



ExtremeCloud™ Orchestrator v3.7.0 Command References

Comprehensive CLI Commands and Syntax

9039073-00 Rev AA
December 2024



Copyright © 2024 Extreme Networks, Inc. All rights reserved.



Table of Contents

What's New in this Document.....	9
New Commands.....	9
Modified Commands.....	9
Updated Command Output.....	9
Using the XCO Application CLI.....	10
User Accounts.....	10
Default Account Credentials.....	11
Log in to XCO.....	11
Completing CLI Commands.....	11
CLI Keyboard Shortcuts.....	12
Unsupported Input Characters.....	12
ExtremeCloud Orchestrator Commands.....	13
efa auth authentication preference.....	20
efa auth execution.....	22
efa auth ldapconfig.....	23
efa auth ldapconfig reset.....	27
efa auth rolemapping.....	29
efa auth settings token.....	31
efa auth tacacsconfig.....	32
efa auth user.....	32
efa certificate device expiry show.....	34
efa certificate device install.....	35
efa certificate expiry show.....	37
efa certificate server.....	37
efa certificate server renew.....	38
efa fabric clone.....	38
efa fabric configure.....	38
efa fabric create.....	39
efa fabric debug clear-config.....	39
efa fabric debug config-gen-reason.....	40
efa fabric debug device drift.....	40
efa fabric debug service lock.....	40
efa fabric debug service smartdrc.....	40
efa fabric debug set.....	41
efa fabric delete.....	41
efa fabric device add.....	42
efa fabric device add-bulk.....	43
efa fabric device remove.....	44
efa fabric error show.....	45
efa fabric execution.....	45


efa fabric health show.....	46
efa fabric migrate.....	50
efa fabric setting.....	51
efa fabric show.....	56
efa fabric show-config.....	56
efa fabric show summary.....	57
efa fabric topology show overlay.....	57
efa fabric topology show underlay.....	57
efa fabric topology show physical.....	58
efa health debug clear.....	58
efa health detail show.....	59
efa health inventory show.....	59
efa health show.....	59
efa inventory admin-state.....	60
efa inventory config-backup.....	61
efa inventory config-replay.....	62
efa inventory debug device-adapter-status.....	63
efa inventory debug devices-lock.....	64
efa inventory debug devices-unlock.....	64
efa inventory debug unblock-from-fwdl.....	65
efa inventory device clear route-all.....	66
efa inventory device compare.....	66
efa inventory device delete.....	67
efa inventory device discovery-time list.....	67
efa inventory device discovery-time update.....	67
efa inventory device execute-cli.....	68
efa inventory device firmware-download activate.....	69
efa inventory device firmware-download commit.....	69
efa inventory device firmware-download execute.....	69
efa inventory device firmware-download prepare add.....	73
efa inventory device firmware-download prepare list.....	78
efa inventory device firmware-download prepare remove.....	79
efa inventory device firmware-download restore.....	81
efa inventory device firmware-download show.....	82
efa inventory device health status.....	84
efa inventory device interface list.....	85
efa inventory device interface list-breakout.....	85
efa inventory device interface redundant-management.....	86
efa inventory device interface set-admin-state.....	87
efa inventory device interface set-breakout.....	89
efa inventory device interface set-description.....	91
efa inventory device interface set-fec.....	93
efa inventory device interface set-link-error-disable.....	94
efa inventory device interface set-mtu.....	94
efa inventory device interface unset-mtu.....	97
efa inventory device interface set-speed.....	98
efa inventory device interface unset-breakout.....	99
efa inventory device interface unset-description.....	101
efa inventory device interface unset-fec.....	102

efa inventory device interface unset-link-error-disable.....	102
efa inventory device list.....	103
efa inventory device lldp list.....	104
efa inventory device ntp disable-server.....	105
efa inventory device ntp server create.....	105
efa inventory device ntp server delete.....	106
efa inventory device ntp server list.....	106
efa inventory device running-config persist.....	107
efa inventory device register.....	108
efa inventory device reload.....	108
efa inventory device secure settings.....	109
efa inventory device secure settings apply.....	109
efa inventory device secure settings reset.....	110
efa inventory device secure settings show.....	110
efa inventory device secure settings update.....	111
efa inventory device setting show.....	112
efa inventory device setting update.....	113
efa inventory device snmp community create.....	119
efa inventory device snmp community delete.....	120
efa inventory device snmp community list.....	120
efa inventory device snmp group list.....	121
efa inventory device snmp group update.....	121
efa inventory device snmp host create.....	122
efa inventory device snmp host delete.....	123
efa inventory device snmp host list.....	124
efa inventory device snmp user create.....	125
efa inventory device snmp user delete.....	126
efa inventory device snmp user list.....	126
efa inventory device snmp use-vrf create.....	126
efa inventory device snmp use-vrf delete.....	127
efa inventory device snmp use-vrf list.....	127
efa inventory device snmp view.....	128
efa inventory device timezone debug-show.....	129
efa inventory device timezone list.....	129
efa inventory device timezone set.....	129
efa inventory device timezone unset.....	130
efa inventory device threshold-monitor list.....	130
efa inventory device threshold-monitor set.....	131
efa inventory device threshold-monitor unset.....	133
efa inventory device tpvm list.....	133
efa inventory device tpvm-upgrade execute.....	134
efa inventory device tpvm-upgrade show.....	135
efa inventory device update.....	136
efa inventory drift-reconcile.....	136
efa inventory execution.....	137
efa inventory firmware-host delete.....	138
efa inventory firmware-host list.....	138
efa inventory firmware-host register.....	138
efa inventory firmware-host update.....	138

efa inventory kvstore.....	139
efa inventory rma.....	139
efa license add.....	140
efa license delete.....	140
efa license show.....	141
efa login.....	141
efa logout.....	141
efa mgmt route create.....	141
efa mgmt route delete.....	142
efa mgmt route show.....	142
efa mgmt subinterface create.....	143
efa mgmt subinterface delete.....	143
efa mgmt subinterface show.....	144
efa mgmt subinterface staticip add.....	144
efa mgmt subinterface staticip remove.....	145
efa mgmt subinterface staticip show.....	145
efa notification subscribers add-https.....	146
efa notification subscribers add-syslog-relp.....	149
efa notification subscribers delete.....	153
efa notification subscribers get.....	153
efa notification subscribers list.....	154
efa policy community-list create.....	154
efa policy community-list update.....	154
efa policy community-list delete.....	157
efa policy community-list list.....	158
efa policy extcommunity-list.....	159
efa policy extcommunity-list create.....	163
efa policy extcommunity-list update.....	164
efa policy extcommunity-list delete.....	166
efa policy extcommunity-list list.....	167
efa policy large-community-list create.....	168
efa policy large-community-list update.....	170
efa policy large-community-list delete.....	172
efa policy large-community-list list.....	172
efa policy prefix-list create.....	173
efa policy prefix-list update.....	174
efa policy prefix-list delete.....	176
efa policy prefix-list list.....	177
efa policy qos map create.....	178
efa policy qos map update.....	179
efa policy qos map delete.....	181
efa policy qos map list.....	181
efa policy qos profile create.....	183
efa policy qos profile update.....	184
efa policy qos profile delete.....	185
efa policy qos profile bind.....	185
efa policy qos profile unbind.....	186
efa policy qos profile list.....	187
efa policy qos service-policy-map create.....	188

efa policy qos service-policy-map delete.....	188
efa policy qos service-policy-map list.....	189
efa policy qos service-policy-map update.....	189
efa policy route-map create.....	190
efa policy route-map update.....	191
efa policy route-map delete.....	192
efa policy route-map list.....	193
efa policy route-map-match create.....	193
efa policy route-map-match delete.....	195
efa policy route-map-set create.....	196
efa policy route-map-set delete.....	197
efa rbac execution.....	198
efa rbac role show.....	199
efa scvmm delete.....	200
efa scvmm links physical.....	200
efa scvmm links virtual.....	200
efa scvmm list.....	200
efa scvmm register.....	201
efa scvmm settings show.....	201
efa scvmm settings update.....	201
efa scvmm update.....	201
efa snmp subscriber.....	202
efa show-running-config.....	203
efa status.....	204
efa system alarm acknowledge.....	205
efa system alarm close.....	205
efa system alarm inventory show.....	206
efa system alarm purge.....	207
efa system alarm show.....	207
efa system alarm summary.....	210
efa system alert.....	210
efa system backup.....	211
efa system backup-list.....	212
efa system cleanup.....	213
efa system feature show.....	213
efa system feature update.....	213
efa system logging set.....	214
efa system logging unset.....	215
efa system logging show.....	216
efa system restore.....	217
efa system service enable.....	217
efa system service disable.....	218
efa system settings reset.....	218
efa system settings show.....	218
efa system settings update.....	219
efa system supportsave.....	222
efa system supportsave-list.....	222
efa tenant create.....	223
efa tenant debug.....	224

efa tenant delete.....	226
efa tenant epg configure.....	226
efa tenant epg create.....	227
efa tenant epg delete.....	233
efa tenant epg detach.....	234
efa tenant epg error show.....	236
efa tenant epg show.....	237
efa tenant epg update.....	238
efa tenant execution.....	243
efa tenant po configure.....	244
efa tenant po create.....	244
efa tenant po delete.....	245
efa tenant po show.....	246
efa tenant po update.....	248
efa tenant service bgp peer configure.....	249
efa tenant service bgp peer create.....	249
efa tenant service bgp peer delete.....	258
efa tenant service bgp peer operational show.....	258
efa tenant service bgp peer show.....	264
efa tenant service bgp peer update.....	267
efa tenant service bgp peer-group configure.....	273
efa tenant service bgp peer-group create.....	274
efa tenant service bgp peer-group delete.....	279
efa tenant service bgp peer-group show.....	279
efa tenant service bgp peer-group update.....	281
efa tenant service mirror session create.....	286
efa tenant service mirror session delete.....	287
efa tenant service mirror session configure.....	288
efa tenant service mirror session show.....	288
efa tenant show.....	288
efa tenant update.....	288
efa tenant vrf create.....	290
efa tenant vrf delete.....	297
efa tenant vrf error show.....	298
efa tenant vrf show.....	298
efa tenant vrf update.....	301
efa vcenter debug.....	308
efa vcenter delete.....	311
efa vcenter links.....	312
efa vcenter list.....	313
efa vcenter register.....	315
efa vcenter update.....	316
efa version.....	317
efa-change-hostname.....	317
efa-change-ip.....	318
efa-change-vip.....	318
efactl.....	320



What's New in this Document

New Commands

[efa inventory device interface set-description](#) on page 91

[efa inventory device interface unset-description](#) on page 101

Modified Commands

[efa system settings reset](#) on page 218

[efa system settings update](#) on page 219

Updated Command Output

[efa system settings show](#) on page 218

[efa version](#) on page 317



Using the XCO Application CLI

- [User Accounts](#) on page 10
- [Default Account Credentials](#) on page 11
- [Log in to XCO](#) on page 11
- [Completing CLI Commands](#) on page 11
- [CLI Keyboard Shortcuts](#) on page 12
- [Unsupported Input Characters](#) on page 12



Note

From release 3.2.0 onwards, Extreme Fabric Automation (EFA) is referred to as ExtremeCloud Orchestrator (XCO). The terms EFA and XCO refer to the same product and are used interchangeably.

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

The following topics describe accessing and using the XCO 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.

XCO 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 [ExtremeCloud Orchestrator Configuration Guide, 3.7.0](#).

Default Account Credentials

As a best practice, log on as the administrator and change the default passwords immediately after the XCO 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 XCO

