF. Valid values are true or false. The default value is false.

--graceful-restart-enable { **true** | **false** }

Turns on or turns off graceful restart for a tenant VRF. Valid values are true or false. The default value is false.

--routing-type { **distributed** | **centralized** }

VRF routing type. Valid values are distributed or centralized. The default value is distributed.

--centralized-router*string*

Comma-separated list of border-leaf IP addresses.

Example: 10.10.1.1,10.10.1.2 | 10.10.1.1-2

--ipv4-network *network*

Specified IPv4 network address for BGP to advertise.

Format:device-ip,network-address

Example: 10.24.80.134,10.20.30.0/30

--ipv4-network-backdoor *stringArray*

Increases the administrative distance of eBGP so that IGP learned routes are preferred for the specified network address. Disabled, by default.

Format: device-ip,network-address,flag

Example: 10.24.80.134,10.21.30.0/30,true

--ipv4-network-weight *weight*

Weight to use to set the BGP weight attribute for the given network address.

Format: device-ip,network-address,weight

Example: 10.24.80.134,10.22.30.0/30,144

--ipv4-network-route-map *route-map*

Route map to use to set the BGP attributes for the given network address, such as MED.

Format: device-ip,network-address,rm

Example: 10.24.80.134,10.23.30.0/30,rmap1

--ipv6-network *stringArray*

BGP advertises given network address.

Format: device-ip,network-address

Example: --ipv6-network 10.24.80.134,11::/128

--ipv6-network-backdoor *stringArray*

Increases the administrative distance of eBGP with the goal of making IGP learned routes preferred for the given network address.

Format: device-ip,network-address

Example: --ipv6-network 10.24.80.134,11::/128

--ipv6-network-weight *stringArray*

Given weight is used to set the BGP weight attribute for the given network address.

Format: device-ip,network-address,weight.

Example: --ipv6-network-weight 10.24.80.134,11::/128,144

--ipv6-network-route-map *stringArray*

Given route map is used to set the BGP attributes for the given network address like MED.

Format: device-ip,network-address,rm.

Example: --ipv6-network-route-map 10.24.80.134,11::/128,rmap1.

--ipv4-static-network *stringArray*

IPv4 Static Network to be configured on the device.

Format: Device IP, IPv4 Static Network

Example: 10.24.80.134,11.10.30.40/30

--ipv4-static-network-distance *stringArray*

IPv4 Static Network with distance to be configured on the device. Valid values are 1 through 255.

Format: Device IP, IPv4 Static Network, Distance

Example: 10.24.80.134,11.10.30.40/30,169

--ipv4-aggregate-advertise-map *stringArray*

Route map name used to filter the BGP attributes to be advertised as part of aggregate-address.

Format: device-ip,ipv4-aggregate-address,ipv4-aggregate-advertise-map

Example: 10.25.25.100,10.20.21.40/30,routeMap1

--ipv6-aggregate-advertise-map *stringArray*

Route map name used to filter the BGP attributes to be advertised as part of aggregate-address.

Format: device-ip,ipv6-aggregate-address,ipv6-aggregate-advertise-map.

Example: 10.25.25.100,10::20/126,routeMap1.

--ipv4-aggregate-suppress-map *stringArray*

Route map name used to suppress the specific routes to be advertised along with the aggregate-address advertisement.

Format: device-ip,ipv4-aggregate-address,ipv4-aggregate-suppress-map

Example: 10.25.25.100,10.20.21.40/30,routeMap1

--ipv4-aggregate-as-set *stringArray*

Enable to set autonomous system set path information as part of aggregate-address advertisement.

Format: device-ip,ipv4-aggregate-address,ipv4-aggregate-as-set

Example: 10.25.25.100,10.20.21.40/30,true

--ipv6-aggregate-as-set *stringArray*

Enable to set AS set path information as part of aggregate-address advertisement.

Format: device-ip,ipv6-aggregate-address,ipv6-aggregate-as-set.

Example: 10.25.25.100,10::20/126,true.

--ipv4-aggregate-summary-only *stringArray*

Enable to advertise only aggregated-address.

Format: device-ip,ipv4-aggregate-address,ipv4-aggregate-summary-only

Example: 10.25.25.100,10.20.21.40/30,true

--ipv6-aggregate-summary-only *stringArray*

Enable to advertise only aggregated-address.

Format: device-ip,ipv6-aggregate-address,ipv6-aggregate-summary-only.

Example: 10.25.25.100,10::20/126,true.

--ipv4-aggregate-address *stringArray*

IPv4 aggregate address.

Format: device-ip,ipv4-aggregate-address

Example: 10.25.25.100,10.20.21.40/30

--ipv6-aggregate-address *stringArray*

IPv6 aggregate address.

Format: device-ip,ipv6-aggregate-address

Example: 10.25.25.100,10::20/126

--ipv6-aggregate-suppress-map *stringArray*

Route map name used to suppress the specific routes to be advertised along with the aggregate-address advertisement.

Format: device-ip,ipv6-aggregate-address,ipv6-aggregate-suppress-map.

Example: 10.25.25.100,10::20/126,routeMap1

Usage Guidelines

The **max-path** and **rh-max-path** parameters can co-exist.

You cannot choose the specific devices on which to configure resilient hashing. Configuration applies to all SLX devices in the tenant VRF.

Examples

This example creates a distributed VRF.

```
$ efa tenant vrf create --tenant tenant11
--name blue11 --local-asn 65001 --rt-type import --rt 100:100 --rt-type export
--rt 100:100 --rt-type import --rt 200:200 --rt-type export --rt 200:200 --rt-type import
--rt 300:300 --rt-type export --rt 400:400 --max-path 50 --redistribute connected
--redistribute static --ipv4-static-route-next-hop
10.20.216.16,192.168.0.0/24,10.10.10.1,5
--ipv4-static-route-next-hop 10.20.216.15,192.168.10.0/24,10.10.10.5,5
--ipv6-static-route-next-hop 10.20.216.16,2020:20::1/128,3001::2,6
--ipv6-static-route-next-hop 10.20.216.15,2020:30::1/128,3001::3,5
--ipv6-static-route-bfd 10.20.216.16,3001::3,3001::1,100,200,5
--ipv6-static-route-bfd 10.20.216.16,3001::2,3001::1
--ipv6-static-route-bfd 10.20.216.16,3001::4,3001::1,100,300,6
--ipv4-static-route-bfd 10.20.216.15,10.10.10.1,10.10.10.254,200,300,6
--ipv4-static-route-bfd 10.20.216.16,10.10.10.5,10.10.10.252 --rh-ecmp-enable
```

```
--rh-max-path 16 --graceful-restart-enable --routing-type distributed
Vrf created successfully.
```

This example creates a centralized VRF.

```
$ efa tenant vrf create --name red13
--tenant tenant21 --max-path 50 --redistribute connected --redistribute static
--local-asn 65002 --ipv4-static-route-next-hop 10.20.216.104,192.168.0.0/24,10.10.10.1,5
--ipv4-static-route-next-hop 10.20.216.104,192.168.10.0/24,10.10.10.5,5
--ipv6-static-route-next-hop 10.20.216.104,2020:20::1/128,3001::2,6
--ipv6-static-route-next-hop 10.20.216.104,2020:30::1/128,3001::3,5
--ipv6-static-route-bfd 10.20.216.104,3001::3,3001::1,100,200,5
--ipv6-static-route-bfd 10.20.216.104,3001::2,3001::1
--ipv6-static-route-bfd 10.20.216.104,3001::4,3001::1,100,300,6
--ipv4-static-route-bfd 10.20.216.104,10.10.10.1,10.10.10.254,200,300,6
--ipv4-static-route-bfd 10.20.216.104,10.10.10.5,10.10.10.252 --rh-max-path 64
--routing-type centralized --centralized-router 10.20.216.103,10.20.216.104
Vrf created successfully.
```

```
$ efa tenant vrf create --name vrf1 --tenant tenant1

--ipv4-aggregate-address 10.24.80.134,10.20.21.40/30
--ipv4-aggregate-summary-only 10.24.80.134,10.20.21.40/30,true
--ipv4-aggregate-as-set 10.24.80.134,10.20.21.40/30,true
--ipv4-aggregate-advertise-map 10.24.80.134,10.20.21.40/30,some
--ipv4-aggregate-suppress-map 10.24.80.134,10.20.21.40/30,some

--ipv6-aggregate-address 10.24.80.135,10::20/126
--ipv6-aggregate-summary-only 10.24.80.135,10::20/126,true
--ipv6-aggregate-as-set 10.24.80.135,10::20/126,true
--ipv6-aggregate-advertise-map 10.24.80.135,10::20/126,some
--ipv6-aggregate-suppress-map 10.24.80.135,10::20/126,some
```

```
$ efa tenant vrf create --name vrf1 --tenant tenant1
--ipv4-network 10.24.80.134,10.20.30.40/30
--ipv4-network 10.24.80.134,10.21.30.40/30
--ipv4-network-backdoor 10.24.80.134,10.21.30.40/30,true

--ipv4-static-network 10.24.80.134,11.10.30.40/30
--ipv4-static-network 10.24.80.134,11.20.30.40/30
--ipv4-static-network-distance 10.24.80.134,11.20.30.40/30,169

--ipv6-network 10.24.80.135,11::22/128
--ipv6-network 10.24.80.135,11::23/128
--ipv6-network-backdoor 10.24.80.134,11::23/128,true
--ipv6-network 10.24.80.135,11::24/128
--ipv6-network-weight 10.24.80.134,11::24/128,144
--ipv6-network 10.24.80.135,11::25/128
--ipv6-network-route-map 10.24.80.134,11::25/128,rmap1
```

efa tenant vrf delete

Deletes a VRF.

Syntax

```
efa tenant vrf delete [--name vrf name | --tenant tenant name ]
```

Parameters

--name *vrf name*

Specifies the name of the VRF.

--tenant *tenant name*

Specifies the name of the tenant.

Examples

This example deletes the specified VRFs.

```
(efa:extreme)extreme@node-1:~$ efa tenant vrf delete --name red12,red13 --tenant tenant21
Vrf: red12 deleted successfully.
Vrf: red13 deleted successfully.
--- Time Elapsed: 1.530242792s ---
```

efa tenant vrf error show

Displays VRF errors.

Syntax

```
efa tenant vrf error show [ --name vrf-name | --tenant tenant-name ]
```

Parameters

--name *vrf-name*

Specifies the name of the VRF.