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

Procedure

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

```
$ 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 XCO, 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.
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.



ExtremeCloud Orchestrator Commands

[efa auth authentication preference](#) on page 20
[efa auth execution](#) on page 22
[efa auth ldapconfig](#) on page 23
[efa auth ldapconfig reset](#) on page 27
[efa auth rolemapping](#) on page 29
[efa auth settings token](#) on page 31
[efa auth tacacsconfig](#) on page 32
[efa auth user](#) on page 32
[efa certificate device expiry show](#) on page 34
[efa certificate device install](#) on page 35
[efa certificate expiry show](#) on page 37
[efa certificate server](#) on page 37
[efa certificate server renew](#) on page 38
[efa fabric clone](#) on page 38
[efa fabric configure](#) on page 38
[efa fabric create](#) on page 39
[efa fabric debug clear-config](#) on page 39
[efa fabric debug config-gen-reason](#) on page 40
[efa fabric debug device drift](#) on page 40
[efa fabric debug service lock](#) on page 40
[efa fabric debug service smartdrc](#) on page 40
[efa fabric debug set](#) on page 41
[efa fabric delete](#) on page 41
[efa fabric device add](#) on page 42
[efa fabric device add-bulk](#) on page 43
[efa fabric device remove](#) on page 44
[efa fabric error show](#) on page 45
[efa fabric execution](#) on page 45
[efa fabric health show](#) on page 46
[efa fabric migrate](#) on page 50
[efa fabric setting](#) on page 51
[efa fabric show](#) on page 56
[efa fabric show-config](#) on page 56
[efa fabric show summary](#) on page 57

[efa fabric topology show overlay](#) on page 57
[efa fabric topology show underlay](#) on page 57
[efa fabric topology show physical](#) on page 58
[efa health debug clear](#) on page 58
[efa health detail show](#) on page 59
[efa health inventory show](#) on page 59
[efa health show](#) on page 59
[efa inventory admin-state](#) on page 60
[efa inventory config-backup](#) on page 61
[efa inventory config-replay](#) on page 62
[efa inventory debug device-adapter-status](#) on page 63
[efa inventory debug devices-lock](#) on page 64
[efa inventory debug devices-unlock](#) on page 64
[efa inventory debug unblock-from-fwdl](#) on page 65
[efa inventory device clear route-all](#) on page 66
[efa inventory device compare](#) on page 66
[efa inventory device delete](#) on page 67
[efa inventory device discovery-time list](#) on page 67
[efa inventory device discovery-time update](#) on page 67
[efa inventory device execute-cli](#) on page 68
[efa inventory device firmware-download activate](#) on page 69
[efa inventory device firmware-download commit](#) on page 69
[efa inventory device firmware-download execute](#) on page 69
[efa inventory device firmware-download prepare add](#) on page 73
[efa inventory device firmware-download prepare list](#) on page 78
[efa inventory device firmware-download prepare remove](#) on page 79
[efa inventory device firmware-download restore](#) on page 81
[efa inventory device firmware-download show](#) on page 82
[efa inventory device health status](#) on page 84
[efa inventory device interface list](#) on page 85
[efa inventory device interface list-breakout](#) on page 85
[efa inventory device interface redundant-management](#) on page 86
[efa inventory device interface set-admin-state](#) on page 87
[efa inventory device interface set-breakout](#) on page 89
[efa inventory device interface set-description](#) on page 91
[efa inventory device interface set-fec](#) on page 93
[efa inventory device interface set-link-error-disable](#) on page 94
[efa inventory device interface set-mtu](#) on page 94
[efa inventory device interface unset-mtu](#) on page 97
[efa inventory device interface set-speed](#) on page 98
[efa inventory device interface unset-breakout](#) on page 99
[efa inventory device interface unset-description](#) on page 101

[efa inventory device interface unset-fec](#) on page 102
[efa inventory device interface unset-link-error-disable](#) on page 102
[efa inventory device list](#) on page 103
[efa inventory device lldp list](#) on page 104
[efa inventory device ntp disable-server](#) on page 105
[efa inventory device ntp server create](#) on page 105
[efa inventory device ntp server delete](#) on page 106
[efa inventory device ntp server list](#) on page 106
[efa inventory device running-config persist](#) on page 107
[efa inventory device register](#) on page 108
[efa inventory device reload](#) on page 108
[efa inventory device secure settings](#) on page 109
[efa inventory device secure settings apply](#) on page 109
[efa inventory device secure settings reset](#) on page 110
[efa inventory device secure settings show](#) on page 110
[efa inventory device secure settings update](#) on page 111
[efa inventory device setting show](#) on page 112
[efa inventory device setting update](#) on page 113
[efa inventory device snmp community create](#) on page 119
[efa inventory device snmp community delete](#) on page 120
[efa inventory device snmp community list](#) on page 120
[efa inventory device snmp group list](#) on page 121
[efa inventory device snmp group update](#) on page 121
[efa inventory device snmp host create](#) on page 122
[efa inventory device snmp host delete](#) on page 123
[efa inventory device snmp host list](#) on page 124
[efa inventory device snmp user create](#) on page 125
[efa inventory device snmp user delete](#) on page 126
[efa inventory device snmp user list](#) on page 126
[efa inventory device snmp use-vrf create](#) on page 126
[efa inventory device snmp use-vrf delete](#) on page 127
[efa inventory device snmp use-vrf list](#) on page 127
[efa inventory device snmp view](#) on page 128
[efa inventory device timezone debug-show](#) on page 129
[efa inventory device timezone list](#) on page 129
[efa inventory device timezone set](#) on page 129
[efa inventory device timezone unset](#) on page 130
[efa inventory device threshold-monitor list](#) on page 130
[efa inventory device threshold-monitor set](#) on page 131
[efa inventory device threshold-monitor unset](#) on page 133
[efa inventory device tpvm list](#) on page 133
[efa inventory device tpvm-upgrade execute](#) on page 134

[efa inventory device tpvm-upgrade show](#) on page 135
[efa inventory device update](#) on page 136
[efa inventory drift-reconcile](#) on page 136
[efa inventory execution](#) on page 137
[efa inventory firmware-host delete](#) on page 138
[efa inventory firmware-host list](#) on page 138
[efa inventory firmware-host register](#) on page 138
[efa inventory firmware-host update](#) on page 138
[efa inventory kvstore](#) on page 139
[efa inventory rma](#) on page 139
[efa license add](#) on page 140
[efa license delete](#) on page 140
[efa license show](#) on page 141
[efa login](#) on page 141
[efa logout](#) on page 141
[efa mgmt route create](#) on page 141
[efa mgmt route delete](#) on page 142
[efa mgmt route show](#) on page 142
[efa mgmt subinterface create](#) on page 143
[efa mgmt subinterface delete](#) on page 143
[efa mgmt subinterface show](#) on page 144
[efa mgmt subinterface staticip add](#) on page 144
[efa mgmt subinterface staticip remove](#) on page 145
[efa mgmt subinterface staticip show](#) on page 145
[efa notification subscribers add-https](#) on page 146
[efa notification subscribers add-syslog-relp](#) on page 149
[efa notification subscribers delete](#) on page 153
[efa notification subscribers get](#) on page 153
[efa notification subscribers list](#) on page 154
[efa policy community-list create](#) on page 154
[efa policy community-list update](#) on page 154
[efa policy community-list delete](#) on page 157
[efa policy community-list list](#) on page 158
[efa policy extcommunity-list](#) on page 159
[efa policy extcommunity-list create](#) on page 163
[efa policy extcommunity-list update](#) on page 164
[efa policy extcommunity-list delete](#) on page 166
[efa policy extcommunity-list list](#) on page 167
[efa policy large-community-list create](#) on page 168
[efa policy large-community-list update](#) on page 170
[efa policy large-community-list delete](#) on page 172
[efa policy large-community-list list](#) on page 172

[efa policy prefix-list create](#) on page 173
[efa policy prefix-list update](#) on page 174
[efa policy prefix-list delete](#) on page 176
[efa policy prefix-list list](#) on page 177
[efa policy qos map create](#) on page 178
[efa policy qos map update](#) on page 179
[efa policy qos map delete](#) on page 181
[efa policy qos map list](#) on page 181
[efa policy qos profile create](#) on page 183
[efa policy qos profile update](#) on page 184
[efa policy qos profile delete](#) on page 185
[efa policy qos profile bind](#) on page 185
[efa policy qos profile unbind](#) on page 186
[efa policy qos profile list](#) on page 187
[efa policy qos service-policy-map create](#) on page 188
[efa policy qos service-policy-map delete](#) on page 188
[efa policy qos service-policy-map list](#) on page 189
[efa policy qos service-policy-map update](#) on page 189
[efa policy route-map create](#) on page 190
[efa policy route-map update](#) on page 191
[efa policy route-map delete](#) on page 192
[efa policy route-map list](#) on page 193
[efa policy route-map-match create](#) on page 193
[efa policy route-map-match delete](#) on page 195
[efa policy route-map-set create](#) on page 196
[efa policy route-map-set delete](#) on page 197
[efa rbac execution](#) on page 198
[efa rbac role show](#) on page 199
[efa scvmm delete](#) on page 200
[efa scvmm links physical](#) on page 200
[efa scvmm links virtual](#) on page 200
[efa scvmm list](#) on page 200
[efa scvmm register](#) on page 201
[efa scvmm settings show](#) on page 201
[efa scvmm settings update](#) on page 201
[efa scvmm update](#) on page 201
[efa snmp subscriber](#) on page 202
[efa show-running-config](#) on page 203
[efa status](#) on page 204
[efa system alarm acknowledge](#) on page 205
[efa system alarm close](#) on page 205
[efa system alarm inventory show](#) on page 206

[efa system alarm purge](#) on page 207
[efa system alarm show](#) on page 207
[efa system alarm summary](#) on page 210
[efa system alert](#) on page 210
[efa system backup](#) on page 211
[efa system backup-list](#) on page 212
[efa system cleanup](#) on page 213
[efa system feature show](#) on page 213
[efa system feature update](#) on page 213
[efa system logging set](#) on page 214
[efa system logging unset](#) on page 215
[efa system logging show](#) on page 216
[efa system restore](#) on page 217
[efa system service enable](#) on page 217
[efa system service disable](#) on page 218
[efa system settings reset](#) on page 218
[efa system settings show](#) on page 218
[efa system settings update](#) on page 219
[efa system supportsave](#) on page 222
[efa system supportsave-list](#) on page 222
[efa tenant create](#) on page 223
[efa tenant debug](#) on page 224
[efa tenant delete](#) on page 226
[efa tenant epg configure](#) on page 226
[efa tenant epg create](#) on page 227
[efa tenant epg delete](#) on page 233
[efa tenant epg detach](#) on page 234
[efa tenant epg error show](#) on page 236
[efa tenant epg show](#) on page 237
[efa tenant epg update](#) on page 238
[efa tenant execution](#) on page 243
[efa tenant po configure](#) on page 244
[efa tenant po create](#) on page 244
[efa tenant po delete](#) on page 245
[efa tenant po show](#) on page 246
[efa tenant po update](#) on page 248
[efa tenant service bgp peer configure](#) on page 249
[efa tenant service bgp peer create](#) on page 249
[efa tenant service bgp peer delete](#) on page 258
[efa tenant service bgp peer operational show](#) on page 258
[efa tenant service bgp peer show](#) on page 264
[efa tenant service bgp peer update](#) on page 267

[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 279
[efa tenant service bgp peer-group show](#) on page 279
[efa tenant service bgp peer-group update](#) on page 281
[efa tenant service mirror session create](#) on page 286
[efa tenant service mirror session delete](#) on page 287
[efa tenant service mirror session configure](#) on page 288
[efa tenant service mirror session show](#) on page 288
[efa tenant show](#) on page 288
[efa tenant update](#) on page 288
[efa tenant vrf create](#) on page 290
[efa tenant vrf delete](#) on page 297
[efa tenant vrf error show](#) on page 298
[efa tenant vrf show](#) on page 298
[efa tenant vrf update](#) on page 301
[efa vcenter debug](#) on page 308
[efa vcenter delete](#) on page 311
[efa vcenter links](#) on page 312
[efa vcenter list](#) on page 313
[efa vcenter register](#) on page 315
[efa vcenter update](#) on page 316
[efa version](#) on page 317
[efa-change-hostname](#) on page 317
[efa-change-ip](#) on page 318
[efa-change-vip](#) on page 318
[efactl](#) on page 320

This document describes commands available in the ExtremeCloud Orchestrator CLI.

efa auth authentication preference

Configure authentication preference.

```
efa auth authentication preference add [--authType {TACACS | LDAP | LOCAL
| HOST} --identifier string --preference {1 | 2 | 3 | 4 | 5} ]
```

```
efa auth authentication preference update [--authType {TACACS | LDAP |
LOCAL | HOST} --identifier string --preference {1 | 2 | 3 | 4 | 5} ]
```

```
efa auth authentication preference delete [--authType {TACACS | LDAP |
LOCAL | HOST} --identifier string --preference {1 | 2 | 3 | 4 | 5} ]
```

```
efa auth authentication preference show
```

```
--authType {TACACS | LDAP | LOCAL | HOST}
```

Specifies the authentication type.

```
--identifier string
```

Specifies the authentication configuration. For TACACS+ it is the host name, and for LDAP it is the name of the configuration. For authType of LOCAL or HOST it is the same as authType.

```
--preference {1 | 2 | 3 | 4 | 5}
```

Specifies the authentication order. Values 1- 5 must be unique, and specify the order used to authenticate the user.

This example adds the specified authentication preference.

```
efa auth authentication preference add --authType=LOCAL --identifier=LOCAL --preference=3
Successfully updated the auth preference.
+-----+-----+-----+
| Auth Type | Identifier | Preference |
+-----+-----+-----+
| LOCAL     | LOCAL     | 3          |
+-----+-----+-----+
```

This example lists the authentication preferences.

```
efa auth authentication preference show
+-----+-----+-----+
| Auth Type | Identifier | Preference |
+-----+-----+-----+
| HOST      | HOST      | 1          |
+-----+-----+-----+
| LOCAL     | LOCAL     | 3          |
+-----+-----+-----+
```

This example updates an authentication preference.

```
efa auth authentication preference update --authType=LOCAL --identifier=LOCAL --
preference=2
```

Successfully updated the auth preference.

```
+-----+-----+-----+
| Auth Type | Identifier | Preference |
+-----+-----+-----+
| LOCAL     | LOCAL     | 2          |
+-----+-----+-----+
```

This example deletes an authentication preference.

```
efa auth authentication preference delete --authType=LOCAL --identifier=LOCAL
```

This example adds the specified authentication preference.

```
efa auth authentication preference add --authType=TACACS --identifier=10.37.135.12 --
preference=2
```

Successfully added the auth preference.

```
+-----+-----+-----+
| Auth Type | Identifier   | Preference |
+-----+-----+-----+
| TACACS    | 10.37.135.12 | 2          |
+-----+-----+-----+
```

This example adds the specified authentication preference.

```
efa auth authentication preference add --authType=LDAP --identifier=kvm12.com --
preference=3
```

Successfully added the auth preference.

```
+-----+-----+-----+
| Auth Type | Identifier   | Preference |
+-----+-----+-----+
| LDAP      | kvm12.com   | 3          |
+-----+-----+-----+
```

This example shows an authentication preference.

```
efa auth authentication preference show
```

```
+-----+-----+-----+
| Auth Type | Identifier   | Preference |
+-----+-----+-----+
| HOST      | HOST        | 1          |
+-----+-----+-----+
| TACACS    | 10.37.135.12 | 2          |
+-----+-----+-----+
| LDAP      | kvm12.com   | 3          |
+-----+-----+-----+
```

efa auth execution

Deletes or displays the authorization execution logs.

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.

--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.

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.

```
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 pword | --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 pword | --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 ]
```

```
efa auth ldapconfig show [--name ldap-name ]
```

```
--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 pword
```

Specifies the password of the bind user.

```
--cacert cert-loc
```

Specifies the location of the Certificate Authority certificate.

```
--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`.

--host *hostname*

Specifies the host name, IPv4, or IPv6 address of the LDAP server.

--name *ldap-name*

Specifies the name of the LDAP connection.

--primary *value*

Specifies 1 when multiple LDAP connections are available.

--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.

--timeout *value*

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

--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.

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

When a user is assigned XCO 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 XCO on a TPVM, see the "TPVM Management" section of the *Extreme SLX-OS Management Configuration Guide*.

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 XCO 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 memberId --group-object-class posixGroup
```

This example maps an LDAP group to an XCO 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 ldapconfig reset

Reset LDAP configuration.



Note

`--name` and at least one other LDAP attribute to reset are required.

```
efa auth ldapconfig reset [ --name ldap-name | --cacert string | --bind-user-name string | --bind-user-password string | --user-search-base string | --user-object-class string | --user-login-attribute string | --user-role-attribute string | --user-role-attribute-key string | --user-member-attribute string | --group-search-base string | --group-object-class string | --group-attribute string | --group-member-user-attribute string | --group-member-mapping-attribute string ]
```

--bind-user-name *string*

Specifies the bind-user-name to reset.

--bind-user-password *string*

Specifies the bind-user-password to reset.

--cacert *string*

Specifies the cacert to reset.

--group-attribute *string*

Specifies the group-attribute to reset.

--group-member-mapping-attribute *string*

Specifies the group-member-mapping-attribute to reset.

--group-member-user-attribute *string*

Specifies the group-member-user-attribute to reset.

--group-object-class *string*

Specifies the group-object-name to reset.

--group-search-base *string*

Specifies the group-search-base to reset.

--name *string*

Specifies the name of the LDAP configuration to reset.

--user-login-attribute *string*

Specifies the user-login-attribute to reset.

--user-member-attribute *string*

Specifies the user-member-attribute to reset.

--user-object-class *string*

Specifies the user-object-class to reset.

--user-role-attribute *string*

Specifies the user-role-attribute to reset.

--user-role-attribute-key *string*

Specifies the user-role-attribute-key to reset.

--user-search-base *string*

Specifies the user-search-base to reset.

```
efa auth ldapconfig reset --name kvm12.com --group-attribute --group-member-mapping-attribute
Reset LDAP configuration is successful
efa auth ldapconfig reset --name kvm12.com --user-member-attribute
Reset LDAP configuration is successful
```

efa auth rolemapping

Assigns XCO roles to a user or an LDAP group.

```
efa auth rolemapping add [--name user-name | --role efa-role | --type user-type | --auth-type string | --auth-identifier string ]
```

```
efa auth rolemapping show [ --auth-type string | --auth-identifier string ]
```

```
efa auth rolemapping remove --id id
```

```
--auth-identifier string
```

Mandatory for LDAP and TACACS role mapping. It represents the target configuration to which the roles are being added to. For LOCAL and HOST auth types it is empty.

```
--auth-type string
```

Possible values are HOST, LOCAL, LDAP or TACACS. Default cli value is HOST.

```
--id id
```

ID of the role mapping.

```
--name user-name
```

Specifies the user name for HOST and LOCAL configuration. For LDAP it can be a group or user name. For TACACS it represents the group name.

```
--role efa-role
```

Specifies the name of 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 rolemapping. Enter either *user* or *group*. For LDAP it has value "group" by default. For local and host role mapping it has value "user" if not input. For TACACS the type value is "group".

These examples assign the role of SystemAdmin.

```
# efa auth rolemapping add --name=user2 --role SystemAdmin --auth-type HOST

# efa auth rolemapping add --name=admin --role SystemAdmin
--auth-type TACACS --auth-identifier 10.x.x.x

# efa auth rolemapping add --name=viewer --role SystemAdmin --type group
--auth-type LDAP --auth-identifier ldapconfig
```

This example assigns a role to a TACACS group.

```
# efa auth rolemapping add --name=admin --role SystemAdmin
--auth-type TACACS --auth-identifier 10.x.x.x

Successfully added the role mapping.
```

```

+-----+-----+
| attribute | value      |
+-----+-----+
| id        | 44         |
+-----+-----+
| name      | admin      |
+-----+-----+
| role      | SystemAdmin |
+-----+-----+
| type      | GROUP      |
+-----+-----+
| auth type | TACACS     |
+-----+-----+
| auth identifier | 10.x.x.x      |
+-----+-----+

# efa auth rolemapping add --name "cn=viewer,dc=extr,dc=com" --role NetworkOperator
--type group --auth-type LDAP --auth-identifier ldapconfig

Successfully added the role mapping.

+-----+-----+
| attribute | value      |
+-----+-----+
| id        | 45         |
+-----+-----+
| name      | cn=viewer,dc=extr,dc=com      |
+-----+-----+
| role      | NetworkOperator |
+-----+-----+
| type      | GROUP      |
+-----+-----+
| auth type | LDAP       |
+-----+-----+
| auth identifier | ldapconfig      |
+-----+-----+

```

```
# efa auth rolemapping show --auth-type LDAP --auth-identifier ldapconfig
# efa auth rolemapping show --auth-type TACACS --auth-identifier 10.x.x.x
```

This example displays all role assignments.

```
# efa auth rolemapping show

+-----+-----+-----+-----+-----+-----+
| ID | Name      | Role           | Type | Auth Type | Auth Identifier |
+-----+-----+-----+-----+-----+-----+
| 1  | user     | SystemAdmin   | USER | HOST      |                 |
+-----+-----+-----+-----+-----+-----+
| 43 | user2    | SystemAdmin   | USER | HOST      |                 |
+-----+-----+-----+-----+-----+-----+
| 44 | admin    | SystemAdmin   | GROUP | TACACS    | 10.x.x.x        |
+-----+-----+-----+-----+-----+-----+
| 45 | cn=viewer,dc=extr,dc=com | NetworkOperator | GROUP | LDAP      | ldapconfig      |
+-----+-----+-----+-----+-----+-----+

# efa auth rolemapping show --auth-type LDAP --auth-identifier ldapconfig
# efa auth rolemapping show --auth-type TACACS --auth-identifier 10.x.x.x
```

This example deletes the role for the user with ID 43.

```
# efa auth rolemapping remove --id 43
Deleted role mapping successfully
```

efa auth settings token

Configures and displays the expiration settings for authentication tokens.

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

```
efa auth settings token show
```

```
--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.

```
--type { cli | access | refresh }
```

Specifies the type of token you want to update.

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 auth tacacsconfig

Configure authentication on TACACS server.

```
efa auth tacacsconfig add [ --host hostname | --port port-num | --secret
sharedsecret | --protocol { CHAP | PAP } ]
```

--host *hostname*

Specifies the host name, IPv4, or IPv6 IP address of the TACACS server.

--port *port -num*

Specifies the port that the TACACS server is running on.

--protocol { CHAP | PAP }

Specifies the protocol used by the TACACS server.

--secret *sharedsecret*

Specifies the secret key used to connect to the TACACS server.

This example configures TACACS connection details.

```
efa auth tacacsconfig add --host 10.24.15.200 --port 49 --secret sharedsecret --protocol
CHAP
```

efa auth user

Configure users with CLI.

```
efa auth user add [ --userName string --emailID string --roles string
--isBlock { true | false} --password string --organization string --
mobileNumber string]
```



Note

--userName, --emailID, --roles and --password are mandatory parameters.

```
efa auth user show [ --userName string ]
```



```

efa auth user show-active [ --authType {TACACS | LDAP | LOCAL | HOST } ]
efa auth user delete [ --userName string ]
efa auth user update [ --userName string | --roles string | --mobileNumber string | organization string ]
efa auth user block-unblock [ --userName string --isBlock { true | false} ]
efa auth user reset-password [ --userName string --emailID string --password string ]
efa auth user change-password [ --userName string --oldPassword string --newPassword string ]
--authType { TACACS | LDAP | LOCAL | HOST }
    Specifies the authentication type of the active users to list.
--emailID string
    Specifies the email address of the user to configure.
--isBlock { true | false}
    Specifies the isBlock status of the user.
--mobileNumber string
    Specifies the mobile number of the user.
--newPassword string
    Specifies the new password for the user.
--oldPassword string
    Specifies the old password for the user.
--organization string
    Specifies the organization the user belong to.
--password string
    Specifies the password for the user.
--roles string
    Specifies the comma separated list of roles to assign to the user.
--userName string
    Specifies the user name to configure.

```

This example adds a new user, specifying user name, user email address, password, roles, organization and mobile number..

```

$ efa auth user add --userName testabc --emailID testabc@gmail.com --password password
--roles SystemAdmin --organization Org --mobileNumber 9122424356

```

This example blocks a specified user.

```

$ efa auth user block-unblock --userName=testUser --isBlock=true

```

This example unblocks a specified user.

```
$ efa auth user block-unblock --userName=testUser --isBlock=false
```

This example resets a password for a specified user.

```
$ efa auth user reset-password --userName testabc
--emailID testabc@gmail.com --password password123
```

```
Successfully updated the user configuration.
```

```
--- Time Elapsed: 2.69181206s ---
```

This example changes a password.

```
$ efa auth user change-password --oldPassword=test1234 --newPassword test456
```

```
The password was changed successfully.
```

```
--- Time Elapsed: 2.469497707s ---
```

This example deletes a specified user.

```
$ efa auth user delete --userName testabc
```

```
Deleted User configuration successfully.
```

```
--- Time Elapsed: 405.54142ms ---
```

This example lists active users with authentication type of HOST.

```
$ efa auth user show-active --authType=HOST
```

```
+-----+-----+-----+
| Username | Auth Type | Active Since |
+-----+-----+-----+
| anoop    | HOST      | 2022-03-07 10:01:52.094972524 +0530 IST |
+-----+-----+-----+
```

efa certificate device expiry show

List device certificates expiry.

```
efa certificate device expiry show [ --ip device-ips | --fabric string |
--fabric-all ]
```

```
--fabric string
```

Specifies the name of the fabric.

```
--fabric-all
```

Specifies all devices of all fabrics.

```
--ip device-ips
```

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

```
efa certificate device expiry show --ip 10.139.44.147
+-----+-----+-----+-----+
+-----+
|      IP      | Type  |      Expiry      | EFA/Device
Sync Status | Status |
+-----+-----+-----+-----+
+-----+
| 10.139.44.147 | HTTPS | 2025-08-06 18:11:47 +0000 GMT |
In Sync      | Success |
+-----+-----+-----+-----+
+-----+
| 10.139.44.147 | SyslogCA | 2033-09-03 17:58:17 +0000 GMT |
In Sync      | Success |
+-----+-----+-----+-----+
+-----+
| 10.139.44.147 | OAuth2  | 2033-09-03 18:07:50 +0000 GMT |
In Sync      | Success |
+-----+-----+-----+-----+
+-----+
Device Certificate Expiry Show
--- Time Elapsed: 53.13444312s ---
```

efa certificate device install

Installs HTTPS and OAuth2 certificates on SLX devices, and GRPC certificates on NPB devices..

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

By default, certificates are not installed.

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

Specifies the type of certificate you are installing.

--fabric *fabric-name*

Specifies the name of the fabric.

--force

Update the certificate even if it is already present.

--grpc-certificate

Specifies the file name of the grpc certificate.

--grpc-key

Specifies the file name of the grpc certificate key.

--https-certificate

Specifies the file name of the certificate.

--https-key

Specifies the file name of the certificate key.

--ip *ip-addr*

Specifies a comma-separated list of the IP addresses of the SLX devices on which you want to install the certificate. Example: 10.1.1.13,10.1.1.50,10.1.1.101.

During the registration of an SLX device in XCO, the following configuration changes are made on the device:

- The public certificate for verifying an XCO token is copied to the device as an OAuth2 certificate.
- XCO 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."

During the registration of a NPB device in XCO, the following configuration changes are made on the device:

- XCO generates the GRPC certificate and key for the NPB device. The certificate and key are imported into the device.
- Configure the grpc-server with certificate id.

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.

This example installs the HTTPS certificate on two devices.

```
efa certificates device install --ip 10.139.44.147-148 --certType https
+-----+-----+
| IP Address  | Status |
+-----+-----+
| 10.139.44.148 | Success |
+-----+-----+
| 10.139.44.147 | Success |
+-----+-----+
```

This example installs the HTTPS certificate the devices in fabric fabric1. It will update the certificates even if already present.

```
efa certificates device install --fabric fabric1 --certType https --force
+-----+-----+
| IP Address  | Status |
+-----+-----+
| 10.139.44.148 | Success |
+-----+-----+
| 10.139.44.147 | Success |
+-----+-----+
```

efa certificate expiry show

Display certificate expiry list.

efa certificate expiry show

```
(efa:root)root@ubuntu:~# efa certificate expiry show
+-----+-----+
|           Name           |           Expiry           |
+-----+-----+
| App Server Certificate   | 2026-03-08T08:29:14Z |
+-----+-----+
| Default Intermediate CA  | 2033-03-06T08:29:13Z |
+-----+-----+
|   Default Root CA      | 2043-03-04T08:29:13Z |
+-----+-----+
|   JWT Certificate      | 2033-03-06T08:37:14Z |
+-----+-----+
| K3s Server Certificate   | 2024-03-08T08:30:21Z |
+-----+-----+
|           K3s CA        | 2033-03-06T08:30:21Z |
+-----+-----+
Certificate Expiry List detail show
--- Time Elapsed: 270.257539ms ---
(efa:root)root@ubuntu:~#
```

efa certificate server

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

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

```
--cacert cert-filename
```

Specifies the file name of the third-party certificate.

```
--certificate cert-filename
```

Specifies the file name of the certificate.

```
--key key-filename
```

Specifies the file name of the certificate key.

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 XCO.

efa certificate server renew [**--cert-type** *string*] This is not applicable if third-party certificates are installed on the system.

--cert-type *string*

Specifies the type of certificate you are renewing. Valid values are (server | token | root-ca | intermediate-ca | k3s-ca | k3s-server) (default "server").

This example renews a certificate of type token.

```
(efa:extreme)extreme@tpvm:/apps$ efa certificate server renew --cert-type=token
Certificate renewal is successful.
--- Time Elapsed: 27.233017418s ---
```

efa fabric clone

Clones source fabric.

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

--destination

New name of the cloned fabric.

--source

Name of the fabric to be cloned.

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.

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.

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

--force

Forces the configuration on the devices.

--name *fabric-name*

Specifies the name of the fabric.

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.

```
efa fabric configure --name common-fabric
WARNING: 'fabric configure' will result in configuration change on the devices which are
in 'cfg refreshed' or 'fabric setting refreshed' or 'cfg refresh error' state.
Please check 'fabric show' to see the 'cfg refreshed' or 'fabric setting refreshed'
or 'cfg refresh error' devices. Please confirm if you want to continue with 'fabric
configure' [y/n]?
y
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: 30.549013784s ---
```

efa fabric create

Creates a fabric.

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

--description *description*

Describes the product.

--name *fabric-name*

Specifies the name of the fabric.

--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.

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

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

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.

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

--device

Device IP address.

--reference-fabric

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

```
$ 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.

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

--device

The device IP address.

--name

Name of the fabric to which the device belongs.

```
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.

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

--force

Force DRC on fabric device. Optional.

--ip *string*

Specifies the device IP.

--name *string*

Specifies the name of the fabric.

--reconcile

Reconciles the device configuration.

efa fabric debug service lock

Lists all lock statuses for fabric services.

```
efa fabric debug service lock
```

efa fabric debug service smartdrc

Enables debug of smartdrc.

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

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

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

```
--enable=true
```



```
--enable=false
```

--help

Help for smartdrc debug.

efa fabric debug set

Set up debug level and debug for a module.

```
efa fabric debug set [ --level { debug | info | debugdb | nodebugdb } ]
```

```
--level { debug | info | debugdb | nodebugdb }
```

Required. Specifies the debug level.

```
$ efa fabric debug set --level debug

Level debug set successfully.
--- Time Elapsed: 104.471ms ---

$ efa fabric debug set --level info
Level info set successfully.

--- Time Elapsed: 104.471ms ---

$ efa fabric debug set --level debugdb

Level debugdb set successfully.
--- Time Elapsed: 91.1052ms ---

$ efa fabric debug set --level nodebugdb

Level nodebugdb set successfully.
--- Time Elapsed: 91.1052ms ---
```

efa fabric delete

Deletes the fabric from inventory.

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

--force

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

```
--name fabric-name
```

Name of the fabric to be deleted.

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.

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.

```
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 ]
```

--asn *local-asn*

Specifies the local ASN.

--hostname *hostname*

Specifies the host name.

--ip *device-ip*

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

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

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

--loopback *port-num*

Specifies the loopback port number.

--name *fabric-name*

Specifies the name of the fabric.

--password *password*

Specifies the password for the device.

--pod *name*

Specifies the name of the pod.

--rack *name*

Specifies the name of the rack.

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

Specifies the device role.

--username *username*

Specifies the user name for the device.

--vtep-loopback *id*

Specifies the VTEP loopback ID.

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.

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.

```
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 ]
```

--border-leaf *list*

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

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

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

--border-leaf-rack *list*

Specifies the border leaf rack name.

--five-stage-pod *name*

Specifies the name of the super-spine pod.

--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.

--name *fabric-name*

Specifies the name of the fabric.

--password *password*

Specifies the password for the list of devices.

--rack *name*

Specifies the rack name.

--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.

--three-stage-pod *name*

Specifies the name of the leaf or spine pod.

--username *username*

Specifies the user name for the list of devices.

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.

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.

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

--ip *device-ips*

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

--name *fabric-name*

Specifies the name of the fabric.

--no-device-cleanup

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

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.

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.

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

--export *string*

Exports error details to a CSV file.

--name *fabric-name*

Specifies the name of the fabric.

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.

```
$ efa fabric error show
+-----+-----+-----+-----+-----+
| FABRIC NAME | IP ADDRESS | ROLE | ERROR TYPE | ERROR REASON |
+-----+-----+-----+-----+-----+
| fab-non-clos| NA         | NA   | clos error | No Devices    |
| fab-clos    | 10.20.246.12 | Leaf | clos error | Leaf Device 10.20.246.12 not |
|             |             |     |           | connected to Spine Device   |
|             |             |     |           | 10.20.246.19                |
| fab-clos    | 10.20.246.12 | Leaf | clos error | Leaf Device 10.20.246.12 not |
|             |             |     |           | connected to Spine Device   |
|             |             |     |           | 10.20.246.20                |
| fab-clos    | 10.20.246.19 | Spine | clos error | Spine Device 10.20.246.19    |
|             |             |     |           | not connected to Leaf Device |
|             |             |     |           | 10.20.246.12                |
| fab-clos    | 10.20.246.20 | Spine | clos error | Spine Device 10.20.246.20    |
|             |             |     |           | not connected to Leaf Device |
|             |             |     |           | 10.20.246.12                |
| fab-clos    | 10.20.246.19 | Spine | clos error | Spine Device 10.20.246.19    |
|             |             |     |           | connected to Spine Device   |
|             |             |     |           | 10.20.246.20                |
| fab-clos    | 10.20.246.20 | Spine | clos error | Spine Device 10.20.246.20    |
|             |             |     |           | connected to Spine Device   |
|             |             |     |           | 10.20.246.19                |
+-----+-----+-----+-----+-----+
```

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

efa fabric execution

Displays or deletes the execution history for the fabric service.

```
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 ]

--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.

```

```
efa fabric execution show
```

efa fabric health show

Display the health status of a fabric.

```

efa fabric health show [ --name string | --detail ]

--detail
  Optional. When specified, displays the health status detail of the fabric.

--name string
  Specifies the name of the fabric.