--tenant *tenant-name*

Specifies the name of the tenant.

Examples

This example shows output of VRF errors for a specific VRF and tenant.

```
efa tenant vrf error show --tenant tnt1 --name vrf1
=====
Name : vrf1
Tenant : tnt1
Errors
+-----+-----+
| MgmtIp | ErrorList |
+-----+-----+
| 10.20.246.29 | Configure RemoteAsn under Router BGP failed for Vrf : |
| | vrf1 due to Netconf <x> error |
+-----+-----+
| 10.20.246.30 | Configure RemoteAsn under Router BGP failed for Vrf : |
| | vrf1 due to Netconf <x> error |
+-----+-----+
--- Time Elapsed: 195.971ms ---
```


efa tenant vrf show

Displays brief or detailed output of the VRF for all tenants, a selected tenant, or a selected VRF.

Syntax

```
efa tenant vrf show [ --name vrf-name | --tenant tenant-name | --detail ]
```

Parameters

--name *vrf-name*

Specifies the name of the VRF.

--tenant *tenant-name*

Specifies the name of the tenant.

--detail

Displays detailed output of the VRF. When this parameter is not used, only brief tabular output is returned.

Examples

This example shows brief output of all VRFs.

```
$ efa tenant vrf show
+-----+-----+-----+-----+-----+-----+
| Name | Tenant | Routing Type | Centralized Routers | Redistribute | Max Path |
+-----+-----+-----+-----+-----+-----+
| blue11 | tenant11 | distributed | | connected,static | 50 |
+-----+-----+-----+-----+-----+-----+
| red13 | tenant21 | centralized | 10.20.216.103 | connected,static | 50 |
| | | | 10.20.216.104 | | |
+-----+-----+-----+-----+-----+-----+
| red11 | tenant21 | distributed | | connected | 8 |
+-----+-----+-----+-----+-----+-----+

+-----+-----+-----+-----+-----+-----+
| Local Asn | Enable GR | State | Dev State | App State |
+-----+-----+-----+-----+-----+-----+
| 65001 | true | vrf-device-created | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| 65002 | false | vrf-create | not-provisioned | cfg-ready |
| | | | | |
+-----+-----+-----+-----+-----+-----+
| | false | vrf-device-created | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
Vrf Details

--- Time Elapsed: 551.819339ms ---
```

This example shows detailed output of a specific VRF.

```
$ efa tenant vrf show --tenant tenant11 --name blue11 --detail
=====
Name : blue11
Tenant : tenant11
Routing Type : distributed
Centralized Routers :
Redistribute : connected,static
```

```

Max Path                : 50
Local Asn               : 65001
L3VNI                  : 14191
EVPN IRB BD            : 4096
EVPN IRB VE            : 8192
BR VNI                 : 14192
BR BD                  : 4095
BR VE                  : 8191
RH Max Path            : 16
Enable RH ECMP         : true
Enable Graceful Restart : true
Route Target           : import 100:100
                       : export 100:100
                       : import 200:200
                       : export 200:200
                       : import 300:300
                       : export 400:400
Static Route           : Switch-IP->Network, Nexthop-IP[Route-Distance], ...
                       : 10.20.216.15->192.168.10.0/24, 10.10.10.5[5]
2020:30::1/128, 3001::3[5]
                       : 10.20.216.16->192.168.0.0/24, 10.10.10.1[5]
2020:20::1/128, 3001::2[6]
Static Route BFD       : Switch-IP->[DestIP, SourceIP][Interval, Min-Rx, Multiplier], ...
                       : 10.20.216.15->10.10.10.1, 10.10.10.254[200, 300, 6]
                       : 10.20.216.16->10.10.10.5, 10.10.10.252
3001::3, 3001::1[100, 200, 5]
                       : 3001::2, 3001::1 3001::4, 3001::1[100, 300, 6]
State                  : vrf-device-created
Dev State              : provisioned
App State              : cfg-in-sync
=====
--- Time Elapsed: 160.103127ms ---

```

This example shows detailed output of all VRFs belonging to a specific tenant.

```

$ efa tenant vrf show --tenant tenant21 --detail
=====
Name                   : red13
Tenant                 : tenant21
Routing Type           : centralized
Centralized Routers    : 10.20.216.103
                       : 10.20.216.104
Redistribute           : connected, static
Max Path               : 50
Local Asn              : 65002
L3VNI                  :
EVPN IRB BD            :
EVPN IRB VE            :
BR VNI                 :
BR BD                  :
BR VE                  :
RH Max Path            : 64
Enable RH ECMP         : false
Enable Graceful Restart : false
Route Target           :
Static Route           : Switch-IP->Network, Nexthop-IP[Route-Distance], ...
                       : 10.20.216.104->192.168.0.0/24, 10.10.10.1[5] 192.168.10.0/24,
                       : 10.10.10.5[5] 2020:20::1/128, 3001::2[6]
2020:30::1/128, 3001::3[5]
Static Route BFD       : Switch-IP->[DestIP, SourceIP][Interval, Min-Rx, Multiplier], ...
                       : 10.20.216.104->10.10.10.1, 10.10.10.254[200, 300, 6]
10.10.10.5, 10.10.10.252
                       : 3001::3, 3001::1[100, 200, 5] 3001::2, 3001::1

```

```
3001::4,3001::1[100,300,6]
State                : vrf-create
Dev State            : not-provisioned
App State            : cfg-ready

=====
Name                  : red11
Tenant                : tenant21
Routing Type          : distributed
Centralized Routers  :
Redistribute          : connected
Max Path              : 8
Local Asn             :
L3VNI                 : 34191
EVPN IRB BD           : 4094
EVPN IRB VE          : 8190
BR VNI                : 34192
BR BD                 : 4093
BR VE                 : 8189
RH Max Path           :
Enable RH ECMP        : false
Enable Graceful Restart : false
Route Target          : import 101:101
                     : export 101:101
Static Route          :
Static Route BFD      :
State                 : vrf-device-created
Dev State             : provisioned
App State             : cfg-in-sync

=====

--- Time Elapsed: 167.110904ms ---
```

efa tenant vrf update

Updates the tenant VRF for various operations and parameters, such as ASN.

Syntax

```
efa tenant vrf update [--name vrf-name | --tenant tenant-name | --
  operation { local-asn-add | local-asn-delete | static-route-bfd-add |
  static-route-bfd-delete | static-route-add | static-route-delete |
  max-path-add | max-path-delete | redistribute-add | redistribute-
  delete | rh-max-path-add | rh-max-path-delete | centralized-router-
  add | centralized-router-delete | rh-ecmp-update | graceful-restart-
  update } | --local-asn local-asn | --ipv4-static-route-bfd route | --
  ipv6-static-route-bfd route --ipv4-static-route-next-hop route | | --
  ipv6-static-route-next-hop route | --ipv6-network-backdoor
  stringArray | --ipv6-network-weight stringArray | --ipv6-network-
  route-map stringArray --max-path unit | --redistribute { static |
  connected } | --rh-max-path { 8 | 16 | 64 } | --rh-ecmp-enable { true
  | false } | --centralized-router | --operation network-add { network-
  delete | static-network-add | static-network-delete } | --ipv4-
  network stringArray | --ipv4-network-backdoor stringArray | --ipv4-
  network-weight stringArray | --ipv4-network-route-map stringArray |
  --ipv4-network-route-map stringArray | --ipv6-network stringArray |
  --ipv4-static-network stringArray | --ipv4-static-network-distance
  stringArray | --graceful-restart-enable { true | false } | --
  operation aggregate-address-add | aggregate-address-delete | --ipv4-
  aggregate-address stringArray | --ipv4-aggregate-summary-only
  stringArray | --ipv6-aggregate-summary-only stringArray | --ipv6-
  aggregate-address stringArray | --ipv4-aggregate-as-set stringArray |
  --ipv6-aggregate-as-set stringArray | --ipv4-aggregate-advertise-map
  stringArray | --ipv6-aggregate-advertise-map stringArray | --ipv4-
  aggregate-suppress-map stringArray | --ipv6-aggregate-suppress-map
  stringArray ]
```

Parameters

--name *vrf-name*

Specifies the VRF name.

--tenant *tenant-name*

Specifies the tenant name.

--operation { local-asn-add | local-asn-delete | static-route-bfd-add | static-route-bfd-delete | static-route-add | static-route-delete | max-path-add | max-path-delete | redistribute-add | redistribute-delete | rh-max-path-add | rh-max-path-delete | centralized-router-add | centralized-router-delete | rh-ecmp-update | graceful-restart-update }

Identifies the operation you want to perform. Valid values are:

- local-asn-add
- local-asn-delete
- redistribute-add
- redistribute-delete
- centralized-router-add
- centralized-router-delete



Note

Idempotency is supported for the following additional valid operations:

- max-path-add
- max-path-delete
- rh-max-path-add
- rh-max-path-delete
- rh-ecmp-update
- graceful-restart-update
- static-route-add
- static-route-delete
- static-route-bfd-add
- static-route-bfd-delete

--local-asn *local-asn*

Specifies the local ASN for the VRF.

--ipv4-static-route-bfd *route*

Specifies the IPv4 static route BFD in the following format: device IP, destination IPv4 address, source IPv4 address[interval, min-rx, multiplier]. For example:

10.25.25.100,1.1.1.1,2.2.2.2,123,456,3.

--ipv6-static-route-bfd *route*

Specifies the IPv6 static route BFD in the following format: device IP, destination IPv6 address, source IP address[interval, min-rx, multiplier]. For example:

10.25.25.100,1:::1,2:::2,300,300,3.

--ipv4-static-route-next-hop *route*

Specifies the IPv4 static route next hop in the following format: device IP, IPv4 static route network, next hop IP, and route distance separated by commas. For example:

10.25.25.100,20.0.0.0/24,16.0.0.2. Valid values for the route distance are 1 through 254.

--ipv6-static-route-next-hop *route*

Specifies the IPv6 static route next hop in the following format: device IP, IPv6 static route network, next hop IP, and route distance separated by commas. For example:

10.25.25.100,2001:1:::/64,3001:::2,3. Valid values for the route distance are 1 through 254.

--max-path *unit*

Specifies the number of load-sharing paths for the VRF. Valid values are 1 through 64.