```

The following example displays the health status of the specified fabric.

```

$ efa fabric health show --name fab3
=====
Fabric Name           : fab3
Fabric Type          : clos
Fabric Health        : black
  Fabric Status      : configure-success
  Fabric Level Physical Topology Health : green
Fabric Device Health
+-----+-----+-----+-----+-----+
| IP ADDRESS | ROLE | CONFIG STATE HEALTH | OPER STATE HEALTH | DEVICE HEALTH |
+-----+-----+-----+-----+-----+
| 10.20.246.6 | Leaf | black                | black              | black          |
| 10.20.246.5 | Leaf | black                | black              | black          |
| 10.20.246.1 | Spine | green                | green              | green          |
| 10.20.246.2 | Spine | green                | green              | green          |
| 10.20.246.4 | Leaf | black                | black              | black          |
| 10.20.246.3 | Leaf | black                | black              | black          |
+-----+-----+-----+-----+-----+

```



```

Dev State : provisioned
App State : cfg refresh error
Operational State Health : black
Cluster Health : black
Operational State : true
Peer Operational State : false
Peer Keepalive Operational State : true
Physical Topology Device Health : red
Device physical topology errors
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE | ERROR |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.20.246.5 | Leaf | | | |
10.20.246.6 | Leaf | | | |
| | | missing-links |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
Underlay Topology Device Health : black
Device underlay topology errors
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN | VRF | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE | ERROR |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 10.20.246.5 | 10.20.246.6 | 172.31.254.156 | 10.20.20.14 |
65001 | 65001 | default-vrf | ipv4 |
unicast | CONN | session_not_established |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
Device IP [Role] : 10.20.246.1 [Spine]
Device Health : green
Configuration State Health : green
Dev State : provisioned
App State : cfg in-sync
Operational State Health : green
Physical Topology Device Health : green
Underlay Topology Device Health : green
-----
Device IP [Role] : 10.20.246.2 [Spine]
Device Health : green
Configuration State Health : green
Dev State : provisioned
App State : cfg in-sync
Operational State Health : green
Physical Topology Device Health : green
Underlay Topology Device Health : green
-----
Device IP [Role] : 10.20.246.4 [Leaf]
Device Health : black
Configuration State Health : black
Dev State : provisioned

```



```

App State : cfg refresh error
Operational State Health : black
Cluster Health : black
  Operational State : true
  Peer Operational State : false
  Peer Keepalive Operational State : true
Physical Topology Device Health : red
Device physical topology errors
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE | ERROR |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.20.246.4 | Leaf | | |
10.20.246.3 | Leaf | | |
| | missing-links |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
Underlay Topology Device Health : black
Device underlay topology errors
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN | VRF | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE | ERROR |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.20.246.4 | 10.20.246.3 | 172.31.254.2 | 10.20.20.16 |
65002 | 65002 | default-vrf | ipv4 |
unicast | CONN | session_not_established |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
Device IP [Role] : 10.20.246.3 [Leaf]
Device Health : black
Configuration State Health : black
  Dev State : provisioned
  App State : cfg refresh error
Operational State Health : black
Cluster Health : black
  Operational State : true
  Peer Operational State : false
  Peer Keepalive Operational State : true
Physical Topology Device Health : red
Device physical topology errors
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| SOURCE NODE IP | SOURCE NODE ROLE | SOURCE NODE POD | SOURCE NODE INTERFACE |
DESTINATION NODE IP | DESTINATION NODE ROLE | DESTINATION NODE POD | DESTINATION NODE
INTERFACE | ERROR |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 10.20.246.3 | Leaf | | |
10.20.246.4 | Leaf | | |

```

```

|                                     | missing-links |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
Underlay Topology Device Health      : black
Device underlay topology errors
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| SOURCE DEVICE IP | DESTINATION DEVICE IP | SOURCE DEVICE ROUTER ID | NEIGHBOR IP |
SOURCE DEVICE ASN | DESTINATION DEVICE ASN | VRF      | NEIGHBOR AFI STATE | NEIGHBOR
SAFI | UNDERLAY STATE |          ERROR          |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| 10.20.246.3      | 10.20.246.4          | 172.31.254.246          | 10.20.20.17 |
65002             | 65002                | default-vrf | ipv4           |
unicast          | CONN                 | session_not_established |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+-----+-----+-----+
=====
--- Time Elapsed: 52.950728ms ---

```

efa fabric migrate

Migrate 3-stage CLOS fabric to 5-stage CLOS fabric.

```
efa fabric migrate [ --type string | --source-fabric string | --destination-3-stage--leaf-spine-pod string | --destination-3-stage--border-leaf-pod string | --super-spine-asn-block string | --super-spine-peer-group string ]
```

--destination-3-stage-border-leaf-pod *string*

Specifies the name of the 3-stage POD into which the border-leaf devices (of the 3-stage CLOS fabric) need to be moved during migration.

--destination-3-stage--leaf-spine-pod *string*

Specifies the name of the 3-stage POD into which the leaf and spine devices (of the 3-stage CLOS fabric) need to be moved during migration.

--source-fabric *string*

Specifies the name of the 3-stage CLOS fabric to migrate.

--super-spine-asn-block *string*

Specifies the ASN block to be used by the super-spine devices of the migrated 5-stage CLOS fabric.

--super-spine-peer-group *string*

Specifies the Peer Group to be used by the spine devices of the migrated 5-stage CLOS fabric.

--type *string*

Specifies the type of migration. Default "3-to-5-stage".

efa fabric setting

Displays or updates the IP fabric settings.

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

```
efa fabric setting update [--name fabric-name | --p2p-link-range fabric-name | --loopback-scheme { uniform|granular } | --loopback-ip-range ip-pool | --rack-l3-backup-ip-range ip-pool | --router-id-loopback-range string | --vtep-loopback-ip-range string | --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 ttl | --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 | --rack-mct-scheme { static | dynamic } | --rack-mct-ports port-list | --rack-ld-mct-ports port-list | --mctlink-ip-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 | --overlay-gateway-broadcast-local-bias-enable { yes | no }]
```

```
--name fabric-name
```

Specifies the name of the fabric.

```
--advanced
```

List advanced fabric parameters.

```
--p2p-link-range ip-pool
```

Specifies the IP address pool used for P2P link configuration.

```
--loopback-scheme { uniform | granular }
```

Specifies the scheme of the IP allocation. Default is uniform. Uniform - common IP pool for RouterID and VTEP; Granular - separate IP pools for RouterID and VTEP.

```
--loopback-ip-range ip-pool
```

IP address pool for unnumbered and VTEP IP if Loopback Scheme is uniform.

```
--rack-l3-backup-ip-range ip-pool
```

Specifies the IP address pool for Layer 3 backup.

--router-id-loopback-ip-range *ip-pool*

IP address pool for RouterID Loopback interface if Lookback Scheme is granular.

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

IP address pool for VTEP Loopback interface if Lookback Scheme is granular.

--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.
- rack-mct-scheme** *scheme-type*
Specifies the rack mct scheme type. Options are static or dynamic.
- rack-mct-ports** *port-list*
Specifies the comma separated list of rack mct ports. Example: 0/46,0/47
- rack-ld-mct-ports** *port-list*
Specifies the comma separated list of low density rack mct ports. Example: 0/30,0/31
- 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.

---overlay-gateway-broadcast-local-bias-enable { **yes** | **no** }

Enable/Disable Local Bias for LVTEP BUM Traffic Handling.

This example shows the output for a fabric with a uniform IP allocation scheme.

```
efa fabric setting show --name fs-1 | grep -i loopback
```

Loopback Scheme	uniform	
Loopback IP Range	172.31.254.0/24	
RouterID Loopback IP Range	172.31.128.0/24	
VTEP Loopback IP Range	172.31.64.0/24	
Loopback Port Number	1	
VTEP Loopback Port Number	2	

This example shows the output for a fabric with a granular IP allocation scheme.

```
efa fabric setting show --name fs-1 | grep -i loopback
```

Loopback Scheme	granular	
Loopback IP Range	172.31.254.0/24	
RouterID Loopback IP Range	172.31.128.0/24	
VTEP Loopback IP Range	172.31.64.0/24	
Loopback Port Number	1	
VTEP Loopback Port Number	2	

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
```

This example displays the advanced mct parameters of the specified fabric.

```
efa fabric setting show --name fabric1 --advanced | grep -i mct
```

MCT Link IP Range	10.20.20.0/24	
-------------------	---------------	--

MCT PortChannel	64	
Rack MCT Scheme	Dynamic	
Rack MCT Ports		
Rack Low Density MCT Ports		

efa fabric show

Displays the details of the fabric.

efa fabric show [--name *fabric-name*]

--name *fabric-name*

Specifies the name of the fabric.

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.

This example displays details for the fabric called common-fabric.

```
efa fabric show -name common-fabric
```

```
Fabric Name: common-fabric, Fabric Description: , Fabric Type: non-clos, Fabric Status:
configure-success, Fabric Health: Green
```

```

+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| IP ADDRESS | RACK | HOST NAME | ASN | ROLE |
| DEVICE STATE | APP STATE | CONFIG GEN REASON | PENDING CONFIGS | VTLB ID | LB ID |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.64.208.22 | r1 | SW-22 | 4200065535 | BorderLeaf | |
| provisioned | cfg-in-sync | NA | NA | 2 | 1 |
| 10.64.208.23 | r1 | SW-23 | 4200065535 | BorderLeaf |
| provisioned | cfg in-sync | NA | NA | 2 | 1 |
| 10.64.208.28 | room1-rack2 | SW-28 | 4200000000 | Leaf |
| provisioned | cfg in-sync | NA | NA | 2 | 1 |
| 10.64.208.29 | room1-rack2 | SW-29 | 4200000000 | Leaf |
| provisioned | cfg in-sync | NA | NA | 2 | 1 |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+

```

efa fabric show-config

Displays the configuration of the specified fabric.

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

--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.

--name *fabric-name*

Specifies the name of the fabric.

```
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.

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

--name *fabric-name*

Specifies the name of the fabric.

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.

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.

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

--name *name*

Specifies the name of the fabric.

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.

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

--name *name*

Specifies the name of the fabric.

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.

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

--name *name*

Specifies the name of the fabric.

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     |
| NODE INTERFACE | MULTI HOMED    | NODE DEVICE     | MULTI HOMED    |
+-----+-----+-----+-----+
| 0/39           | TRUE           | FALSE           |
| 0/40           | TRUE           | FALSE           |
| 0/46           | TRUE           | TRUE            |
| 0/46           | TRUE           | TRUE            |
+-----+-----+-----+-----+
```

efa health debug clear

Clear the health of the specified subtree and resource

```
efa health debug clear [ --resource string ]
```

--resource *string*

Specifies the subtree of the resource to clear. If not specified, the health of the tree from Root node is cleared.

```
efa health debug clear --resource /App/Component/Asset/Device?device_ip=<ip of Device>
```

efa health detail show

Display the list of health detail.

```
efa health detail show [ --resource string ]
```

--resource *string*

Specifies the health resource path associated to the alert being sent. Example:
--resource /App/System/Security/Certificate.

efa health inventory show

Display the list of health inventory.

```
efa health inventory show [ --resource string ]
```

--resource *string*

Specifies the health resource path associated to the alert being sent. Example:
--resource /App/System/Security/Certificate.

```
(efa:root)root@ubuntu:~/certs891# efa health inventory show
+-----+
|                               Resource                               |
+-----+
|                               /App/System/HA/Nodes/Services         |
+-----+
|                               /App/System/HA/Nodes/Node             |
+-----+
|                               /App/System/Storage?mount_point=&node_ip= |
+-----+
|                               /App/System/Storage?mount_point=/&node_ip=10.21.89.1 |
+-----+
|                               /App/System/Security/Certificate?type=unknown_certificate |
+-----+
|                               /App/System/Security/Certificate?device_ip=&type=unknown_certificate |
+-----+
|                               /App/System/Security/Certificate?type=app_server_certificate |
+-----+
|                               /App/System/Security/Certificate?type=third_party_ca |
+-----+
|                               /App/System/Security/Authentication?server= |
+-----+
|                               /App/Component/Asset/Device?device_ip= |
+-----+
--- Time Elapsed: 62.125206ms ---
```

efa health show

Display the list of health resources.

```
efa health show [ --resource string | --detail ]
```

--detail

Detailed information for health resources.

--resource *string*

Specifies the health resource path associated to the alert being sent. Example:
--resource /App/System/Security/Certificate.

efa inventory admin-state

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

```
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 ]
```

--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.

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 .

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.

```
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 ]
```

```
--file-dump filename
```

Saves configuration text to the specified file.

```
--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.

--show-config

Displays configuration text.

--uuid *cb-uuid*

Specifies the ID of the configuration backup.

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.

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.

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*]

--ip *ip-addr*

Specifies the IP address of the device.

--key *ip-addr*

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

--no-reboot

Configures replay without rebooting the device.

--ssid *ssid*

Specifies the SSID of the config-backup job.

--startup-config

Copies the backup file to the startup-config file.

--uuid *cr-uuid*

Specifies the config replay ID or the UUID of config backup job.

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.

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 device-adapter-status

List status of all device adapters in inventory

```
efa inventory debug device-adapter-status [ --ip device-ips | --fabric
string | --fabric-all ]
```

--fabric *string*

Specifies the name of the fabric.

--fabric-all

All devices for all fabrics.

--ip *device-ips*

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

```
efa inventory debug device-adapter-status --ip 10.24.80.58
+-----+-----+-----+-----+-----+
|      IP      | Adapter      | Protocol | State  | Error Status |
+-----+-----+-----+-----+-----+
| 10.24.80.58 | SwitchRestAdapter | https   | Connected | Success      |
+-----+-----+-----+-----+-----+
| 10.24.80.58 | SwitchSSHAdapter  | SSH     | Connected | Success      |
+-----+-----+-----+-----+-----+
| 10.24.80.58 | SwitchNcAdapter  | NewConf | Connected | Success      |
+-----+-----+-----+-----+-----+

Device Adapter Status

--- Time Elapsed: 5.867081913s ---
```

efa inventory debug devices-lock

Displays lock status of the devices.

efa inventory debug devices-lock

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.

```
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.

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

--fabric *fabric-name*

Fabric with which devices are associated.

--ip *ip-list*

List of IPs to unlock.

--servicelock

Specify to unlock all service locks.

```
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.

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

--fabric *fabric-name*

Fabric with which devices are associated.

--ip *ip-list*

List of IPs to unlock.

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.

```
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-fwld --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.

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

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

```
--ip device-ips
```

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

```
--ipv4-only
```

Optional flag to only clear IPv4 routes.

```
--ipv6-only
```

Optional flag to only clear IPv6 routes.

```
--fabric fabric-name
```

Specifies the name of the fabric.

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.

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

```
--ip
```

Specifies the IP address of the device to be compared.

efa inventory device delete

Deletes the device.

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

```
--fabric fabric-name
```

Fabric name for which devices needs to deleted.

```
--ip ip-addr
```

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

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.

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

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

```
--fabric fabric-name
```

Specifies the name of the fabric.

```
--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.

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.

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

The default discovery interval is one hour.

```
--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.

--hour *hours*

Specifies the interval in hours.

--ip *ip-addr*

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

--min *minutes*

Specifies the interval in minutes.

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.

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.

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

--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`.

--fabric *fabric-devices*

Specifies all devices in the specified fabric.

--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`.

--role *role-devices*

Specifies devices in the fabric, based on their role.

```
# 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 activate

This command will trigger the reloading of devices for a given fabric, devices or a prepared list group.

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

--commit

Configure commit operation to be performed after firmware activate.

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *ip-addr*

Specifies a comma-separated list of device IP addresses.

--prepared-list-name *prepared-list-name*

Prepared list name.

efa inventory device firmware-download commit

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

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

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *ip-addr*

Specifies a comma-separated list of device IP addresses.

--prepared-list-name *prepared-list-name*.

Prepared list name.

```
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 enable the run to be retried.

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

--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 	

noMaint	drc	enable_on_reboot	FWDL Reboot Behavior	Comments
True	False	False	<ul style="list-style-type: none"> No maintenance mode NO DRC 	
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	

--fabric *fabric-name*

Specifies the name of the fabric.

--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.

--noActivate

Configure Activation in Firmware Download. Device will not reload after the downloading the firmware. To activate the firmware, you must run the `efa inventory device firmware-download activate` command on the device.

--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.

--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.

--prepared-list-name *prepared-list-name*

Prepared list name.

You can 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.

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

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

--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.

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

Path to the target firmware build.

--firmware-host *firmware download host*

Firmware download host IP address.

--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.

--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.

--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 XCO hosted node's SLX firmware to be upgraded simultaneously.

--prepared-list-name *prepared-list-name*

Prepared list name.



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.
- XCO 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.

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 XCO 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]

```

This example prepares an NPB device for firmware upgrade.

```

efa inventory device firmware-download prepare add --firmware-host=10.37.135.12 --
firmware-directory=/var/lib/tftpboot/SLXOS3C/slxos18s.1.03c --prepared-list-name testList
--ip 10.37.128.72
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+

```

```

| Group | IP Address | Host Name | Model | Chassis Name |
ASN | Role | Current Firmware | Firmware Host | Firmware Directory
| Target Firmware | Force | Full Install | Last Update Time |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 1 | 10.37.128.72 | SLX | 3000 | BR-SLX9240 | 0
| | 18s.1.03e | | 10.37.135.12 | /var/lib/tftpboot/SLXOS3C/slxos18s.1.03c
| 18s.1.03c | false | false | 2022-11-02 04:01:32 -0400 EDT |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
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.

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

--fabric *fabric-name*

Specifies the name of the fabric.

--prepared-list-name *prepared-list-name*

Prepared list name.

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.

This example prepares the fabric and lets XCO 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.

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

--fabric *fabric-name*

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

--group *group-number*

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

--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.

--prepared-list-name *prepared-list-name*

Prepared list name.

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

Firmware Download Prepare Add (Fabric-based)

Prepare the fabric and let XCO 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]

```

Firmware Download Prepare Delete (NPB)

```

efa inventory device firmware-download prepare remove --prepared-list-name testList
+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+
| Group | IP Address | Host Name | Model | Chassis Name |
ASN | Role | Current Firmware | Firmware Host | Firmware
Directory | Target Firmware | Force | Last Update Time |
+-----+-----+-----+-----+-----+-----+
+-----+
| 1      | 10.37.128.72 | SLX      | 3000 | BR-SLX9240 | 0
|      | 18s.1.03e | 10.37.135.12 | /var/lib/tftpboot/SLXOS3C/slxos18s.1.03c
| 18s.1.03c | false | 2022-11-02 04:01:32 -0400 EDT |
+-----+-----+-----+-----+-----+-----+
+-----+
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.

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

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *ip-addr*

Specifies a comma-separated list of device IP addresses.

--prepared-list-name *prepared-list-name*

Prepared list name.

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.

```
efa inventory device firmware-download show [ --fabric fabric-name | --ip device-ip-address | --prepared-list-name prepared-list-name | --execution-id execution-id | --type { fabric | npb } ]
```

--execution-id *execution-id*

Execution id provided when running the efa inventory device firmware-download activate command.

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *device-ip-address*

Specifies a comma-separated list of device IP addresses.

--prepared-list-name *prepared-list-name*

Prepared list name.

--type { **fabric** | **npb** }

Specifies the device type. Options are fabric or npb.

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]

```

This example shows a firmware download with --type fabric

```

$ efa inventory device firmware-download show --type fabric
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| Group | IP Address | Host
Name | Model | Chassis Name | ASN | Role | Current Firmware |
Target Firmware | Update State | Status | Detailed Status | Force | Full
Install | Last Update Time | User Input Flags | Group Execution Policy |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1      | 10.20.246.7 | SN1      | 4001
| BR-SLX9640 | 4200000000 | Leaf | 20.4.2slxos20.4.2c_230602_1733 | 20.4.2a
| Completed | Firmware Not Activated | | true | false
| 2023-06-13 07:50:49 +0530 IST | D,C | stop-on-error |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+
| 1      | 10.20.246.8 | SN2      | 4001
| BR-SLX9640 | 4200000000 | Leaf | 20.4.2slxos20.4.2c_230602_1733 | 20.4.2a

```

```

| Completed | Firmware Not Committed | | true | false
| 2023-06-13 08:33:38 +0530 IST | A | continue-on-error |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
Firmware Download Show Details

User Input Flags
D - Firmware Download, A - Firmware Activate, C - Firmware Commit,
R - Firmware Restore, MM - Firmware Maintenance Mode, DRC - Firmware Drift and Reconcile
Group Execution Policy: per-device-policy
Firmware Download Show Overall Status [Device Update In Progress]
If Firmware download execution initiated
with '--noAutoCommit' flag, please commit/restore after completion using:
efa inventory device firmware-download commit --ip 10.20.246.7, 10.20.246.8
efa inventory device firmware-download restore --ip 10.20.246.7, 10.20.246.8
If Firmware download execution initiated
with '--noActivate' flag, please activate after completion using:
efa inventory device firmware-download activate --ip 10.20.246.7, 10.20.246.8
--- Time Elapsed: 53.3182ms ---

```

This example shows a firmware download for a specified fabric.

```

$ efa inventory device firmware-download show --fabric fab
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| Group | IP Address |
Host Name | Model | Chassis Name | ASN | Role | Current Firmware |
| Target Firmw are | Update State | Status | Detailed
Status | Force | Full Install | Last Update Time | User Input Flags |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 1 | 10.20.246.15 | Avalanche-01
| 4000 | BR-SLX9540 | 0 | Leaf | 20.4.3b |
20.4.3b | Completed | Firmware Not Committed |
| false | false | 2023-05-20 14:46:27 +0530 IST | D,A,C,MM |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 1 | 10.20.246.16 |
Avalanche-02 | 4000 | BR-SLX9540 | 0 | Leaf | 20.4.2slxos20.4.2c_230516_1439 |
20.4.3a | Completed | Firmware Not Committed |
| false | false | 2023-05-22 16:06:17 +0530 IST | D,A,MM |
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
Firmware Download Show Details
User Input Flags
D - Firmware Download, A - Firmware Activate, C - Firmware Commit,
R - Firmware Restore, MM - Firmware Maintenance Mode, DRC - Firmware Drift and Reconcile

```

efa inventory device health status

Displays device health status.

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

--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 XCO or has drifted (refreshed or deleted).

```
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 ]
```

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

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

--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.

--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`.

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

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

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 XCO.

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

--ip *device-ip*

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

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.

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

```
--enable { true | false }
```

Enables redundant management.

```
--if-name
```

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

```
--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.

The configuration you set is persisted in the XCO 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, XCO 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.

```
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.

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

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

--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.

--if-type *eth*

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

--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.

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

State. Can be up or down.

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.

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.

```
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 } ]
```

--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.

--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.

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

Configures breakout mode for the ports.

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.

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-description

Set the description for the interface on the device

```
efa inventory device interface set-description [ --ip string | if-type
string | if-name string | --description string ]
```

--description *string*

String value to set the description on device interface. Maximum length 63.

if-name *string*

Comma separated range of Interface names addresses and should be unique.

Example: 0/50-52,0/45,0/5:4,0/5:1-2

if-type *string*

Optional. Interface type default value is 'eth'. Current support is only for 'eth'.

--ip *string*

Comma separated range of device IP addresses. For example:
10.1.1.13,10.1.1.50-53,10.1.1.101.

The following example sets the description on the specified device interface. Only a partial output from the list command is shown.

```
$ efa inventory device interface set-description -ip 10.64.164.47 -if-name 0/14 --description "interface eth 0/14"
```

DeviceIP	ID	Name	Interface Type	Description	Result	Reason
10.64.164.47	95	0/14	ethernet	interface eth 0/14	Success	

```
$ efa inventory device interface list --ip 10.64.164.47
```

DeviceIP	Name	Interface Type	Admin Status	Oper Status	Description	Line Speed	MAC	Switchport Mode	IP Address	App State
10.64.164.47	0/1	ethernet	up	down						
unknown	00:04:96:b8:ce:0e	unknown						cfg-in-sync		
unknown	0/2	ethernet	up	down						
unknown	00:04:96:b8:ce:0f	unknown						cfg-in-sync		
unknown	0/3	ethernet	up	down						
unknown	00:04:96:b8:ce:10	unknown						cfg-in-sync		
unknown	0/4	ethernet	up	down						
unknown	00:04:96:b8:ce:11	unknown						cfg-in-sync		
100Gbps	0/5	ethernet	up	up						
unknown	00:04:96:b8:ce:12	unknown						cfg-in-sync		
unknown	0/13	ethernet	up	down						
unknown	00:04:96:b8:ce:1a	unknown						cfg-in-sync		
interface eth 0/14	0/14	ethernet	up	down						
unknown	00:04:96:b8:ce:1b	unknown						cfg-in-sync		
unknown	0/15	ethernet	up	down						
unknown	00:04:96:b8:ce:1c	unknown						cfg-in-sync		

```

|          | 0/16 | ethernet      | up      | down      |
| unknown  | 00:04:96:b8:ce:1d | unknown      |         |         | cfg-in-sync |
+-----+-----+-----+-----+-----+
.
.
.
+-----+-----+-----+-----+-----+
|          | 0/30 | ethernet      | up      | down      |
| unknown  | 00:04:96:b8:ce:2b | unknown      |         |         | cfg-in-sync |
+-----+-----+-----+-----+-----+
|          | 0/31 | ethernet      | up      | up        |
clusterPeerIntfMember | 100Gbps | 00:04:96:b8:ce:2c | unknown      |         |         | cfg-in-sync |
+-----+-----+-----+-----+-----+
|          | 0/32 | ethernet      | up      | up        |
clusterPeerIntfMember | 100Gbps | 00:04:96:b8:ce:2d | unknown      |         |         | cfg-in-sync |
+-----+-----+-----+-----+-----+
.
.
.
+-----+-----+-----+-----+-----+
|          | 0/40 | ethernet      | up      | down      |
| unknown  | 00:04:96:b8:ce:35 | unknown      |         |         | cfg-in-sync |
+-----+-----+-----+-----+-----+
Interface Details

```

efa inventory device interface set-fec

Configures forward 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 }]
```

--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.

--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.

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

Specifies the FEC mode for the ports.

The configuration you set is persisted in the XCO 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.

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

--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.

--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.

--sampling-time *value*

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

--toggle-threshold *value*

Valid values are 1 through 50.

--wait-time *value*

Valid values, in seconds, are 0 through 65565.

The configuration you set is persisted in the XCO database.

DRC and idempotency are supported.

```
efa inventory device interface set-link-error-disable --ip 10.20.246.10
--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.

```
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 ]
```

--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.

--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.

--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.

--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.

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.

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 unset-mtu

Unsets the device interface MTU.

```
efa inventory device interface unset-mtu [ --ip device-ip | --if-type eth
  | --if-name if-list | --mtu | --ip-mtu | --ipv6-mtu ]
```

--if-name *if-list*

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

--if-type **eth**

Optionally specifies the interface type. Default value is 'eth'. Currently only interface of 'eth' is supported.

--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.

--ip-mtu

Specifies that IP MTU is to be removed.

--ipv6-mtu

Specifies that IPv6 MTU is to be removed.

--mtu

Specifies that MTU is to be removed.

This example unsets both mtu and ip-mtu from an interface.

```
efa inventory device interface unset-mtu --ip 10.20.24.10 --if-name=0/6 --mtu -ip-mtu

+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| DeviceIP   | ID | Name | Interface Type | MTU |
IP MTU | IPv6 MTU | Result | Reason |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.24.10 | 451 | 0/6 | ethernet      | true |
true  | false  | Success |              |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+

Interface MTU Details

--- Time Elapsed: 6.469215243s ---
```

efa inventory device interface set-speed

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

```
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 }]
```

--if-name *if-list*

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

--if-type **eth**

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

--ip *device-ip*

Specifies a comma-separated list of device IP addresses.

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

Specifies the speed for the port.

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.

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 |
+-----+-----+-----+-----+-----+-----+
| 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.

```

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

```

--if-name *if-list*

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

--if-type *eth*

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

--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.

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.

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 |
+-----+-----+-----+-----+-----+-----+
| 10.10.10.10 | | | | 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-description

Unset the description for the interface on the device

```
efa inventory device interface unset-description [ --ip string | if-type
string | if-name string ]
```

if-name *string*

Comma separated range of Interface names addresses and should be unique.

Example: 0/50-52,0/45,0/5:4,0/5:1-2

if-type *string*

Optional. Interface type default value is 'eth'. Current support is only for 'eth'.

--ip *string*

Comma separated range of device IP addresses. For example:

10.1.1.13,10.1.1.50-53,10.1.1.101.

```
(efa:extreme)extreme@SLX12TPVM95:~$ efa inventory device interface unset-description --
if-name "0/21:3" --ip 10.64.212.11
+-----+-----+-----+-----+-----+-----+
| DeviceIP | ID | Name | Interface Type | Result | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.64.212.11 | 863 | 0/21:3 | ethernet | Success | |
+-----+-----+-----+-----+-----+-----+
Interface Description Details
--- Time Elapsed: 6.117291034s ---
```

efa inventory device interface unset-fec

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

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

```
--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.
```

```
--if-type eth
```

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

```
--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 configuration you set is persisted in the XCO 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.

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

```
--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.
```

```
--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 configuration you set is persisted in the XCO database.

DRC and idempotency are supported.

```
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.

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

--fabric *name*

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

--ips *device-ips*

Specifies comma separated device IPs.

--orphan

Fetches all devices that are not associated to a fabric.

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

Specifies device role.

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 } ]
```

```
--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 XCO.

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

```
--enable { yes | no }
```

Turns the NTP server on or off.

```
--fabric name
```

Fabric name.

```
--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.

```
--list
```

Lists disable-server on devices.

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.

```
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 ]
```

```
--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. .

```
--fabric name
```

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

```
--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.

--trusted-key *key*

Specifies the trusted key.

The command creates an NTP server. By default, XCO 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, XCO enforces a unique key id per NTP server. Here, device IP addresses and fabric options are mutually exclusive.

```
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*]

--fabric *name*

Specifies the name of the fabric.

--ip *device-ip*

Specifies a comma-separated range of device IP addresses. For example:
1.1.1-3,1.1.2,2.2.2.

--ntp-ip *ntp-ip*

Specifies the IP address of the NTP server.

```
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 XCO.

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

--fabric *name*

Specifies the name of the fabric.

--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.

```
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.

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

```
--fabric fabric-name
```

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

```
--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.

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.

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.

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



Note

IPv6 is not supported in Fabric suite.

--ip *ip-list*

Specifies a comma separated range of device IP addresses. Can be either IPv4 or IPv6 IP addresses. Example: 10.10.10.1,2000::1.

--location *string*

Location to connect to the device.

--maintmode-enable-on-reboot

Configures a device to reboot in maintenance mode.

--password *password*

Specifies password to connect to the device.

--username *username*

Specifies the user name to connect to the device.

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.

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

--fabric *fabric-name*

Specifies the name of the fabric.

--ip *device-ips*

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

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.

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 secure settings

Enable or disable security settings

```
efa inventory device secure settings [ enable | disable ]
```

```
efa inventory device secure settings disable
secure settings disable[Success]
--- Time Elapsed: 57.000421492s ---
```

efa inventory device secure settings apply

Apply security hardening configuration on device

```
efa inventory device secure settings apply [ --ip device-ips | --fabric fabric ]
```

```
--fabric fabric
```

Specifies the fabric name.

```
--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.

```
efa inventory device secure settings apply --ip 1.1.1.1-3,2.2.2.2
+-----+-----+
+-----+-----+
| IP Address                                     |
| Status |                                     Warning |
+-----+-----+
+-----+-----+
| 1.1.1.1 |                                     |
Success |                                     |
+-----+-----+
+-----+-----+
| 1.1.1.2 |                                     |
Success |                                     |
+-----+-----+
+-----+-----+
| 1.1.1.3 |                                     |
Success |                                     |
+-----+-----+
+-----+-----+
| 2.2.2.2 |                                     |
Success |                                     |
+-----+-----+
+-----+-----+
Secure settings apply
--- Time Elapsed: 120.505169258s ---

efa inventory device secure settings apply -fabric fabric1
+-----+-----+
+-----+-----+
| IP Address                                     |
| Status |                                     Warning |
+-----+-----+
+-----+-----+
```

```

+-----+
| 1.1.1.6      |
Success |
+-----+
+-----+
| 1.1.1.8      |
Success |
+-----+
+-----+
Secure settings apply
--- Time Elapsed: 99.548236s ---

```

efa inventory device secure settings reset

Reset security setting to default vale.

```
efa inventory device secure settings reset [ --min-tls-version | --mac-
algorithm | --key-exchange-algorithm | --telnet | --cipher | --max-
password-age | --force-default-password-change ]
```

--cipher

Reset Ciphers to default values.

--force-default-password-change

Enable force password change.

--key-exchange-algorithm

Reset Key-Exchange Algorithms list to default values.

--mac-algorithm

Reset MAC Algorithms list to default values.

--max-password-age

Reset password age to default value.

--min-tls-version

Reset minimum TLS version to the default value.

--telnet

Disable telnet.

```

efa inventory device secure settings reset --telnet --cipher --max-password-age
secure settings reset[Success]
--- Time Elapsed: 83.826371ms ---

```

efa inventory device secure settings show

Display current security setting on device.

```
efa inventory device secure settings show [ --ip device-ips ]
```

--ip device-ips

Specifies a device IP address. Example: 10.24.80.56

```

efa inventory device secure settings show --ip 10.24.80.56
+-----+-----+
|          NAME          | VALUE          |
+-----+-----+

```

```

| Min-tls-version          | 1.2          |
+-----+-----+
| Mac-algorithm           | hmac-sha2-512-etm@openssh.com |
|                         | hmac-sha2-256-etm@openssh.com |
|                         | hmac-sha2-512          |
|                         | hmac-sha2-256          |
+-----+-----+
| Key-exchange-algorithm  | curve25519-sha256      |
|                         | curve25519-sha256@libssh.org |
|                         | diffie-hellman-group14-sha256 |
|                         | diffie-hellman-group16-sha512 |
|                         | diffie-hellman-group18-sha512 |
|                         | diffie-hellman-group-exchange-sha256 |
+-----+-----+
| Cipher                  | non-cbc              |
+-----+-----+
| Telnet                  | Disable              |
+-----+-----+
| Max-password-age        | 365                  |
| Force-default-password-change | Enable              |
+-----+-----+
Current security settings on device 10.24.80.56
--- Time Elapsed: 64.12891ms ---

```

efa inventory device secure settings update

Update security settings on devices.