--redistribute { **static** | **connected** }

Specifies the redistribute type for routes. Valid values are static or connected.

--rh-max-path { **8** | **16** | **64** }

Specifies the maximum number of resilient hashing paths allowed per tenant VRF. Valid values are 8, 16, or 64.

--rh-ecmp-enable { **true** | **false** }

Turns on or turns off resilient hashing for a tenant VRF. Valid values are true or false.

--routing-type { **distributed** | **centralized** }

VRF routing type. Default value is distributed.

--centralized-router

Comma-separated list of border-leaf IP addresses.

--graceful-restart-enable { **true** | **false** }

Turns on or turns off graceful restart for a tenant VRF. Valid values are true or false.

--ipv4-network *stringArray*

BGP advertises given network address. Format: device-ip,network-address. Example: `--ipv4-network 10.24.80.134,10.20.30.0/30`.

--ipv4-network-backdoor *stringArray*

Increases the administrative distance of eBGP with the goal of making IGP learned routes preferred for a given network address. Format: device-ip,network-address,flag. Example: `--ipv4-network-backdoor 10.24.80.134,10.21.30.0/30,true`.

--ipv4-network-weight *stringArray*

Given weight is used to set the BGP weight attribute for the given network address. Format: device-ip,network-address,weight. Example: `--ipv4-network-weight 10.24.80.134,10.22.30.0/30,144`.

--ipv4-network-route-map *stringArray*

Given route-map is used to set the BGP attributes for the given network address like MED. Format: device-ip,network-address,rm. Example: `--ipv4-network-route-map 10.24.80.134,10.23.30.0/30,routeMap1`.

--ipv6-network *stringArray*

BGP advertises given network address. Format: device-ip,network-address. Example: `--ipv6-network 10.24.80.134,11::/128`.

--ipv6-network-backdoor *stringArray*

Increases the administrative distance of eBGP with the goal of making IGP learned routes preferred for given network address. Format:device-ip,network-address,flag. Example: `--ipv6-network-backdoor 10.24.80.134,11::/128,true`.

--ipv6-network-weight *stringArray*

Given weight is used to set the BGP weight attribute for the given network address. Format: device-ip,network-address,weight. Example: `--ipv6-network-weight 10.24.80.134,11::/128,144`.

--ipv6-network-route-map *stringArray*

Given route map is used to set the BGP attributes for the given network address like MED.
Format:device-ip,network-address,rm. Example: --ipv6-network-route-map
10.24.80.134,11::/128,rmap1.

--ipv4-static-network *stringArray*

IPv4 static network to be configured on the device. Format: Device IP, IPv4 Static Network.
Example: 10.24.80.134,11.10.30.40/30.

--ipv4-static-network-distance *stringArray*

IPv4 static network with distance to be configured on the device. Format: Device IP, IPv4 Static Network, Distance(valid values are <1-255>). Example:
10.24.80.134,11.10.30.40/30,169.

--ipv4-aggregate-address *stringArray*

IPv4 aggregate-address in the format device-ip,ipv4-aggregate-address Example:
10.25.25.100,10.20.21.40/30.

--ipv6-aggregate-address *stringArray*

IPv6 aggregate-address in the format device-ip,ipv6-aggregate-address Example:
10.25.25.100,10::20/126.

--ipv4-aggregate-summary-only *stringArray*

Enable to advertise only aggregated-address. Format: device-ip,ipv4-aggregate-address,ipv4-aggregate-summary-only. Example: 10.25.25.100,10.20.21.40/30,true.

--ipv6-aggregate-summary-only *stringArray*

Enable to advertise only aggregated-address. Format: device-ip,ipv6-aggregate-address,ipv6-aggregate-summary-only. Example: 10.25.25.100,10::20/126,true.

--ipv4-aggregate-as-set *stringArray*

Enable to set AS set path information as part of aggregate-address advertisement. Format:
device-ip,ipv4-aggregate-address,ipv4-aggregate-as-set. Example:
10.25.25.100,10.20.21.40/30,true.

--ipv6-aggregate-as-set *stringArray*

Enable to set AS set path information as part of aggregate-address advertisement. device-
ip,ipv6-aggregate-address,ipv6-aggregate-as-set. Example:
10.25.25.100,10::20/126,true.

--ipv4-aggregate-advertise-map *stringArray*

Route map name used to filter the BGP attributes to be advertised as part of aggregate-address.
Format: device-ip,ipv4-aggregate-address,ipv4-aggregate-advertise-map. Example:
10.25.25.100,10.20.21.40/30,routeMap1.

--ipv6-aggregate-advertise-map *stringArray*

Route map name used to filter the BGP attributes to be advertised as part of aggregate-address.
Format: device-ip,ipv6-aggregate-address,ipv6-aggregate-advertise-map. Example:
10.25.25.100,10::20/126,routeMap1.

--ipv4-aggregate-suppress-map *stringArray*

Route map name used to suppress the specific routes to be advertised along with the
aggregate-address advertisement, in the following format: device-ip,ipv4-aggregate-

address,ipv4-aggregate-suppress-map. For example:
10.25.25.100,10.20.21.40/30,routeMap1.

--ipv6-aggregate-suppress-map *stringArray*

Route map name used to suppress the specific routes to be advertised, along with the aggregate-address advertisement in the following format: device-ip,ipv6-aggregate-address,ipv6-aggregate-suppress-map. For example:
10.25.25.100,10::20/126,routeMap1.

Usage Guidelines

The **--max-path** and **--rh-max-path** parameters can co-exist.

You cannot choose the specific devices on which to configure resilient hashing. Configuration applies to all SLX devices in the tenant VRF.

Examples

This example adds redistribution connected to a VRF.

```
$ efa tenant vrf update
--name blue11 --tenant tenant11 --operation redistribute-add --redistribute connected

Vrf updated successfully.
```

This example removes the max-path from the VRF.

```
$ efa tenant vrf update
--name blue12 --tenant tenant1 --operation rh-max-path-delete

Vrf updated successfully.
```

```
$ efa tenant vrf update --name vrf1 --tenant tenant1
--operation aggregate-address-add
--ipv4-aggregate-address 10.24.80.134,10.21.21.40/30
--ipv4-aggregate-summary-only 10.24.80.134,10.21.21.40/30,true
--ipv4-aggregate-as-set 10.24.80.134,10.21.21.40/30,true
--ipv4-aggregate-advertise-map 10.24.80.134,10.21.21.40/30,some
--ipv4-aggregate-suppress-map 10.24.80.134,10.21.21.40/30,some
```

```
$ efa tenant vrf update --name vrf1 --tenant tenant1
--operation network-add
--ipv4-network 10.24.80.134,10.22.30.40/30
--ipv4-network-weight 10.24.80.134,10.22.30.40/30,144
--ipv4-network 10.24.80.134,10.23.30.40/30
--ipv4-network-route-map 10.24.80.134,10.23.30.40/30,rmap1
```


efa vcenter debug

A set of debug commands for troubleshooting EFA vCenter issues. This command also sets configuration parameters for vCenter tenants.

Syntax

efa vcenter debug tenant show --host *string*

Shows a list of tenants for a particular VMware vCenter server host.

efa vcenter debug event show --host *string* [**--page** *int32* | **--page-size** *int32*]

Shows a list of events recorded for a particular VMware vCenter server host.

efa vcenter debug setting show

Shows the different settings configured on the VMware vCenter server host. These settings are global in nature and are applicable to all VMware vCenter Server hosts registered with this EFA vCenter Service instance.

efa vcenter debug setting update [**--poll-frequency** *string* | **--dead-link-clearing-time** *string*]

Updates the different settings for the added VMware vCenter server hosts registered with this EFA vCenter Service instance. These settings are applicable to all the VMware vCenter Server hosts.

efa vcenter debug set --level [*info* | *debug*]

Sets the debug level for this instance of EFA vCenter service. Debug levels can be set to receiving information level messages or complete debugging messages.

Command Default

This command has no defaults.

Parameters

--host *string*

IP address or hostname of the VMware vCenter Server to connect to.

--page *int32*

Events are fetched for the page number specified in this parameter. When this parameter is not passed, page number 1 is always fetched by default.

--page-size *int32*

The number of events to display per page. When this parameter is not passed, twenty (20) records are shown per page.

--poll-frequency *string*

The poll duration in hours. The EFA vCenter Service polls all the VMware vCenter servers registered with it after this time duration has expired.

--dead-link-clearing-time *string*

The time duration in days. This is the time duration for which dead links are removed from the EFA vCenter service database.

--level [*info* | *debug*]

Sets the debug level for this EFA vCenter service instance. Can be one of *info* or *debug*. Use *info* to view only messages of the level *info*. Use *debug* to view all debug messages.

Examples

The following example returns the debug information for tenant creation failure (and any other reported errors) for a vCenter server with IP address 10.24.85.111. This example shows only a portion of the possible output.

```
$ efa vcenter debug tenant --host 10.24.85.111
Tenant Configuration Details for vCenter 10.24.85.111
+-----+-----+-----+-----+
+-----+
| Tenant Name          | Ports | VLANs | Status          |
Reason                |
+-----+-----+-----+-----+
| vcenter-10.24.85.111 |      | 2-4089 | Creation Failed | Tenant vcenter-10.24.85.111
is missing the      |
|                    |      |      |                  | following interfaces:
10.25.225.46="ethernet-0/5"
+-----+-----+-----+-----+
+-----+
Tenant Details

EPG Configuration Details for vCenter 10.24.85.111
+-----+-----+-----+-----+
+-----+
| EPG Name              | Ports                | Switchport
Mode |
+-----+-----+-----+-----+
| vCenter_10.24.85.111_10.24.82.20_vSwitch1_VMNetwork2 | 10.24.82.20[0/11] |
trunk |
+-----+-----+-----+-----+
+-----+
```

This example shows the vCenter service settings.

```
$ efa vcenter debug setting show

vCenter service settings
+-----+-----+-----+-----+
| Poll Frequency Hours | Dead Link Clearing Time Days |
+-----+-----+-----+-----+
| 4                    | 3                            |
+-----+-----+-----+-----+
```

This example shows tenant and endpoint group information for the specified IP address.

```
$ efa vcenter debug tenant --host 10.24.85.111
```

Tenant Name	Ports	VLANS	Status	Reason
vcenter-10.24.85.111		2-4089	Created Successfully	

Tenant Details

EPG Configuration Details for vCenter 10.24.85.111

EPG Name	Ports	POs	Switchport Mode	VLANS	Status	Reason
----------	-------	-----	-----------------	-------	--------	--------

This example shows event details for the specified host. This example shows only a portion of the possible output.

```
$ efa vcenter debug event --host 10.24.85.111
```

Task: Update network configuration	10.24.81.10
dvPort group VM network 200-220 in SRA-Dev-DC was reconfigured. Modified: config.defaultPortConfig.vlan: (inherited = false, vlanId = 200) -> (inherited = false, vlanId = ((start = 0, end = 200))); Added: Deleted:	VM network 200-220
Task: Update network configuration	10.24.82.20
dvPort group VM network 200-220 in SRA-Dev-DC was reconfigured. Modified: config.defaultPortConfig.vlan: (inherited = false, vlanId = ((start = 200, end = 220))) -> (inherited = false, vlanId = 200); Added: Deleted:	VM network 200-220

The following example shows typical output for the specified host.

```
$ efa vcenter debug tenant --host 10.24.85.111
```

Tenant Configuration Details for vCenter 10.24.85.111

Tenant Name	Ports	VLANS	Status	Reason
vcenter-10.24.85.111		2-4089	Created Successfully	

Tenant Details

EPG Configuration Details for vCenter 10.24.85.111

EPG Name	Ports	POs	Switchport Mode	VLANS	Status	Reason
----------	-------	-----	-----------------	-------	--------	--------

This example shows typical, but truncated, output for debugging events for the specified host.

```
$ efa vcenter debug event --host 10.24.85.111

vCenter Events Total Count (4) for host 10.24.85.111
+-----+-----+
| Description | Target |
+-----+-----+
| Task: Update network configuration | 10.24.81.10 |
+-----+-----+
| dvPort group VM network 200-220 in SRA-Dev-DC | VM network 200-220 |
| was reconfigured. Modified: | |
| config.defaultPortConfig.vlan: (inherited = false, vlanId | |
| = 200) -> (inherited = false, vlanId = ((start = 0, end = | |
| 200)); Added: Deleted: | |
+-----+-----+
| Task: Update network configuration | 10.24.82.20 |
+-----+-----+
| dvPort group VM network 200-220 in SRA-Dev-DC | VM network 200-220 |
| was reconfigured. Modified: | |
| config.defaultPortConfig.vlan: (inherited = false, vlanId = | |
| ((start = 200, end = 220)) -> (inherited = false, vlanId = | |
| 200); Added: Deleted: | |
+-----+-----+
```

This command sets the polling frequency to 5 minutes.

```
$ efa vcenter debug setting update --poll-frequency=5

vCenter service settings updated successfully
```

This example command sets the dead-link-clearing-time value to 4 days.

```
$ efa vcenter debug setting update --dead-link-clearing-time=4

vCenter service settings updated successfully
```

This example sets the log level to debug for this EFA instance.

```
$ efa vcenter debug set --level=debug

Level debug set successfully.
```

efa vcenter delete

Deletes the VMware vCenter server registered as a tenant with this EFA instance.

Syntax

```
efa vcenter delete --host string [ --cleanup-tenants ]
```

Parameters

--host *string*

IP address or hostname of the VMware vCenter server tenant to delete.

--cleanup-tenants

When included, the tenants and EPGs associated to the VMware vCenter Server are deleted.

When not included, the tenants and EPGs associated with the VMware vCenter Server are retained when the server is removed.

When the `cleanup-tenants` parameter is supplied, the command removes the registered ESXi hosts and EPGs for this VMware vCenter Server tenant.

Examples

This example deletes the VMware vCenter Server with host IP 10.24.85.111 while retaining the associated tenants and EPGs.

```
$ efa vcenter delete --host 10.24.85.111
```

This example deletes the VMware vCenter Server with host IP 10.24.85.111 and removes the associated tenants and EPGs.

```
$ efa vcenter delete --host 10.24.85.111 --cleanup-tenants
```

efa vcenter links

Displays the links for VMware vCenter server tenants and their ESXi hosts.

Syntax

```
efa vcenter links [ physical | unconnected ] [ --host string | --esxi string ]
```

```
efa vcenter links virtual --esxi string
```

Parameters

physical

Displays the physical links for the selected device. Device can be a VMware vCenter Server or an ESXi server.

unconnected

Displays the unconnected (unused) links for the selected device. Device can be a VMware vCenter Server or an ESXi server.

virtual

Displays the virtual links created on a ESXi server.

--host

The IP address or host name of the VMware vCenter Server for which to view the physical and unconnected (unused) links.

--esxi

The IP address or host name of the ESXi tenant for which to view the physical, virtual, and unconnected (unused) links.

For VMware vCenter server tenants, physical and unconnected links are displayed. For ESXi hosts, physical, unconnected, and virtual links are displayed.

Examples

This example is list of ESXi device managed by the VMware vCenter Server with the IP 10.24.85.111.

```
$ efa vcenter links physical --host 10.24.85.111
Physical Links for vCenter 10.24.85.111
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
| ESXi Host   | PNIC   | Driver | PNIC MAC           | Interface   | PO Number | PO
Name | Device       | Missing | Missing Time |              |           |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
| 10.24.81.10 | vmnic3 | ntg3   | 40:f2:e9:bb:4c:83 | Ethernet 0/10 |
|             | 10.9.9.20 |       |                   |               |           |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
| 10.24.82.20 | vmnic2 | ntg3   | 40:f2:e9:bb:58:b2 | Ethernet 0/10 |
|             | 10.9.9.30 |       |                   |               |           |
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+
| 10.24.83.30 | vmnic3 | ntg3   | 40:f2:e9:bb:58:b3 | Ethernet 0/10 |
```

```

|          | 10.9.9.40 |          |          |
+-----+-----+-----+-----+-----+-----+
Physical Link Details
--- Time Elapsed: 359.1528ms ---

```

This example is a list of virtual machines and their virtual links on the ESXi host with the IP 10.24.81.10.

```

$ efa vcenter links virtual --esxi 10.24.81.10
Virtual Links for ESXi server 10.24.81.10
+-----+-----+-----+-----+-----+-----+
| Virtual Machine | vNIC          | MAC          | VLAN | Port Group |
vSwitch          | Distributed |
+-----+-----+-----+-----+-----+-----+
| 20200221 v1     | Network adapter 1 | 00:50:56:9c:75:29 | 3048 | VM Network |
vSwitch0         | false          |
+-----+-----+-----+-----+-----+-----+
| test_vm2       | Network adapter 2 | 00:50:56:9c:7e:83 | 3048 | VM Network |
vSwitch0         | false          |
+-----+-----+-----+-----+-----+-----+
| 20200221 v1     | Network adapter 2 | 00:50:56:9c:f6:13 | 3048 | VM Network |
vSwitch0         | false          |
+-----+-----+-----+-----+-----+-----+
| vMotionVm      | Network adapter 1 | 00:50:56:9c:8f:fd | 10   | DPortGroup | Real-DV-
Switch | true          |
+-----+-----+-----+-----+-----+-----+
| test_vm2       | Network adapter 1 | 00:50:56:9c:1d:09 | 10   | DPortGroup | Real-DV-
Switch | true          |
+-----+-----+-----+-----+-----+-----+
vNIC Link Details
--- Time Elapsed: 500.4792ms ---

```

efa vcenter list

Lists the VMware vCenter servers registered as tenants.

Syntax

```
efa vcenter list [ --host string ]
```

Parameters

--host *string*

IP address or hostname of the VMware vCenter Server to connect to.

When the `host` parameter is supplied, the command lists the ESXi hosts for that VMware vCenter Server.

Examples

This example lists a VMware vCenter Server that has just been added as a tenant with this EFA instance.

```
$ efa vcenter list
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Host Name      | Name                | Version | Discovery Status | Datacenters Count |
| Host Count    | VM Count           |         |                  |                   |
+-----+-----+-----+-----+-----+
| 10.24.85.111  | VMware vCenter Server | 6.7.0   | Pending Complete Discovery | 1 |
| 1              | 1                   | 0       |                          | 0 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```

This example lists a VMware vCenter Server with its status updated to this EFA instance.

```
$ efa vcenter list
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Host Name      | Name                | Version | Discovery Status | Datacenters Count |
| Host Count    | VM Count           |         |                  |                   |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.24.85.111  | VMware vCenter Server | 6.7.0   | Update Completed | 5 |
| 4              | 13                  |         |                          | 4 |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
```

This example lists the ESXi devices managed by the VMware vCenter Server with IP address 10.24.85.111.

```
$ efa vcenter list --host 10.24.85.111
```


efa vcenter register

Registers a VMware vCenter server as a tenant with the EFA instance.

Syntax

```
efa vcenter register [ --host string | --username string | --password string | --tenant string ]
```

Command Default

This command has no defaults. All flags are mandatory.

Tenant details must be registered before running the command.

Parameters

--host *string*

IP address or hostname of the VMware vCenter Server to connect to.

--username *string*

Username to connect to the VMware vCenter Server.

--password *string*

Password to connect to the VMware vCenter Server.

--tenant *string*

Tenant associated with the VMware vCenter Server host.

Examples

This example registers a VMware vCenter Server as a tenant with EFA vCenter service.

```
$ efa vcenter register --host 10.24.85.111 --username administrator@vmvcenter.local  
--password 'aBc123#' --tenant tenant-10.24.85.111
```

efa vcenter update

Updates the local database with configuration changes made to the VMware vCenter server and the ESXi servers managed by it.

Syntax

```
efa vcenter update [ --host string | --username string | --password string ]
```

Command Default

This command has no defaults. All flags are mandatory.

Parameters

--host *string*

IP address or host name of the VMware vCenter server from which the local EFA vCenter service database needs to be updated.

--username *string*

Username to connect to the VMware vCenter Server.

--password *string*

Password to connect to the VMware vCenter Server.

Examples

This example updates the local EFA vCenter service database with the latest state of the VMware vCenter server with the IP address 10.24.85.111.

```
$ efa vcenter update --host 10.24.85.111 --username administrator@vmvcenter.local --password 'aBc123#'
```

efa version

Displays the version of EFA that is installed.

Syntax

```
efa version
```

Examples

The following example shows typical output for the command.