```

efa inventory device secure settings update [ --min-tls-version min-tls-version | --mac-algorithm mac-algorithm-list | --key-exchange-algorithm key-exchange-algorithm-list | --telnet disable | --cipher ciphers | --max-password-age max-password-age | --force-default-password-change enable ]

```

--cipher *ciphers*

Comma separated list of Ciphers.

--force-default-password-change *enable*

Force the user to change default passwords.

--key-exchange-algorithm *key-exchange-algorithm-list*

Comma separated list of Key-Exchange Algorithms.

--mac-algorithm *mac-algorithm-list*

Comma separated list of MAC Algorithms.

--max-password-age *max-password-age*

Maximum number of days before password expiry

--min-tls-version *min-tls-version*

Minimum TLS version to be configured on the server. Example: 1.1|1.2

--telnet *disable*

Disable or Enable telnet on mgmt-vrf

This command does not automatically apply changes to these settings. Run *efa inventory device secure settings apply* to manually apply them.

```
efa inventory device secure settings update --min-tls-version 1.2
secure settings update[Success]
--- Time Elapsed: 33.26355ms ---

efa inventory device secure settings update
--mac-algorithm hmac-sha2-512-etm@openssh.com,hmac-sha2-256-etm@openssh.com,hmac-
sha2-512,hmac-sha2-256
secure settings update[Success]
--- Time Elapsed: 18.82952ms ---

efa inventory device secure settings update
--key-exchange-algorithm curve25519-sha256,curve25519-sha256@libssh.org,diffie-hellman-
group14-sha256,
diffie-hellman-group16-sha512,diffie-hellman-group18-sha512,diffie-hellman-group-exchange-
sha256
secure settings update[Success]
--- Time Elapsed: 83.826371ms ---

efa inventory device secure settings update --telnet enable --cipher non-cbc --max-
password-age 365
secure settings update[Success]
--- Time Elapsed: 55.88571ms ---
```

efa inventory device setting show

Displays device settings.

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

--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	APP STATE
Maintenance Mode Enable	On	No
Reboot		
Maintenance Mode Enable	No	
Maintenance Convergence Time		
MCT Bring-up Delay		
Health Check Enabled	No	
Health Check Interval	6m	
Health Check Heartbeat Miss Threshold	2	
Periodic Backup Enabled	Yes	
Config Backup Interval	24h	
Config Backup Count	4	


```

--urpf string | --ip-dhcp-relay-disable-flooding { Yes | No } | --
ip-option-disable { Yes | No } | --ip-option-disable-cpu { Yes |
No } | --ipv6-option-disable { Yes | No } | --peer-group-ipv6-prefix-
over-ipv4-peer { Yes | No } | --crypto-cert-expiry-info string | --
crypto-cert-expiry-minor string | --crypto-cert-expiry-major string |
--crypto-cert-expiry-critical string | --password-expiry-info string
| --password-expiry-minor string | --password-expiry-major string |
--password-expiry-critical string | -h, --help ]

--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. Example: 99m or 1h20m or 20m, the default is 24h.

--crypto-cert-expiry-info string
Specifies number of days before certificate expiry after which SNMP trap is
generated. Range is 1 to 90 days, 0 to deconfigure.

--crypto-cert-expiry-minor string
Specifies number of days before certificate expiry after which SNMP trap is
generated. Range is 1 to 90 days, 0 to deconfigure.

--crypto-cert-expiry-major string
Specifies number of days before certificate expiry after which SNMP trap is
generated. Range is 1 to 90 days, 0 to deconfigure.

--crypto-cert-expiry-critical string
Specifies number of days before certificate expiry after which SNMP trap is
generated. Range is 1 to 90 days, 0 to deconfigure.

--fabric fabric-name
Specify the name of the fabric.

--health-check-enable { Yes | No }
Specify Yes to enable health check and No to disable it.

--health-check-interval minutes
Health check interval in seconds/minutes, valid values for Fabric device 6m-24h,
valid values for NPB device 30s-24h Example. 30s or 99m or 1h20m or 20m,
default 6m for Fabric device, 30s for NPB device.

--health-check-heartbeat-miss-threshold { 2 | 3 | 4 | 5 }
Specifies the health check miss threshold. Default is 2.

--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.

--ip-dhcp-relay-disable-flooding { Yes | No }

```

Enter Yes to disable flooding for DHCP broadcast packets from CPU. Valid values are Yes/No. Default value is "". This configuration is supported from the firmware version 20.5.3.

--ip-option-disable { **Yes** | **No** }

Enter Yes to disable processing IP packets with ip options. Enter No to enable processing IP packets with ip options.

--ip-option-disable-cpu { **Yes** | **No** }

Enter Yes to disable processing IP packets with IP options destined to the CPU. Enter No to enable processing IP packets with IP options destined to the CPU.

--ipv6-option-disable { **Yes** | **No** }

Enter Yes to disable processing IPV6 packets with IP options. Enter No to enable processing IPV6 packets with IP options. This parameter is not supported on SLX9540 and SLC9640.

--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.

--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.

--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.

--maximum-ipv6-prefix-length-64 *string*

Enter Yes to configure the maximum route prefix length of 64. This configuration is applicable for Extreme 8820, 8720, 8520 hardware.

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

Specifies the route load-sharing maximum paths. Valid values are 8,16,32,64,128, and 0 to deconfigure.

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

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

--password-expiry-critical *string*

Number of days before expiry to generate critical level alert. Valid range is between 1-90 and 0 to de-configure.

--password-expiry-info *string*

Number of days before expiry to generate info level alert. Valid range is between 1-90 and 0 to de-configure.

--password-expiry-major *string*

Number of days before expiry to generate major level alert. Valid range is between 1-90 and 0 to de-configure.

--password-expiry-minor *string*

Number of days before expiry to generate minor level alert. Valid range is between 1-90 and 0 to de-configure.

--peer-group-ipv6-prefix-over-ipv4-peer { **Yes** | **No** }

Enter Yes to enable the peer group ipv6prefix over ipv4peer. Enter No to disable. This configuration is supported from firmware version 20.5.2a.

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

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

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

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

--urpf *string*

Enter Yes to configure the maximum route prefix length of 64 along with unicast reverse path forwarding. This configuration is applicable for Extreme 8820, 8720, 8520 hardware.

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: The best practice is to clear all routes after configuring PIC. You can also
reload the device.
Execute the CLI to clear : efa inventory device clear route-all --ip 10.20.24.1
```

This example enables prefix independent convergence static on the device.

```
efa inventory device setting update --ip 10.20.24.18 --prefix-independent-convergence-
static yes
+-----+-----+-----+-----+-----+
| IP ADDRESS | NAME | STATUS | VALUE | ERROR |
+-----+-----+-----+-----+-----+
| 10.20.24.18 | Static Prefix Independent | Success | Yes | |
| | Convergence Enabled | | | |
+-----+-----+-----+-----+-----+
Warning: The best practice is to clear all routes after configuring Static PIC. You can
also reload the device.
Execute the CLI to clear : efa inventory device clear route-all --ip 10.20.24.18
Execute the CLI to reload : efa inventory device reload --ip 10.20.24.18
```

This example configures the maximum route prefix length of 64.

```
efa inventory device setting update --ip 10.24.80.135 --maximum-ipv6-prefix-length-64
Warning: Configuration will not take effect until reloaded.
Execute the CLI to reload : efa inventory device reload --ip 10.24.80.135
```

This example specifies disabling ip option processing and disabling ipv6 option processing while allowing ip option processing destined to the cpu.

```
efa inventory device setting update --ip 10.10.10.1 --ip-option-disable yes
--ip-option-disable-cpu no --ipv6-option-disable yes
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.24.2.120	ip-option-disable	Success	Yes	
	ip-option-disable-cpu	Success	No	
	ipv6-option-disable	Success	Yes	

This example enables the peer group ipv6prefix over ipv4peer.

```
efa inventory device setting update --ip 10.20.246.16 --peer-group-ipv6-prefix-over-ipv4-peer yes
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.20.246.16	peer-group-ipv6-prefix-over-ipv4-peer	Success	Yes	

```
efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-info 50
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.20.24.10	Crypto Cert Expiry Info	Success	50	

```
efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-minor 30
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.20.24.10	Crypto Cert Expiry Minor	Success	30	

```
efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-major 10
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.20.24.10	Crypto Cert Expiry Major	Success	10	

```
efa inventory device setting update --ip 10.20.24.10,10.20.24.12 --crypto-cert-expiry-critical 5
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.20.24.10	Crypto Cert Expiry Critical	Success	5	

This example enables flooding for the IP DHCP relay on the specified device.

```
efa inventory device setting update --ip 10.20.246.15 --ip-dhcp-relay-disable-flooding No
```

IP ADDRESS	NAME	STATUS	VALUE	ERROR
10.20.246.15	Ip Dhcp Relay Disable Flooding	Success	No	

This example enables flooding for the IP DHCP relay on the specified devices.

```
efa inventory device setting update --ip-dhcp-relay-disable-flooding No --ip
10.20.246.15,10.20.246.16
+-----+
| IP ADDRESS | NAME | STATUS | VALUE | ERROR |
+-----+
| 10.20.246.15| Ip Dhc | Success | No | |
+-----+
| 10.20.246.16| Ip Dhc | Success | No | |
+-----+
```

This example disables flooding for the IP DHCP relay on the specified fabric.

```
efa inventory device setting update --ip-dhcp-relay-disable-flooding Yes --fabric fs
+-----+
| IP ADDRESS | NAME | STATUS | VALUE | ERROR |
+-----+
| 10.20.246.15| Ip Dhc | Success | Yes | |
+-----+
| 10.20.246.16| Ip Dhc | Success | Yes | |
+-----+
```

This example sets the password expiry values for the specified parameters and values on the specified device.

```
# efa inventory device setting update --ip 10.64.192.72 --password-expiry-info 30
--password-expiry-minor 20 --password-expiry-major 10 --password-expiry-critical 200
+-----+
| IP ADDRESS | NAME | STATUS | VALUE | ERROR |
+-----+
| 10.64.192.72 | Password Expiry For | Success | 30 | |
| Info Level |
+-----+
| Password Expiry for | Success | 20 | |
| Minor Level |
+-----+
| Password Expiry for | Success | 10 | |
| Major Level |
+-----+
| Password Expiry for | Success | 2 | |
| Critical Level |
+-----+
--- Time Elapsed: 53.849675ms ---
```

This example sets the password expiry values for the specified parameters and values on the specified fabric.

```
# efa inventory device setting update --fabric myclos --password-expiry-info 30
--password-expiry-minor 20 --password-expiry-major 10 --password-expiry-critical 200
+-----+
| IP ADDRESS | NAME | STATUS | VALUE | ERROR |
+-----+
| 10.64.192.72 | Password Expiry For | Success | 30 | |
| Info Level |
+-----+
```

```

+-----+-----+-----+-----+
| Password Expiry for | Success | 20 |
| Minor Level |
+-----+-----+-----+-----+
| Password Expiry for | Success | 10 |
| Major Level |
+-----+-----+-----+-----+
| Password Expiry for | Success | 2 |
| Critical Level |
+-----+-----+-----+-----+
--- Time Elapsed: 53.849675ms ---

```

efa inventory device snmp community create

Creates an SNMP community and SNMP group.

```

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

```

--enable-notify-access

Sets notify access for the view.

--enable-read-access

Sets read access for the view.

--enable-write-access

Sets write access for the view.

--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.

--group *group*

Specifies an SNMP group name.

--name *community*

Specifies an SNMP community name.

--view *view-name*

Optionally specify a SNMP view name. Default view `efav3View` used when not specified.

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

XCO 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.

```

efa inventory device snmp community create --ip 10.139.44.153 --name community1
--group group1 --enable-read-access --view view1

```

efa inventory device snmp community delete

Deletes an SNMP community.

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

--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.

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

XCO 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.

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

efa inventory device snmp community list

Lists SNMP communities.

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

--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.

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

+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+
| IP Address |      Community Name      | Group |
| Read view | Write view | Notify view | AppState |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| 10.20.246.2 | $9$jcpLTpu9FlqFFcWN7uhHig== |      |
|           |           | cfg-not-managed |         |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| 10.20.246.2 | $9$smklvisSgh0ZEQvXJKBDeA== | cgrp1 |
efav3View | efav3View |           | cfg-in-sync |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
```



```

| 10.20.246.1 | $9$Y4Cy6Icf1H9B0KM932JUNw== | v2group |
efav3View | efav3View | | cfg-not-managed |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.1 | $9$yS/jSs6dBjyG5o0yy+coKw== | |
| | | cfg-not-managed |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.1 | $9$kR+plantb3TRIGkiBkV83Q== | v2group
| efav3View | efav3View | | cfg-not-managed |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.1 | $9$mdDwWNPg2OmO4Fqnodl+Bw== | gtest1
| efav3View | efav3View | | cfg-not-managed |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.1 | $9$smklvisSghOZEQvXJKBDeA== | cgrp1
| efav3View | efav3View | | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
+-----+
Snmp community details

```

efa inventory device snmp group list

Lists SNMP groups

efa inventory device snmp group list [--ip *device-ip*]

--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.

This examples lists groups on a specific device.

```

efa inventory device snmp group list --ip 10.139.44.159
+-----+-----+-----+-----+-----+-----+
+-----+
| IP Address | Group Name | Version | Read view
| Write view | Notify view | Auth Level | AppState |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.139.44.159 | efav3Group | v3 |
| efav3View | noauth | cfg-not-managed |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.139.44.159 | test-grp-v2 | v2c | test-view-1
| efav3View | | noauth | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
+-----+
Snmp group details

```

efa inventory device snmp group update

Update a SNMP group

efa inventory device snmp group update [--ip *device-ips* | --name *group-name* | --version *version* | --read-view *view-name* | --write-view *view-name* | --notify-view *view-name* | --auth-level *auth-level*]

--auth-level *auth-level*

Sets group auth level. Valid values are noauth, auth and priv.

--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.

--name *group-name*

Specifies an SNMP name.

--notify-view *view-name*

Sets group notify view name.

--read-view *view-name*

Sets group read view name.

--version *version*

Specifies an SNMP group version. Valid values are v2c and v3.

--write-view *view-name*

Sets group write view name.

This example updates a group.

```
efa inventory device snmp group update --name "test-grp-v2" --version v2c --read-view
"test-view-1" --ip 10.139.44.159
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| IP Address | Group Name | Version | Read view |
| Write view | Notify view | Auth Level | Status | Reason |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.139.44.159 | test-grp-v2 | v2c | test-view-1 |
| | | | Success | |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
Snmp group details
```

efa inventory device snmp host create

Creates an SNMP v2c or v3 host.

```
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 | --udpport ]
```

--community *community*

Specifies a community name. Applicable for v2c 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.

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

Specifies a host IP address.

--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.

--notify-type { **traps** | **informs** }

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

--udpport

Optional port number used to send notifications. Range: 0-65535, Default value is 162.

--user *user*

Specifies an SNMP v3 user.

--version { **v2c** | **v3** }

Specifies the SNMP version.

The command creates an SNMP v2c or v3 host.

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

```
efa inventory device snmp host create --ip 10.139.44.153 --host-ip 1.1.1.1
--user user1 --version v3 --notify-type traps --udp-port 163
```

```
efa inventory device snmp host list --ip 10.139.44.153
```

efa inventory device snmp host delete

Deletes an SNMP v2c or v3 host.

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

--community *community*

Specifies an SNMP community. Applicable for v2c only.

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

Specifies an SNMP host address.

--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.

--user *username*

Specifies an SNMP v3 user.

```
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.