```
# efa version
Version : 2.2.0
Build: 41
Time Stamp: 20-06-03:03:15:47
Mode: Insecure
Deployment Type: single-node
Deployment Platform: SERVER
--- Time Elapsed: 8.986817ms ---
```

efa-change-hostname

Changes the host name in the EFA system after EFA is deployed.

Syntax

```
efa-change-hostname old-host-name
```

Parameters

old-host-name

Specifies the host name that you are changing.

Usage Guidelines

Host name changes are supported in single-node and multi-node deployments.

Running this command by itself does not change the host name. For the complete procedure, see the "Change the Host Name or IP Address" topic in the "System Management" section of the *Extreme Fabric Automation Administration Guide*.

Run this command as a root user or as a user with sudo privileges.

In a TPVM deployment, run this command from `/apps/bin/`.

In a single-node deployment, EFA is not operational while this command is running. In a multi-node deployment, EFA remains operational if the command is running on the standby node. EFA is not operational if the command is running on the active node.

Examples

This example shows standard output for the command.

```
$ sudo efa-change-hostname <old-host-name>

Reading host name of the system
Restarting mariadb service
Restarting k3s service
Checking k3s for the new host name
Host is in ready state in k3s
Setting current host as active node
Deleting old host name references
Waiting for EFA containers to start
Successfully updated host name in EFA
```

efa-change-ip

Changes the IP address after EFA is deployed.

Syntax

```
efa-change-ip
```

Usage Guidelines

IP address changes are supported in single-node deployments.

Run this command as a root user or as a user with sudo privileges.

In a TPVM deployment, run the command from `/apps/bin/`.

EFA is not operational while this command is running.

After the IP address is changed, either run `source /etc/profile` or open a new EFA session to log in.

Examples

This example shows standard output for the command.

```
$ sudo efa-change-ip

Updating IP in EFA
Restarting k3s service
Updating all files with new IP
Deleting EFA services: gonotification-service gofabric-service gotenant-service
goauth-service gorbac-service goinventory-service goopenstack-service
govcenter-service gohyperv-service goraslog-service efa-api-docs gosystem-service
Waiting for EFA containers to start
Successfully updated IP in EFA
```

efa-change-vip

Changes the virtual IP address (VIP) after EFA is deployed.

Syntax

```
efa-change-vip new-vip
```

Parameters

new-vip

Specifies the new VIP for the multi-node deployment.

Usage Guidelines

VIP changes are supported in multi-node deployments.

Run this command as a root user or as a user with sudo privileges.

In a TPVM deployment, run this command from `/apps/bin/`.

EFA is not operational while this command is running.

After the VIP is changed, run `source /etc/profile` or open a new EFA session to log in.

Examples

This example shows standard output for the command.

```
$ sudo efa-change-vip <new vip>
Updating all files with new VIP.
Restarting services on nodes
Waiting for EFA containers to start
Updating services with new VIP
Waiting for EFA containers to start
Successfully updated VIP for the installation.
```

efactl

Starts, stops, restarts, and displays the status services for EFA.

Syntax

```
efactl start  
efactl stop  
efactl restart  
efactl start-service service-name  
efactl restart-service service-name  
efactl clean  
efactl status  
efactl stop-service service-name  
efactl db-status  
efactl rabbit-status
```

Examples

The following example starts all services.

```
# efactl start  
Are you sure you want to start all services? [Y/n]  
Y  
gosystem-service has been started  
goauth-service has been started  
gorbac-service has been started  
godb-service has been started  
gonotification-service has been started  
goinventory-service has been started  
gofabric-service has been started  
gotenant-service has been started  
goopenstack-service has been started  
gohyperv-service has been started  
govcenter-service has been started  
goswitch-service has been started  
goraslog-service has been started  
efa-api-docs has been started  
rabbitmq has been started  
Services have been started
```

The following example stops all services.

```
# efactl stop  
Are you sure you want to stop all services? [Y/n]  
Y  
gosystem-service has been stopped  
goauth-service has been stopped  
gorbac-service has been stopped  
godb-service has been stopped  
gonotification-service has been stopped  
goinventory-service has been stopped  
gofabric-service has been stopped
```

```
gotenant-service has been stopped
goopenstack-service has been stopped
gohyperv-service has been stopped
govcenter-service has been stopped
goswitch-service has been stopped
goraslog-service has been stopped
efa-api-docs has been stopped
rabbitmq has been stopped
Services have been stopped
root@administrator-08:~#
```

The following example restarts all services.

```
efactl restart
Are you sure you want to restart all services? [Y/n]
Y
gosystem-service has been stopped
goauth-service has been stopped
gorbac-service has been stopped
godb-service has been stopped
gonotification-service has been stopped
goinventory-service has been stopped
gofabric-service has been stopped
gotenant-service has been stopped
goopenstack-service has been stopped
gohyperv-service has been stopped
govcenter-service has been stopped
goswitch-service has been stopped
goraslog-service has been stopped
efa-api-docs has been stopped
rabbitmq has been stopped
Services have been stopped
gosystem-service has been started
goauth-service has been started
gorbac-service has been started
godb-service has been started
gonotification-service has been started
goinventory-service has been started
gofabric-service has been started
gotenant-service has been started
goopenstack-service has been started
gohyperv-service has been started
govcenter-service has been started
goswitch-service has been started
goraslog-service has been started
efa-api-docs has been started
rabbitmq has been started
Services have been started
All services have been restarted
```

The following example starts one service.

```
# efactl start-service gofabric-service
Are you sure you want to start gofabric-service? [Y/n]
Y
gofabric-service has been started
```

The following example stops one service.

```
# efactl stop-service gorbac-service
Are you sure you want to stop gorbac-service? [Y/n]
Y
gorbac-service has been stopped
```


The following example restarts one service.

```
# efactl restart-service rabbitmq
Are you sure you want to restart rabbitmq? [Y/n]
Y
rabbitmq has been stopped
rabbitmq has been started
rabbitmq has been restarted
```

The following example deletes all logs.

```
# efactl clean
This will delete all logs. Are you sure [Y/n]?
Y
Cleaning efa Logs...
Cleaned logs
```

The following example shows EFA status.

```
(efa:extreme)extreme@tpvm:~$ efactl status
NAME      STATUS  ROLES    AGE  VERSION          LABELS
tpvm      Ready   primary  50m  v1.18.6+k3s1    beta.kubernetes.io/
arch=amd64,beta.kubernetes.io/os=linux,keepalived=active,kubernetes.io/
arch=amd64,kubernetes.io/hostname=tpvm,kubernetes.io/os=linux,node-role.kubernetes.io/
primary=true
tpvm2     Ready   primary  50m  v1.18.6+k3s1    beta.kubernetes.io/
arch=amd64,beta.kubernetes.io/os=linux,keepalived=standby,kubernetes.io/
arch=amd64,kubernetes.io/hostname=tpvm2,kubernetes.io/os=linux,node-role.kubernetes.io/
primary=true
NAME      READY  STATUS    RESTARTS  AGE
IP        NODE   NOMINATED NODE  READINESS GATES
pod/gofabric-service-cb547d7db-1k6r9          0/1      Init:0/2   0          45m
10.42.2.8   tpvm2 <none>      <none>
pod/gonotification-service-ffccd5bf6-wl96v    0/1      Init:0/2   0          45m
10.20.255.114 tpvm2 <none>      <none>
pod/goauth-service-5fdd67574-g6jw5           0/1      Init:0/2   0          45m
10.42.2.9   tpvm2 <none>      <none>
pod/gotenant-service-7db6648b97-rvft9         0/1      Init:0/3   0          45m
10.42.2.10  tpvm2 <none>      <none>
pod/efa-api-docs-5886cb6fb4-vnq8x            1/1      Running    0          45m
10.42.2.12  tpvm2 <none>      <none>
pod/efa-api-docs-5886cb6fb4-h7lrb            1/1      Running    0          45m
10.42.0.11  tpvm  <none>      <none>
pod/gosnmp-service-56dc6c46c8-hd8nk          0/1      Init:0/1   0          45m
10.20.255.114 tpvm2 <none>      <none>
pod/goopenstack-service-798b445769-sgdc6     0/1      Init:0/2   0          45m
10.42.2.13  tpvm2 <none>      <none>
pod/goinventory-service-5c7f7b57c4-txr9m      0/1      Init:0/2   0          45m
10.42.2.14  tpvm2 <none>      <none>
pod/gosnmp-service-56dc6c46c8-vbtbv          1/1      Running    0          45m
10.20.255.230 tpvm  <none>      <none>
pod/govcenter-service-6ff99bf8b8-dx2dk       0/1      Init:0/2   0          45m
10.42.2.15  tpvm2 <none>      <none>
pod/goopenstack-service-798b445769-v7dzzr    1/1      Running    0          45m
10.42.0.13  tpvm  <none>      <none>
pod/gohyperv-service-7f8d87d55-xspkg         0/1      Init:0/2   0          45m
10.42.2.16  tpvm2 <none>      <none>
pod/gosystem-service-67767c4796-n7rc4        1/1      Running    0          45m
10.20.255.114 tpvm2 <none>      <none>
pod/gosystem-service-67767c4796-14w4m        1/1      Running    0          45m
10.20.255.230 tpvm  <none>      <none>
pod/rabbitmq-0                                1/1      Running    0          46m
10.42.0.6   tpvm  <none>      <none>
pod/gorbac-service-69c665497b-bdlkh          0/1      Init:1/2   0          45m
```

10.42.2.11	tpvm2	<none>	<none>			
pod/goauth-service-5fdd67574-5h4rx			1/1	Running	0	45m
10.42.0.9	tpvm	<none>	<none>			
pod/gorbac-service-69c665497b-slkr2			1/1	Running	0	45m
10.42.0.10	tpvm	<none>	<none>			
pod/goinventory-service-5c7f7b57c4-68b7w			1/1	Running	0	45m
10.42.0.12	tpvm	<none>	<none>			
pod/gonotification-service-ffccd5bf6-7551d			1/1	Running	0	45m
10.20.255.230	tpvm	<none>	<none>			
pod/goraslog-service-6d89f667bd-6d46n			1/1	Running	0	45m
10.20.255.114	tpvm2	<none>	<none>			
pod/gofabric-service-cb547d7db-nnt1b			1/1	Running	0	45m
10.42.0.7	tpvm	<none>	<none>			
pod/gotenant-service-7db6648b97-6wdht			1/1	Running	0	45m
10.42.0.8	tpvm	<none>	<none>			
pod/goraslog-service-6d89f667bd-jgkww			1/1	Running	0	45m
10.20.255.230	tpvm	<none>	<none>			
pod/rabbitmq-1			1/1	Running	0	43m
10.42.2.17	tpvm2	<none>	<none>			
pod/govcenter-service-6ff99bf8b8-lvm46			1/1	Running	2	45m
10.42.0.14	tpvm	<none>	<none>			
pod/gohyperv-service-7f8d87d55-bqz1z			1/1	Running	2	45m
10.42.0.15	tpvm	<none>	<none>			

NAME	AGE	TYPE	CLUSTER-IP	EXTERNAL-IP
service/rabbitmq		NodePort	10.43.215.4	<none>
15672:31672/TCP,5672:30672/TCP	46m	app=rabbitmq		
service/gofabric-service		ClusterIP	10.43.76.230	<none>
8081/TCP	46m	app=gofabric-service		
service/goinventory-service		ClusterIP	10.43.63.133	<none>
8082/TCP	46m	app=goinventory-service		
service/goinventory-cluster-service		ClusterIP	10.43.55.0	<none>
8082/TCP	46m	app=goinventory-service		
service/gotenant-service		ClusterIP	10.43.51.68	<none>
8083/TCP	46m	app=gotenant-service		
service/efa-api-docs		ClusterIP	10.43.243.140	<none>
80/TCP	46m	app=efa-api-docs		
service/goraslog-service		ClusterIP	10.43.207.123	<none>
8091/TCP	46m	app=goraslog-service		
service/gosnmp-service		ClusterIP	10.43.122.218	<none>
8092/TCP	46m	app=gosnmp-service		
service/gonotification-service		ClusterIP	10.43.95.30	<none>
8088/TCP	46m	app=gonotification-service		
service/goauth-service		ClusterIP	10.43.106.23	<none>
8080/TCP	46m	app=goauth-service		
service/gorbac-service		ClusterIP	10.43.10.15	<none>
8089/TCP	46m	app=gorbac-service		
service/goopenstack-service		ClusterIP	10.43.91.98	<none>
8085/TCP	45m	app=goopenstack-service		
service/govcenter-service		ClusterIP	10.43.25.206	<none>
8086/TCP	45m	app=govcenter-service		
service/gohyperv-service		ClusterIP	10.43.44.224	<none>
8087/TCP	45m	app=gohyperv-service		
service/gosystem-service		ClusterIP	10.43.210.54	<none>
8090/TCP	45m	app=gosystem-service		

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
CONTAINERS				
IMAGES				
deployment.apps/efa-api-docs	2/2	2	2	46m
docs				
efa-api-docs:2.4.0		app=efa-api-docs		
deployment.apps/gosystem-service	2/2	2	2	45m
gosystem				
gosystem:2.4.0		app=gosystem-service		
deployment.apps/goraslog-service	2/2	2	2	45m
goraslog-				

service	goraslog:2.4.0	app=goraslog-service				
deployment.apps/gosnmp-service		1/2 2 1	45m	gosnmp-		
service	gosnmp:2.4.0	app=gosnmp-service				
deployment.apps/goopenstack-service		1/2 2 1	45m			
openstack	goopenstack:2.4.0	app=goopenstack-service				
deployment.apps/goauth-service		1/2 2 1	46m	go-		
auth	goauth:2.4.0	app=goauth-service				
deployment.apps/gorbac-service		1/2 2 1	46m	go-		
rbac	gorbac:2.4.0	app=gorbac-service				
deployment.apps/goinventory-service		1/2 2 1	46m			
goinventory-service	goinventory:2.4.0	app=goinventory-service				
deployment.apps/gonotification-service		1/2 2 1	46m			
gonotification-service	gonotification:2.4.0	app=gonotification-service				
deployment.apps/gofabric-service		1/2 2 1	46m	gofabric-		
service	gofabric:2.4.0	app=gofabric-service				
deployment.apps/gotenant-service		1/2 2 1	46m	gotenant-		
service	gotenant:2.4.0	app=gotenant-service				
deployment.apps/govcenter-service		1/2 2 1	45m			
vcenter	govcenter:2.4.0	app=govcenter-service				
deployment.apps/gohyperv-service		1/2 2 1	45m			
hyperv	gohyperv:2.4.0	app=gohyperv-service				
NAME		DESIRED	CURRENT	READY	AGE	
CONTAINERS	IMAGES	SELECTOR				
replicaset.apps/efa-api-docs-5886cb6fb4		2	2	2	45m	efa-
api-docs	efa-api-docs:2.4.0	app=efa-api-docs,pod-template-hash=5886cb6fb4				
replicaset.apps/gosnmp-service-56dc6c46c8		2	2	1	45m	
gosnmp-service	gosnmp:2.4.0	app=gosnmp-service,pod-template-				
hash=56dc6c46c8						
replicaset.apps/goopenstack-service-798b445769		2	2	1	45m	
openstack	goopenstack:2.4.0	app=goopenstack-service,pod-template-				
hash=798b445769						
replicaset.apps/gosystem-service-67767c4796		2	2	2	45m	
gosystem	gosystem:2.4.0	app=gosystem-service,pod-template-				
hash=67767c4796						
replicaset.apps/goauth-service-5fdd67574		2	2	1	46m	go-
auth	goauth:2.4.0	app=goauth-service,pod-template-				
hash=5fdd67574						
replicaset.apps/gorbac-service-69c665497b		2	2	1	45m	go-
rbac	gorbac:2.4.0	app=gorbac-service,pod-template-				
hash=69c665497b						
replicaset.apps/goinventory-service-5c7f7b57c4		2	2	1	46m	
goinventory-service	goinventory:2.4.0	app=goinventory-service,pod-template-				
hash=5c7f7b57c4						
replicaset.apps/gonotification-service-ffccd5bf6		2	2	1	46m	
gonotification-service	gonotification:2.4.0	app=gonotification-service,pod-template-				
hash=ffccd5bf6						
replicaset.apps/gofabric-service-cb547d7db		2	2	1	46m	
gofabric-service	gofabric:2.4.0	app=gofabric-service,pod-template-				
hash=cb547d7db						
replicaset.apps/gotenant-service-7db6648b97		2	2	1	46m	
gotenant-service	gotenant:2.4.0	app=gotenant-service,pod-template-				
hash=7db6648b97						
replicaset.apps/goraslog-service-6d89f667bd		2	2	2	45m	
goraslog-service	goraslog:2.4.0	app=goraslog-service,pod-template-				
hash=6d89f667bd						
replicaset.apps/govcenter-service-6ff99bf8b8		2	2	1	45m	
vcenter	govcenter:2.4.0	app=govcenter-service,pod-template-				
hash=6ff99bf8b8						
replicaset.apps/gohyperv-service-7f8d87d55		2	2	1	45m	
hyperv	gohyperv:2.4.0	app=gohyperv-service,pod-template-				
hash=7f8d87d55						
NAME	READY	AGE	CONTAINERS	IMAGES		

```
statefulset.apps/rabbitmq 2/2 46m rabbitmq-node rabbitmq:2.4.0
(efa:extreme)extreme@tpvm:~$
```

The following example shows EFA db-status.

```
$ efactl db-status
• mariadb.service - MariaDB 10.4.17 database server
  Loaded: loaded (/lib/systemd/system/mariadb.service; enabled; vendor preset: enabled)
  Drop-In: /etc/systemd/system/mariadb.service.d
           └─migrated-from-my.cnf-settings.conf
  Active: active (running) since Thu 2022-04-28 09:37:31 PDT; 5h 50min ago
  Docs: man:mysql(8)
        https://mariadb.com/kb/en/library/systemd/
  Process: 5355 ExecStartPost=/etc/mysql/debian-start (code=exited, status=0/SUCCESS)
  Process: 5349 ExecStartPost=/bin/sh -c systemctl unset-environment
  _WSREP_START_POSITION (code=exited, status=0/SUCCESS)
  Process: 4730 ExecStartPre=/bin/sh -c [ ! -e /usr/bin/galera_recovery ] && VAR= ||
  VAR=`cd /usr/bin/..; /usr/bin/galera_recovery`; [ $?
  Process: 4719 ExecStartPre=/bin/sh -c systemctl unset-environment _WSREP_START_POSITION
  (code=exited, status=0/SUCCESS)
  Process: 4717 ExecStartPre=/usr/bin/install -m 755 -o mysql -g root -d /var/run/mysql
  (code=exited, status=0/SUCCESS)
  Main PID: 4992 (mysqld)
  Status: "Taking your SQL requests now..."
  Tasks: 36 (limit: 4630)
  CGroup: /system.slice/mariadb.service
          └─4992 /usr/sbin/mysqld --wsrep_start_position=cd3b7a5b-
  c670-11ec-9dbc-2e615cfda7d0:2502

Apr 28 09:37:24 tpvm2 systemd[1]: Starting MariaDB 10.4.17 database server...
Apr 28 09:37:27 tpvm2 mysqld[4730]: WSREP: Recovered position cd3b7a5b-
c670-11ec-9dbc-2e615cfda7d0:2502
Apr 28 09:37:27 tpvm2 mysqld[4992]: 2022-04-28 9:37:27 0 [Note] /usr/sbin/mysqld (mysqld
10.4.17-MariaDB-1:10.4.17+maria~bionic) starting
Apr 28 09:37:28 tpvm2 -wsrep-sst-joiner[5248]: Streaming with xstream
Apr 28 09:37:28 tpvm2 -wsrep-sst-joiner[5293]: Evaluating timeout -k 310 300 socat -u TCP-
LISTEN:4444,reuseaddr stdio | mbstream -x; RC=( $
Apr 28 09:37:31 tpvm2 -wsrep-sst-joiner[5312]: Removing the sst_in_progress file
Apr 28 09:37:31 tpvm2 systemd[1]: Started MariaDB 10.4.17 database server.
```

The following example shows EFA rabbit-status.