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

--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.

```
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 | | traps
| 00:00:00:00:00:00:00:00:00 | management chassis-ip | mgmt-vrf | 162
| None | cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.10 | efav3User | |
traps | 00:00:00:00:00:00:00:00:00 | management chassis-ip | mgmt-vrf
| 162 | None | cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.24 | efav3User | |
traps | 00:00:00:00:00:00:00:00:00 | management chassis-ip | mgmt-vrf
| 162 | None | cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.7 | efav3User | |
traps | 00:00:00:00:00:00:00:00:00 | management chassis-ip | mgmt-vrf
| 162 | None | cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.241.85 | efav3User | |
traps | 00:00:00:00:00:00:00:00:00 | management chassis-ip | mgmt-vrf
| 162 | None | cfg-not-managed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.2 | 10.20.255.115 | efav3User | |
traps | 00:00:00:00:00:00:00:00:00 | management chassis-ip | mgmt-vrf
| 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.

```
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 | --viewview-name ]
```

--auth-pass *passphrase*

Authentication password.

--auth-protocol { md5 | sha }

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

--enable-notify-access

Sets notify access for the view.

--enable-read-access

Sets read access for the view.

--enable-write-access

Sets write access for the view.

--group *group*

Specifies an SNMP group name.

--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.

--priv-pass *privphrase*

Privacy password.

--priv-protocol { AES128 | DES }

Privacy protocol.

--view*view-name*

Optionally specify a SNMP view name. Default view efa3View used when not specified.

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, XCO encrypts all passwords.

XCO 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.

This example creates user using a specified view.

```
efa inventory device snmp user create --ip 10.139.44.153 --name user1
--group group1 --view view1
```

efa inventory device snmp user delete

Deletes an SNMP user.

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

```
--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.

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

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

efa inventory device snmp user list

Lists SNMP users.

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

```
--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.

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

This example lists snmp users.

```
efa inventory device snmp user list --ip 10.139.44.153
```

efa inventory device snmp use-vrf create

Create a SNMP use-vrf

```
efa inventory device snmp use-vrf create [ --ip device-ips | --name VRF-  
name | --shutdown ]
```

```
--ip device-ips
```

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

```
--name VRF-name
```

VRF name.

--shutdown

Shutdown SNMP service in the VRF.

This example creates a use-vrf on specified devices.

```
efa inventory device snmp use-vrf create --name vrf1 --ip 10.139.44.159-160 --shutdown
+-----+-----+-----+-----+-----+
| IP Address | Use VRF | Shutdown | Status | Reason |
+-----+-----+-----+-----+-----+
| 10.139.44.159 | vrf1 | true | Success | |
+-----+-----+-----+-----+-----+
| 10.139.44.160 | vrf1 | true | Success | |
+-----+-----+-----+-----+-----+
Snmp use-vrf details
```

efa inventory device snmp use-vrf delete

Delete a SNMP use-vrf.

efa inventory device snmp use-vrf delete [**--ip** *device-ips* | **--name** *VRF-name*]

--ip *device-ips*

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

--name *VRF-name*

VRF name.

This example deletes a use-vrf on specified devices.

```
efa inventory device snmp use-vrf delete --name vrf1 --ip 10.139.44.159-160
+-----+-----+-----+-----+
| IP Address | User Name | Status | Reason |
+-----+-----+-----+-----+
| 10.139.44.160 | vrf1 | Success | |
+-----+-----+-----+-----+
| 10.139.44.159 | vrf1 | Success | |
+-----+-----+-----+-----+
Snmp user details
```

efa inventory device snmp use-vrf list

List SNMP use-vrfs

efa inventory device snmp use-vrf list [**--ip** *device-ips*]

--ip *device-ips*

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

This example shows the current SNMP use-vrfs for the specified devices.

```
efa inventory device snmp use-vrf list --ip 10.139.44.159-160
+-----+-----+-----+-----+
| IP Address | Use VRF | Shutdown | AppState |
+-----+-----+-----+-----+
| 10.139.44.159 | default-vrf | false | cfg-not-managed |
+-----+-----+-----+-----+
```

```

| 10.139.44.159 | mgmt-vrf | false | cfg-not-managed |
+-----+-----+-----+-----+
| 10.139.44.159 | vrf1 | true | cfg-in-sync |
+-----+-----+-----+-----+
| 10.139.44.160 | default-vrf | false | cfg-not-managed |
+-----+-----+-----+-----+
| 10.139.44.160 | mgmt-vrf | false | cfg-not-managed |
+-----+-----+-----+-----+
| 10.139.44.160 | vrf1 | true | cfg-in-sync |
+-----+-----+-----+-----+
Snmp use-vrf details

```

efa inventory device snmp view

SNMP View commands.

```
efa inventory device snmp view create [ --ip device-ips | --name view-name | --mib-tree mib-oid | --mib-tree-access access ]
```

```
efa inventory device snmp view delete [ --ip device-ips | --name view-name | --mib-tree mib-oid ]
```

```
efa inventory device snmp view list [ --ip device-ips ]
```

--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.

--name *view-name*

Specifies the name of the view.

--mib-tree *mib-oid*

Specifies the MIB subtree in the form of Object identifier. Example: 1.3.6.1

--mib-tree-access *access*

Specifies the mib-tree access. Valid values are: included, excluded.

This example creates a view on a specified device.

```

efa inventory device snmp view create --ip 10.139.44.153-154 --name view1
--mib-tree 1.3.6.1 --mib-tree-access included

+-----+-----+-----+-----+-----+-----+
| IP Address | Name | MIB-Tree | MIB-Tree-Access | Status | Reason |
+-----+-----+-----+-----+-----+-----+
| 10.139.44.153 | view1 | 1.3.6.1 | included | Success | |
+-----+-----+-----+-----+-----+-----+
| 10.139.44.154 | view1 | 1.3.6.1 | included | Success | |
+-----+-----+-----+-----+-----+-----+
Snmp view details

```

This example deletes a view on a specified device.

```

efa inventory device snmp view delete --ip 10.139.44.153-154 --name view1
--mib-tree 1.3.6.1

+-----+-----+-----+-----+-----+
| IP Address | Name | MIB-Tree | Status | Reason |
+-----+-----+-----+-----+-----+

```



```

| 10.139.44.153 | view1 | 1.3.6.1 | Success |      |
+-----+-----+-----+-----+-----+
| 10.139.44.154 | view1 | 1.3.6.1 | Success |      |
+-----+-----+-----+-----+-----+
Snmp view details

```

This example shows the current SNMP view for the specified device

```

efa inventory device snmp view list --ip 10.139.44.153-154

+-----+-----+-----+-----+-----+
| IP Address | Name | MIB-Tree | MIB-Tree_Access | AppState |
+-----+-----+-----+-----+-----+
| 10.139.44.153 | view1 | 1.2.3.4 | included |      |
+-----+-----+-----+-----+-----+
|                | view2 | 1.2.3.4 | included |      |
+-----+-----+-----+-----+-----+
| 10.139.44.154 | view3 | 1.2.3.5 | excluded |      |
+-----+-----+-----+-----+-----+
Snmp view details

```

efa inventory device timezone debug-show

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

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

efa inventory device timezone list

Lists timezones configured on an SLX device.

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

--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.

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

--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.

```
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.

efa inventory device timezone unset [**--ip** *device-ip* | **--fabric** *fabric-name*]

--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

```
efa inventory device timezone unset --ip 10.20.246.13
```

efa inventory device threshold-monitor list

Display the threshold monitor settings for a device.

efa inventory device threshold-monitor list [**--ip** *device-ips* | **--fabric** *fabric-name*]

--fabric *fabric-name*

Specifies the name of the fabric.

--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.

This example shows the current inventory threshold-monitor settings for the specified device.

```
efa inventory device threshold-monitor list --ip 10.10.10.75
```

```
Hardware Resource Configuration
+-----+-----+-----+
| IP Address | Type | AppState |
+-----+-----+-----+
| 10.10.10.75 | hardware-resources | cfg-in-sync |
+-----+-----+-----+
```

```
Basic Monitoring Type Configuration
```

```

+-----+-----+-----+-----+-----+-----+
| IP Address | Type | Actions | High Limit | Low Limit | AppState |
+-----+-----+-----+-----+-----+-----+
| 10.10.10.75 | bfd-session | snmp | 80 | 50 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | vxlan-tunnel | | 80 | | | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | lif | snmp | 80 | 50 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | mac-table | snmp | 80 | 50 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | route | snmp | 80 | 50 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | host | | 80 | | | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | nexthop | snmp | 80 | 50 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | ecmp | snmp | 80 | 50 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+

```

CPU/Memory Type Configuration

```

+-----+-----+-----+-----+-----+-----+
| IP Address | Type | Actions | High Limit | Retry | AppState |
+-----+-----+-----+-----+-----+-----+
| 10.10.10.75 | cpu | all | 80 | 3 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
| | memory | raslog | 80 | 3 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+

```

Threshold Monitor details

```
--- Time Elapsed: 55.652194ms ---
```

efa inventory device threshold-monitor set

Set threshold monitor settings.

```

efa inventory device threshold-monitor set [ --ip device-ips | --fabric
fabric-name | --type { cpu | memory | bfd-session | lif | mac-table
| vxlan-tunnel | route | host | nexthop | ecmp | hardware-resources }
| --actions { all | none | raslog | snmp | loginfo } | --high-limit
percentage | --low-limit percentage | --interval seconds | --retry
retry-interval | --count max-events | -h, --help ]

```



Note

--count is valid only when --type of hardware-resources is specified.

--actions *actions*

Specifies the action to take upon exceeding the high-limit or low-limit. Valid actions are all, none, raslog, snmp, and loginfo. Loginfo is only valid for type cpu or memory. [all] Both RASLOG and SNMP trap will be generated. [none] No action will be taken. [raslog] RASLOG will be generated. [snmp] SNMP trap will be generated. [loginfo] Diagnostic data collection along with RASLOG.

--count *max-events*

Specifies the maximum number of event actions that can be taken in an interval. See `interval`. A valid value is in the range of 1-60 with Default 4. Valid only for type `hardware-resources`, which sets the common/shared value for all hardware resource types of `bfd-session`, `lif`, `mac-table`, `vxlan-tunnel`, `route`, `host`, `nexthop` and `ecmp`. (Default -1)

--fabric *fabric-name*

Specifies the name of the fabric.

--high-limit *percentage*

Specifies the high limit threshold percentage. Valid values for type `cpu` is 0-95 with Default 50. Valid values for type `memory` is 0-95 with Default 70. Valid values for types `bfd-session/lif/mac-table/vxlan-tunnel` is 10-100 with Default 90.

--interval *seconds*

For type `cpu/memory` it specifies the polling interval, in seconds after which a sample will be taken. A valid value is in the range of 10-3600 seconds with Default 120 seconds.

For type `hardware-resources`, which sets the common/shared value for all hardware resource types of `bfd-session`, `lif`, `mac-table`, `vxlan-tunnel`, `route`, `host`, `nexthop` and `ecmp`, it specifies the number of notifications within an interval. A valid value is in the range of 30 to 900 seconds with Default 30 seconds.

--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.

--low-limit *percentage*

Specifies the low limit threshold percentage. Valid for types `bfd-session`, `lif`, `mac-table`, `vxlan-tunnel`, `route`, `host`, `nexthop`, and `ecmp`. Valid values are in the range of 10-99 with Default 70.

--retry *retry-interval*

Specifies the number of polling interval retries that will be attempted (see `--interval`) before desired action is taken. Only applicable for `--type` `cpu` or `memory`. Valid values are 1-100 with default 3.

--type *type*

Specifies the type of threshold being configured. Valid types are `cpu`, `memory`, `bfd-session`, `lif`, `mac-table`, `vxlan-tunnel`, `route`, `host`, `nexthop`, `ecmp`, and `hardware-resources`. Default is `all`.

This example sets a `cpu` threshold-monitor to trigger an `snmp` event when the `cpu` threshold of 80% is exceeded. The threshold must be exceeded 3 times (`retry`) for the event to be triggered, which would take 3 minutes before the notification was sent.

```
efa inventory device threshold-monitor set --type cpu --ip 10.10.10.153-154
--actions snmp --high-limit 80 --retry 3
```

This example sets a mac-table threshold-monitor to trigger an snmp event when the high-limit threshold of 80% is exceeded or the low-limit threshold of 50% is exceeded. Events can be for the threshold exceeding 80% (high) or 50% (low) limits.

```
efa inventory device threshold-monitor set --type mac-table --fabric fabric1
--actions snmp --high-limit 80 --low-limit 50
```

This example sets the threshold-monitor hardware-resources interval and count to 60 and 3 respectively. These values apply to the types bfd-session, lif, mac-table, and vxlan-tunnel, route, host, nexthop and ecmp.

```
efa inventory device threshold-monitor set --type hardware-resources
--ip 10.10.10.153-154 --interval 60 --count 3
```

efa inventory device threshold-monitor unset

Unset threshold monitor settings

```
efa inventory device threshold-monitor unset [ --ip device-ips | --fabric
fabric-name | --type { cpu | memory | bfd-session | lif | mac-table |
vxlan-tunnel | route | host | nexthop | ecmp | hardware-resources ]
```

--fabric *fabric-name*

Specifies the name of the fabric.

--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.

--type *type*

Specifies the type of threshold being configured. Valid types are cpu, memory, bfd-session, lif, mac-table, vxlan-tunnel, route, host, nexthop, ecmp, and hardware-resources.

This example unsets all current inventory threshold-monitor settings for the specified devices.

```
efa inventory device threshold-monitor unset --ip 10.10.10.75-76
```

This example unsets a single inventory threshold-monitor setting for the specified device.

```
efa inventory device threshold-monitor unset --ip 10.10.10.75 --type cpu
```

efa inventory device tpvm list

Lists all the known tpvm nodes.

```
efa inventory device tpvm list
```

This example shows the output from the command.

```
efa inventory device tpvm list
-----+-----
+-----
Device IP Address | TPVM IP Address | TPVM Hostname
```

```

| SLX Firmware Version | TPVM Version |
+-----+-----+-----+-----+-----+
+-----+
| 10.24.80.56 | 10.24.80.180 | node180 | 20.4.2slxos20.4.2_220614_1000 | 4.5.0 |
+-----+-----+-----+-----+-----+
+-----+
| 10.24.80.58 | 10.24.80.181 | node181 | 20.4.1 | 4.5.0 |
+-----+-----+-----+-----+-----+
+-----+

```

efa inventory device tpvm-upgrade execute

Asynchronously launches the TPVM upgrade process.

```

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 | --auto-reboot ]

```

--auto-reboot

If this option is given then automatically both standby and active TPVMs will be rebooted sequentially (if required). Default is false.

--firmware-host *firmware-host-ip-address*

Specifies the IP address of the firmware host. The host must be registered beforehand.

--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.

--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-password *sudo-user-password*

Specifies the password of the trusted peer sudo user to be reconfigured after TPVM is upgraded.

--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*.

Upgrade one TPVM node:

```

efa inventory device tpvm-upgrade execute --ip 10.24.80.58
--firmware-host 10.31.2.101
--tpvm-image /buildsjc/sre_fusion/Nightly/raphael/slxos20.4.2/slxos20.4.2_220511_1000/
dist/SWBD2900/tpvm_inc_upg-4.5.0-4.amd64.deb

```

Upgrade two TPVM nodes:

```

efa inventory device tpvm-upgrade execute --ip 10.24.80.58,10.20.80.56
--firmware-host 10.31.2.101
--tpvm-image /buildsjc/sre_fusion/Nightly/raphael/slxos20.4.2/slxos20.4.2_220511_1000/
dist/SWBD2900/tpvm_inc_upg-4.5.0-4.amd64.deb

```

This example shows the tpvm upgrade configuration with auto-reboot option.

```
$ efa inventory device tpvm-upgrade execute --ip 1.1.1.1,1.1.1.2 --firmware-host 2.2.2.2
--tpvm-image /buildsjc/sre_fusion/Nightly/tpvm/ci_tpvm/tpvm_inc_upg-4.5.5-5.amd64.deb --
auto-reboot

TPVM Upgrade Execute [success]
Monitor TPVM upgrade execution progress using:

efa inventory device tpvm-upgrade show --ip 1.1.1.1,1.1.1.2
efa inventory device tpvm-upgrade show --execution-id
aac8b30e-6911-4bea-8f67-986bd40528a9

Please do not execute other commands on the device until process is completed
```

efa inventory device tpvm-upgrade show

Polls and monitors the progress of the TPVM upgrade workflow.

```
efa inventory device tpvm-upgrade show [ --execution-id id | --ip ip-
addresses ]
```

--execution-id *id*

Specifies the ID provided from the **tpvm-upgrade execute** command.