```
$ efactl rabbit-status
RABBITMQ_ERLANG_COOKIE env variable support is deprecated and will be REMOVED in a future
version. Use the $HOME/.erlang.cookie file or the --erlang-cookie switch instead.
Cluster status of node rabbit@rabbitmq ...
Basics

Cluster name: rabbit@rabbitmq

Disk Nodes

rabbit@rabbitmq

Running Nodes

rabbit@rabbitmq

Versions

rabbit@rabbitmq: RabbitMQ 3.8.14 on Erlang 23.3.2

Maintenance status

Node: rabbit@rabbitmq, status: not under maintenance
```

```

Alarms

(none)

Network Partitions

(none)

Listeners

Node: rabbit@rabbitmq, interface: [::], port: 15672, protocol: http, purpose: HTTP API
Node: rabbit@rabbitmq, interface: [::], port: 25672, protocol: clustering, purpose: inter-
node and CLI tool communication
Node: rabbit@rabbitmq, interface: [::], port: 5672, protocol: amqp, purpose: AMQP 0-9-1
and AMQP 1.0

Feature flags

Flag: drop_unroutable_metric, state: enabled
Flag: empty_basic_get_metric, state: enabled
Flag: implicit_default_bindings, state: enabled
Flag: maintenance_mode_status, state: enabled
Flag: quorum_queue, state: enabled
Flag: user_limits, state: enabled
Flag: virtual_host_metadata, state: enabled
extreme@tpvm2:~$ efactl db-status
• mariadb.service - MariaDB 10.4.17 database server
  Loaded: loaded (/lib/systemd/system/mariadb.service; enabled; vendor preset: enabled)
  Drop-In: /etc/systemd/system/mariadb.service.d
           └─migrated-from-my.cnf-settings.conf
  Active: active (running) since Thu 2022-04-28 09:37:31 PDT; 5h 43min ago
  Docs: man:mysql(8)
        https://mariadb.com/kb/en/library/systemd/
  Process: 5355 ExecStartPost=/etc/mysql/debian-start (code=exited, status=0/SUCCESS)
  Process: 5349 ExecStartPost=/bin/sh -c systemctl unset-environment
  _WSREP_START_POSITION (code=exited, status=0/SUCCESS)
  Process: 4730 ExecStartPre=/bin/sh -c [ ! -e /usr/bin/galera_recovery ] && VAR= ||
  VAR=`cd /usr/bin/..; /usr/bin/galera_recovery`; [ $?
  Process: 4719 ExecStartPre=/bin/sh -c systemctl unset-environment _WSREP_START_POSITION
  (code=exited, status=0/SUCCESS)
  Process: 4717 ExecStartPre=/usr/bin/install -m 755 -o mysql -g root -d /var/run/mysql
  (code=exited, status=0/SUCCESS)
  Main PID: 4992 (mysqld)
  Status: "Taking your SQL requests now..."
  Tasks: 36 (limit: 4630)
  CGroup: /system.slice/mariadb.service
          └─4992 /usr/sbin/mysqld --wsrep_start_position=cd3b7a5b-
  c670-11ec-9dbc-2e615cfda7d0:2502

Apr 28 09:37:24 tpvm2 systemd[1]: Starting MariaDB 10.4.17 database server...
Apr 28 09:37:27 tpvm2 mysqld[4730]: WSREP: Recovered position cd3b7a5b-
c670-11ec-9dbc-2e615cfda7d0:2502
Apr 28 09:37:27 tpvm2 mysqld[4992]: 2022-04-28 9:37:27 0 [Note] /usr/sbin/mysqld (mysqld
10.4.17-MariaDB-1:10.4.17+maria~bionic) starting
Apr 28 09:37:28 tpvm2 -wsrep-sst-joiner[5248]: Streaming with xstream
Apr 28 09:37:28 tpvm2 -wsrep-sst-joiner[5293]: Evaluating timeout -k 310 300 socat -u TCP-
LISTEN:4444,reuseaddr stdio | mbstream -x; RC=( $
Apr 28 09:37:31 tpvm2 -wsrep-sst-joiner[5312]: Removing the sst_in_progress file
Apr 28 09:37:31 tpvm2 systemd[1]: Started MariaDB 10.4.17 database server.

```



Openstack Controller EFA Commands

[efa-blpair-mapping](#) on page 327

[efa-health show](#) on page 328

[efa-journal clear](#) on page 330

[efa-journal list](#) on page 332

[efa-journal list deps](#) on page 333

[efa-journal reset](#) on page 334

[efa-sync execute](#) on page 335

[openstack network efa-bl-pair-map create](#) on page 337

[openstack network efa-topology-link-map create](#) on page 338

[check syntax openstack network efa-topology-link-map delete](#) on page 339

[openstack network efa-topology-link-map list](#) on page 340

[openstack router create](#) on page 341

This chapter describes commands available for Openstack Controller integrated with EFA Neutron plugin.

efa-blpair-mapping

Adds, removes, or lists border-leaf pairs.

Syntax

```
efa-blpair-mapping add [ --name name | --switch-ips switches ]
```

Adds the specified border leaf pair.

```
efa-blpair-mapping remove [ --name name ]
```

Removes the specified border leaf pair.

```
efa-blpair-mapping list
```

Lists border leaf pairs.

Parameters

--name *name*

Specifies the IP address of the device for which you want to see a list of interfaces.

--switch-ips

Comma-separated IP addresses of border-leaf on fabric. For example, 10.2.2.2,10.2.2.3.

Examples

The following example adds the specified border leaf pair.

```
$ efa-blpair-mapping add --name blp2 --switch-ips 10.2.2.4,10.2.2.3
```

The following example removes the specified border leaf pair.

```
$ efa-blpair-mapping remove --name blp2
```

The following example lists the configured border leaf pairs.

```
$ efa-blpair-mapping list
+-----+-----+-----+
|              ID              | Name |      Switches      |
+-----+-----+-----+
| 8f3d786b-70e3-4c8e-888c-4f95adc162f4 | blp2 | 10.2.2.4,10.2.2.3 |
+-----+-----+-----+
```

efa-health show

Verifies the following: connectivity with EFA, database write operations, OpenStack journal thread status, and the status of the EFA Kubernetes cluster.

Syntax

```
efa-health show [--advanced ]
```

Parameters

--advanced

Provides detailed status about EFA pods and the cluster. Omit this parameter to print basic status information.

Usage Guidelines

In the output, the EFA Health entry represents EFA reachability and EFA database health.

In the output, the EFA High Availability Status entry shows DOWN even if only one node is DOWN. However, even with one node down, EFA Health remains UP because EFA is still operational.

Examples

This example shows basic output for the command.

```
$ sudo efa-health show
EFA Host                : efa.extremenetworks.com
EFA Health              : UP
EFA Health Failure Reason :
EFA High Availability Status : UP
EFA Master              : UP
EFA Slave               : UP

EFA Neutron Journal Size within Threshold : NO
EFA Neutron Journal Size Threshold Setting : 5
EFA Neutron Current Journal Size         : 7
EFA Neutron Journal Maintenance Status   : UP
```

Run `efa-health show` with the advanced option to print EFA POD or cluster status, in detail. The command can be used to infer the pods running status on EFA.

```
$ sudo efa-health show --advanced
EFA Host                : efa.extremenetworks.com
EFA Health              : UP
EFA Health Failure Reason :
EFA High Availability Status : UP
EFA Master              : UP
EFA Slave               : UP

EFA Neutron Journal Size within Threshold : NO
EFA Neutron Journal Size Threshold Setting : 5
EFA Neutron Current Journal Size         : 7
EFA Neutron Journal Maintenance Status   : UP

EFA Cluster Status      : 18/28 Pods Running
```



```
rabbitmq-0 : Running
efa-api-docs-67b8c76ddb-2hg7s : Running
gosnmp-service-546d76b6f9-7wb6t : Running
goopenstack-service-5b687f9f8c-srlxx : Running
gosystem-service-76bbbc6d-jzkng : Running
goauth-service-75c88f4986-vqppm : Running
goraslog-service-7467fb7759-jm28m : Running
gotenant-service-5f8bc9f458-zbz21 : Running
gohyperv-service-b546d647f-ws7t8 : Running
gorbac-service-b546dbdd5-v9dml : Running
gofabric-service-65b6f7d4cf-rfq6f : Running
gotification-service-5fb74fc959-2x91t : Running
goinventory-service-79d6545c69-s66tx : Running
govcenter-service-7c6cb944dd-2d81t : Running
goinventory-service-79d6545c69-gxtn6 : Init:0/2
efa-api-docs-67b8c76ddb-n8zfp : Running
gosnmp-service-546d76b6f9-fv8tz : Init:0/1
goraslog-service-7467fb7759-rvppq : Running
gorbac-service-b546dbdd5-s9qxh : Init:1/2
gosystem-service-76bbbc6d-x8x8g : Running
govcenter-service-7c6cb944dd-rx7gd : Init:0/2
gohyperv-service-b546d647f-zjq97 : Init:0/2
goauth-service-75c88f4986-9fbn7 : Init:0/2
gofabric-service-65b6f7d4cf-q7qns : Init:0/2
gotification-service-5fb74fc959-rr724 : Init:0/2
gotenant-service-5f8bc9f458-n7gm8 : Init:0/3
goopenstack-service-5b687f9f8c-bfssf : Init:0/2
rabbitmq-1 : Running
```

efa-journal clear

Removes all entries or entries specified according to state from the journal.

Syntax

```
efa-journal clear [ pending | failed | completed | processing | all ]
```

Parameters

pending

Clears all entries in Pending state.

failed

Clears all entries in Failed state.

completed

Clears all entries in Completed state.

processing

Clears all entries in Processing state.

all

Clears all entries.

Usage Guidelines

**Important**

Stop the Neutron service before clearing entries using the **all**, **pending**, **failed**, or **processing** parameters. Otherwise, the processing thread can repopulate the entries in the journal.

Examples

The following example clears all entries in the Failed state.

```
$ efa-journal clear failed
```

The following example clears all entries. Note the warning to stop the Neutron service.

```
$ efa-journal clear all
Warning!!! Clearing states pending/processing/all should be executed after shutting
down neutron service. Otherwise the entries can be re-populated by the processing thread.
```

The following example clears all entries in Processing state.

```
$ efa-journal clear processing
Warning!!! Clearing states pending/processing/all should be executed after shutting
down neutron service.
Otherwise the entries can be re-populated by the processing thread.
```

The following example clears all entries in Pending state.

```
$ efa-journal clear pending
Warning!!! Clearing states pending/processing/all should be executed after shutting down
neutron service.
Otherwise the entries can be re-populated by the processing thread.
```

efa-journal list

Lists all journal entries that are present in the extreme journaling database, along with their statuses.

Syntax

```
efa-journal list
```

Examples

This example lists all journal entries and their statuses.

```
$ efa-journal list
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| Id | Object Type |          Object UUID          | Operation | state |
Retry Count |          Reason          |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| 77 | network | 208421ad-483e-4e45-8c6f-d135146alee7 | create | failed |
5 | EFA connection timed out |
| 78 | subnet | 0443d876-eafb-4b8c-9c83-7d32c2cccad6 | create | failed |
0 | Parent Journal ID: 77 failed |
| 79 | port | 9cb6d418-d916-4be8-aad5-1cb07d388d5f | update | failed |
0 | Parent Journal ID: 77 failed |
| 80 | network | b22415ef-afdc-46fa-834f-8506c132da08 | create | failed |
5 | EFA connection timed out |
| 81 | subnet | 2a1e3aea-c996-45f1-94f9-84b604597dc3 | create | failed |
0 | Parent Journal ID: 80 failed |
| 82 | port | 4f3fe359-8de9-45f6-be4b-b9d7c565399f | update | failed |
0 | Parent Journal ID: 80 failed |
| 83 | network | 2237342e-95ec-44c2-92e6-706092bleda3 | create | processing |
0 | |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
```

efa-journal list deps

Debug command: dumps the journal dependency table in user-readable format.

Syntax

efa-journal list deps

Examples

This example dumps the journal dependency table.

```
$ efa-journal list deps
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Parent | Parent-obj |           Parent-uuid           | Parent-op | Parent-state |
Child | Child-obj |           Child-uuid           | Child-op | Child-state |
+-----+-----+-----+-----+-----+-----+
| 77 | network | 208421ad-483e-4e45-8c6f-d135146alee7 | create | failed |
| 78 | subnet | 0443d876-eafb-4b8c-9c83-7d32c2cccad6 | create | failed |
| 77 | network | 208421ad-483e-4e45-8c6f-d135146alee7 | create | failed |
| 79 | port | 9cb6d418-d916-4be8-aad5-1cb07d388d5f | update | failed |
| 78 | subnet | 0443d876-eafb-4b8c-9c83-7d32c2cccad6 | create | failed |
| 79 | port | 9cb6d418-d916-4be8-aad5-1cb07d388d5f | update | failed |
| 80 | network | b22415ef-afdc-46fa-834f-8506c132da08 | create | failed |
| 81 | subnet | 2a1e3aea-c996-45f1-94f9-84b604597dc3 | create | failed |
| 80 | network | b22415ef-afdc-46fa-834f-8506c132da08 | create | failed |
| 82 | port | 4f3fe359-8de9-45f6-be4b-b9d7c565399f | update | failed |
| 81 | subnet | 2a1e3aea-c996-45f1-94f9-84b604597dc3 | create | failed |
| 82 | port | 4f3fe359-8de9-45f6-be4b-b9d7c565399f | update | failed |
| 83 | network | 2237342e-95ec-44c2-92e6-706092b1eda3 | create | failed |
| 84 | subnet | 51c32ef1-c15e-4a1a-ad89-6e2adfdf8da5 | create | failed |
| 83 | network | 2237342e-95ec-44c2-92e6-706092b1eda3 | create | failed |
| 85 | port | aab84828-5e08-4a95-aef6-2498f6ef5991 | update | failed |
| 84 | subnet | 51c32ef1-c15e-4a1a-ad89-6e2adfdf8da5 | create | failed |
| 85 | port | aab84828-5e08-4a95-aef6-2498f6ef5991 | update | failed |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
```

efa-journal reset

Resets the retry count for all or specified failed journal entries and retries the operation.

Syntax

```
efa-journal reset [Journal_ID ]
```

Parameters

Journal_ID

Specifies the journal ID that you want to reset.

Usage Guidelines

You can return journal IDs by running the `efa-journal list` command.

Examples

The following example resets specific failed entries - 77 and 78.

```
$ efa-journal reset 77 78  
5dc0bf8b-674c-4c9f-944e-dc048bd184f0
```

The following example resets all failed entries.

```
$ efa-journal reset
```

efa-sync execute

Syncs EFA with Neutron. Displays summary or detailed information about entities that are out of sync between Neutron and EFA configurations.

Syntax

```
efa-sync execute
```

```
efa-sync check-summary
```

```
efa-sync check-detail
```

Usage Guidelines

The command also logs console output into efa-sync-console.log.

The efa-sync tool uses Neutron APIs and keystone authentication. The authentication parameters are selected from the neutron.conf file. Before running efa-sync, make sure that all the parameters under the [keystone-authtoken] section are set to the correct values.



Note

In case of default OpenStack settings in neutron.conf, the following parameters need to be edited to 'default' -- lowercase 'd' under [keystone_authtoken]:

```
project_domain_name = default
user_domain_name = default
```

Journal entries that are completed or failed are cleared from the journal list if the execution of efa-sync is successful. In case of failure, the journal entries are not cleared.

Examples

This example syncs EFA with Neutron.

```
$ efa-sync execute
Starting Sync
Syncing Networks..
  Add Network to EFA Id: 9e63bc57-568f-488e-9214-21db3fd3cd12
  (Succeeded)
Syncing Ports..
  Add Port to EFA Id:
  9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.11-port-channel-lag_1
  (Succeeded)
  Add Port to EFA Id:
  9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.46-port-channel-lag_1
  (Succeeded)
Completed Sync
```

This example displays summary information about entities that are out of sync between Neutron and EFA configurations.

```
$ efa-sync check-summary
05-November-2020 11:04:15 : Starting Check
Summary:
  Neutron Networks to be added to EFA : 1
```

```
Neutron Networks to be deleted from EFA : 0
Neutron Ports to be added to EFA : 2
Neutron Ports to be deleted from EFA : 0
Total # of resources to be resynced : 3
05-November-2020 11:04:16 : Completed Check
```

This example displays detailed information about entities that are out of sync between Neutron and EFA configurations.

```
$ efa-sync check-detail
05-November-2020 11:04:22 : Starting Check
  Neutron Networks to be added to EFA :
    Id: 9e63bc57-568f-488e-9214-21db3fd3cd12
  Neutron Ports to be added from EFA :
    Id: 9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.11-port-channel-
lag_1
    Id: 9e63bc57-568f-488e-9214-21db3fd3cd12-10.25.225.46-port-channel-lag_1
05-November-2020 11:04:22 : Completed Check
```


openstack network efa-bl-pair-map create

Maps a border-leaf pair and specifies the border-leaf pair on which the route will be installed.

Syntax

```
openstack network efa-bl-pair-map create [ --name name | --switch-ips  
      switches]
```

Parameters

--name *name*

Specifies the name of the border-leaf pair.

--switch-ips *switches*

Comma-separated IP addresses of border-leaf on fabric. For example, 10.2.2.2,10.2.2.3.

Examples

```
openstack network efa-bl-pair-map create --name blp1 --switch-ips 10.2.2.2,10.2.2.3  
+-----+  
| Field      | Value                                     |  
+-----+  
| id         | 0d19abc5-d31b-4682-b86d-cc1617ea0c97 |  
| name       | blp1                                     |  
| switch_ips | 10.2.2.2,10.2.2.3                       |  
+-----+
```

openstack network efa-topology-link-map create

Creates a topology link.

Syntax

```
openstack network efa-topology-link-map create [--host host | --nic NIC |  
--provider-network provider-network | --port port | --po-name po-  
name ]
```

Parameters

--host *host*

Host name of the compute node that connects to the switch.

--nic *NIC*

Physical NIC on the compute host which connects to the switch.

--provider-network *provider-network*

Provider network name on compute host. For example, physnet1 (default).

--port *port*

Switch port to which the physical NIC on the compute host is connected.

--po-name *po-name*

Switch PO name to which the compute host NICs are connected.

Examples

```
# openstack network efa-topology-link-map create --host  
Openstack115 --nic eth1 --pn default --switch 10.24.14.133 --port 0/1 --po-name  
lag_1
```

openstack network efa-topology-link-map delete

Deletes a topology link.

Syntax

```
openstack network efa-topology-link-map delete <uuid-link1> [<uuid-link2>
... ]
```

Specify the UUIDs of the links to delete.

Examples

```
openstack network efa-topology-link-map delete
08bb90b6-ac37-4b7f-aa80-d1d08de7e54b
012b90b6-ac12-427f-a220-d1d08de7e123
```

openstack network efa-topology-link-map list

Lists topology links.

Syntax

```
openstack network efa-topology-link-map list
```

Examples

The following example displays the created topology links on the controller.

```
# openstack network efa-topology-link-map list
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| ID                                     | Host           | Nic | Provider Network |
Switch      | Port | Po Name |                                     |
+-----+-----+-----+-----+-----+-----+
| 08fa6b52-8c25-4193-8e4b-00ce3d81f392 | DCGW1          |     | external_nw      |
10.20.246.8 | 0/14 | lag_4   |                                     |
| 0977490d-f07a-4b2c-b683-8359bde19fcc | niyer-devstack | eth3 | physnet3         |
10.20.246.8 | 0/15 |         |                                     |
| 14486fd7-1bd7-4ba0-a159-b2e76f775fdd | niyer-devstack | eth4 | default          |
10.20.246.8 | 0/16 |         |                                     |
| 4ce69af9-03ce-49d7-a0dc-0b5c64e25ded | niyer-devstack | eth1 | physnet11       |
10.20.246.8 | 0/13 | lag_1   |                                     |
| c3879465-ecdd-492b-918e-981c104005ba | DCGW1          |     | external_nw      |
10.20.246.7 | 0/14 | lag_4   |                                     |
| c9663e8d-f91c-44ff-b4a1-fb97e603bda2 | niyer-devstack | eth0 | physnet11       |
10.20.246.7 | 0/13 | lag_1   |                                     |
| daf3c58a-9b24-4069-afe8-94d2a20004c9 | niyer-devstack | eth2 | physnet2         |
10.20.246.7 | 0/15 |         |                                     |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
```

openstack router create

Creates a router in centralized mode.

Syntax

```
openstack router create [ router-name --centralized ]
```

Parameters

router-name **--centralized**

Creates a router in centralized mode.

Usage Guidelines

In centralized mode, routing is configured only on the border-leaf pairs.

When you create a router, the default mode is centralized. You can create a router with or without the `centralized` option and achieve the same result.

This example creates one centralized router called R2.

```
$ openstack router create R2  
$ openstack router create R2 --centralized
```