--ip *ip-addresses*

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

Specify either the **--execution-id** or the **--ip** option.

TPVM upgrade show with one TPVM IP address:

This example shows only a portion of the possible output.

```
$ efa inventory device tpvm-upgrade show --ip 10.24.80.58
+-----+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model | Chassis Name | ASN          | Role   |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX      | 4001 | BR-SLX9640  | 4200000000  | Leaf  |
+-----+-----+-----+-----+-----+-----+
```

TPVM upgrade show with two TPVM IP addresses:

This example shows only a portion of the possible output.

```
$ efa inventory device tpvm-upgrade show --ip 10.24.80.58,10.24.80.56
+-----+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model | Chassis Name | ASN          | Role   |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX      | 4001 | BR-SLX9640  | 4200000000  | Leaf  |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.56 | SLX      | 4001 | BR-SLX9640  | 4200000000  | Leaf  |
+-----+-----+-----+-----+-----+-----+
```

TPVM upgrade show with execution id:

This example shows only a portion of the possible output.

```
$ efa inventory device tpvm-upgrade show --execution-id 670cb89e-d8d1-4213-
ac97-20403458627f
+-----+-----+-----+-----+-----+-----+
| IP Address | Host Name | Model | Chassis Name | ASN | Role |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.58 | SLX | 4001 | BR-SLX9640 | 4200000000 | Leaf |
+-----+-----+-----+-----+-----+-----+
| 10.24.80.56 | SLX | 4001 | BR-SLX9640 | 4200000000 | Leaf |
+-----+-----+-----+-----+-----+-----+
```

efa inventory device update

Updates the Asset service with the latest information from the switch.

```
efa inventory device update [ --ip device-ip | --fabric fabric-name |
--username username | --password password ]
```

--fabric *fabric-name*

Specifies the fabric of the devices that you want to synchronize.

--ip *device-ip*

Specifies the IP address of the device that you want to synchronize.

--password *password*

Specifies the password to connect to the device. If the password must be updated, use this option.

--username *username*

Specifies the user name to connect to the device. If the user name must be updated, use this option.

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.

```
efa inventory drift-reconcile history [ --ip string | --reverse ]
```

```
efa inventory drift-reconcile execute [ --ip string | --reconcile ]
```

```
efa inventory drift-reconcile delete [ --key { ip-address | dr-uuid } ]
```

```
efa inventory drift-reconcile detail [ --uuid dr-uuid ]
```

--ip *string*

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

--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.

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.

The following examples show **efa inventory drift-reconcile** commands.

```
# efa inventory drift-reconcile execute --ip 10.24.14.133 --reconcile
# efa inventory drift-reconcile execute --ip 10.24.14.133-134,10.24.14.140 --reconcile
# efa inventory drift-reconcile history --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.

efa inventory execution show [**--id** *id* | **--limit** *value* | **--status** { **failed** | **succeeded** | **all** }]

--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.

```
efa inventory firmware-host delete [ --ip host-ip ]
```

```
--ip host-ip
```

Specifies the IP address of the firmware host.

efa inventory firmware-host list

Displays all registered firmware hosts.

```
efa inventory firmware-host list [ --ip host-ips ]
```

```
--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.

```
efa inventory firmware-host register [--ip host-ip | --protocol  
  { scp,ftp,sftp,http } | --username username | --password password ]
```

```
--ip host-ip
```

Specifies the IP address of the firmware host.

```
--protocol { scp,ftp,sftp,http }
```

Specifies the comma-separated protocols used by the firmware host for downloading 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.

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.

```
efa inventory firmware-host update [ --ip host-ip | --protocol  
  { scp,ftp,sftp,http } | --username username | --password password ]
```

```
--ip host-ip
```

Specifies the IP address of the firmware host.

```
--protocol { scp,ftp,sftp,http }
```

Specifies the comma-separated protocols used by the firmware host for downloading 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.

efa inventory kvstore create [**--key** *name* | **--value** *value* | **--encrypt**]

efa inventory kvstore delete [**--key** *name*]

efa inventory kvstore list [**--decrypt** | **--key** *name* | **--prefix** *prefix*]

--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).

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*]

--config-backup-id *uuid*

Specifies the UUID of the configuration to be replayed.

--ip *ip-address*

Specifies the IP address of the device.

--key *ip-address*

Specifies IP address of the device or RMA UUID to be deleted.

--uuid *uuid*

Specifies the UUID of the drift reconcile.

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.

The following example shows **efa inventory rma** commands.

```
# efa inventory rma execute --ip 10.24.14.133
--config-backup-id 1111-1111-111
# efa inventory rma history -ip 10.24.14.133
# efa inventory rma detail -uuid 123e4567-e89b-12d3-a456-426614174000
```

efa license add

Add license.

```
efa license add [ --filepath string ]
```

```
--filepath string
```

Specifies the name of the license file with path.

```
# efa license add --filepath=lservlcXIQSE-1111111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic
+-----+-----+-----+-----+-----+-----+
| License File | License Type | License Features | Start Date | Expiration Date | AID |
+-----+-----+-----+-----+-----+-----+
| lservlcXIQSE-1111111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1 | Normal Standalone | PRD-XCO-EFA-D | at 0:00 hrs (UTC) on Oct 12, 2023 | at 23:59 hrs (UTC) on Oct 11, 2024 | 5272f34d-elfd-488c-bf6c-7e2b2d831db1
+-----+-----+-----+-----+-----+-----+
License details
--- Time Elapsed: 108.919628ms ---
```

efa license delete

Delete license

```
efa license delete [ --filename string ]
```

```
--filename string
```

Specifies the name of the license file to delete.

This example deletes the specified license file.

```
# efa license delete --filename=lservlcXIQSE-1111111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic
License file 'lservlcXIQSE-1111111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1.lic' deleted successfully.
--- Time Elapsed: 53.849675ms ---
```

efa license show

Show information for all valid XCO licenses.

efa license show

```
# efa license show
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| License File                               | License Type | License Features | Start
Date                                         | Expiration Date | AID                                     |
+-----+-----+-----+-----+-----+
| lservlcXIQSE-11111111111111111111111111111111103020233-16AMPRD-XCO-EFA-D1 | Normal Standalone | PRD-XCO-EFA-D | at 0:00 hrs (UTC) on Oct
12, 2023 | at 23:59 hrs (UTC) on Oct 11, 2024 | 5272F34d-e1fd-488c-bf6c-7e2b2d831db1 |
+-----+-----+-----+-----+-----+
License details
--- Time Elapsed: 108.919628ms ---
```

efa login

Logs into the XCO application.

efa login [**--username** | **--password**]

--password

Specifies password of the user.

--username

Specifies name of the user to login.

efa logout

Logs out of the XCO application.

efa logout

efa mgmt route create

Creates a virtual management route in a multi-node deployment for the Multiple Management IP Networks feature.

efa mgmt route create [**--src** *mmip-vip* | **--to** *dest-cidr* | **--via** *next-hop-ip*]

--src *mmip-vip*

Identifies the source IP address for the route. Example: 192.168.34.40 or 2000::1.

--to *dest-cidr*

Identifies the destination network for outbound traffic, as CIDR. Example: 10.20.0.0/16 or 3000::/64.

--via *next-hop-ip*

Identifies the next-hop or gateway IP address through which access to the destination network is provided. Example: 192.168.34.45 or 2000::2.

The following example creates an IPv4 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
```

The following example creates an IPv6 virtual route.

```
efa mgmt route create --src=2000::1 --via=2000::2 --to=4000::/64
```

efa mgmt route delete

Deletes a virtual management route in a multi-node deployment for the Multiple Management IP Networks feature.

```
efa mgmt route delete [ --src mmip-vip | --to dest-cidr | --via next-hop-ip ]
```

--src *mmip-vip*

Identifies the source IP address for the route. Example: 192.168.34.40 or 2000::1.

--to *dest-cidr*

Identifies the destination network for outbound traffic, as CIDR. Example: 10.20.0.0/16 or 3000::/64.

--via *next-hop-ip*

Identifies the next-hop or gateway IP address through which access to the destination network is provided. Example: 192.168.34.45 or 2000::2.

The following example deletes a virtual IPv4 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
```

The following example deletes an IPv6 virtual route.

```
efa mgmt route delete --src=2000::1 --via=2000::2 --to=4000::/64
```

efa mgmt route show

Displays a list of all virtual management routes in a multi-node deployment for the Multiple Management IP Networks feature.

```
efa mgmt route show
```

```
efa mgmt route show ipv4
```

```
efa mgmt route show ipv6
```

This example lists both IPv4 and IPv6 routes.

```
$ efa mgmt route show
+-----+-----+-----+
| Route-Src | Route-To | Route-Via |
+-----+-----+-----+
| 2000::1   | ffee::/64 | 2000::2   |
+-----+-----+-----+
| 2000::1   | 4000: :/64 | 2000::3   |
```

```
+-----+
| 10.10.10.1 | 1.1.1.0/24 | 10.10.10.2 |
+-----+
```

This example lists only IPv4 routes.

```
$ efa mgmt route show ipv4
+-----+
| Route-Src | Route-To | Route-Via |
+-----+
| 10.10.10.1 | 11.11.11.0/24 | 10.10.10.2 |
+-----+
| 10.10.10.1 | 12.12.12.0/24 | 10.10.10.2 |
+-----+
```

This example lists only IPv6 routes.

```
$ efa mgmt route show ipv6
+-----+
| Route-Src | Route-To | Route-Via |
+-----+
| 3000::1 | 4000::/64 | 3000::2 |
+-----+
| 3000::1 | 5000::/64 | 3000::2 |
+-----+
```

efa mgmt subinterface create

Creates a subinterface in a Multiple Management IP network deployment.

```
efa mgmt subinterface create [ --name sub | --vlan-id vlan-id | --ip-  
address ip-addr | --ipv6-address ipv6-addr ]
```

--ip-address *ip-addr*

Specifies the IP address of the subinterface, including the subnet mask. For example: 10.24.80.150/24.

--ipv6-address *ipv6-addr*

Specifies the IPv6 address of the subinterface, including the subnet mask. For example: 3000::1/64.

--name *sub*

Specifies the subinterface name.

--vlan-id *vlan-id*

Specifies the ID of the VLAN for the subinterface.

```
efa mgmt subinterface create --name=sub4 --vlan-id=400  
--ip-address=14.14.14.1/24 --ipv6-address=3000::1/64
```

efa mgmt subinterface delete

Deletes a specified subinterface in a Multiple Management IP network deployment.

```
efa mgmt subinterface delete --name sub
```

--name *sub*

Specifies the name of the subinterface that you want to delete.

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.

```
efa mgmt subinterface show [ --name name ]
```

--name *name*

Specifies the name of the subinterface for which you want to see details.

The following example shows the details for all configured subinterfaces.

```
+-----+-----+-----+-----+-----+
| Name | Parent Interface | Vlan | IP Subnet | IPv6 Subnet |
+-----+-----+-----+-----+-----+
| sub1 | ens160          | 100 | 10.10.10.1/24 |              |
+-----+-----+-----+-----+-----+
| sub2 | ens160          | 200 | 11.11.11.1/24 | 2000::1/64   |
+-----+-----+-----+-----+-----+
| sub3 | ens160          | 300 | 13.13.13.1/24 |              |
+-----+-----+-----+-----+-----+
| sub4 | ens160          | 400 | 14.14.14.1/24 | 3000::1/64   |
+-----+-----+-----+-----+-----+
```

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

Set static IP addresses for a subinterface.

For IPv4 Stack

```
efa mgmt subinterface staticip add [ --subinterface sub | --ip1 ipv4-  
address-1 | --ip2 ipv4-address-2 ]
```

For IPv6 Stack

```
efa mgmt subinterface staticip add [ --subinterface sub | --ip1 ipv6-  
address-1 | --ip2 ipv6-address-2 ]
```

For Dual Stack

```
efa mgmt subinterface staticip add [ --subinterface sub | --ip1 ipv4-  
address-1 | --ip2 ipv4-address-2 | --ip3 ipv6-address-1 | --ip4 ipv6-  
address-2 ]
```

```
--ip1 ipv4-address-1
```


Required. First static IPv4 address. Example: 192.168.10.1/32.

--ip2 *ipv4-address-2*

Optional second static IPv4 address. Example: 192.168.10.2/32.

--ip3 *ipv6-address-1*

Optional first static IPv6 address. Example: 2000::1/64

--ip4 *ipv6-address-2*

Optional second static IPv6 address. Example: 2000::2/64

--subinterface *sub*

Name of the sub-interface for which you want to set static IP addresses.

You can assign a maximum of one pair of static IP addresses. Only one subinterface at a time can have static IP addresses. To assign static IPv4/IPv6 address, subinterface must already have IPv4/IPv6 address configured.

```
efa mgmt subinterface staticip add -subinterface sub1
--ip1 10.10.10.1/24 --ip2 10.10.10.2/24
```

```
efa mgmt subinterface staticip add -subinterface sub2
--ip1 2000::1/64 --ip2 2000::2/64
```

```
efa mgmt subinterface staticip add -subinterface sub3
--ip1 10.10.10.1/24 --ip2 10.10.10.2/24 -ip3 2000::1/64 --ip4 2000::2/64
```

efa mgmt subinterface staticip remove

Removes the IPs that are attached to the specified subinterface.

efa mgmt subinterface staticip remove [**--subinterface** *sub*]

--subinterface *sub*

Specifies the subinterface for which to remove IPs.

```
(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.

efa mgmt subinterface staticip show

```
(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.

```
efa notification subscribers add-https [ --url address | --username name
  | --password password | --insecure | --cacert string | --filter
  strings | --conn-timeout int | --device-event strings | --minimum-
severity string ]
```

--cacert *string*

(Optional) Local path to the cacert pem file for SSL verification.

--conn-timeout *int*

Timeout to open a connection to the server (default 10). (Optional)

--device-event *strings*

Comma separated filter values of allowed device event notifications. Possible values are "raslog", "audit-configuration", "audit-firmware" and "audit-security". If no sub-filters are provided it is equivalent to including all message types, e.g. --device-event raslog,audit-configuration,audit-firmware,audit-security. (Optional)

--filter *strings*

(Optional) Comma separated filter values. Possible values are "DEVICE_EVENTS" - RAS/syslog events from devices, "APP_ALARMS" - fault alarms from application, "APP_ALERTS" - fault alerts from application, "APP_EVENTS" - task events from application. If no filters are provided it means all types. Example: --filter DEVICE_EVENTS,APP_ALARMS,APP_ALERTS,APP_EVENTS.

--insecure

(Optional) Indicates that insecure SSL connection and transfers are used for sending notifications. By default, the SSL connection and transfers are secure.

--minimum-severity *string*

Value of the minimum severity level value for all previously unfiltered DEVICE_EVENTS, APP_ALARMS, APP_ALERTS, and APP_EVENTS. Possible values are "critical", "error", "warning" and "info". If "info" or no sub-filter is provided then no filtering is done. Input of a higher level severity such as "critical", "error" or "warning" results in all messages of lower severity to be filtered out. (Optional)

--password *password*

(Optional) Specifies the password for the user name.

--url *address*

(Required) Specifies the URL of the subscriber. Can be either IPv4 or IPv6 address.

--username *name*

(Required) Specifies the user name for access to the URL.

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 ---
```

The following example enables webhook subscriber with only APP_ALERTS.

```
#efa notification subscribers add-https --url https://127.0.0.1:5000 --username jarvis
--password vision --insecure --filter APP_ALERTS
Successfully registered subscriber.
```

attribute	value
id	18
handler	http
endpoint	https://127.0.0.1:5000
config	{"cacert":"","filters":["APP_ALERTS"],"insecure":true,"password":"vision","usern ame":"jarvis"}

```
Notification Subscriber ID=18
--- Time Elapsed: 2.148580068s ---
```

The following example enables webhook subscriber with only APP_ALARM and APP_ALERTS.

```
#efa notification subscribers add-https --url https://127.0.0.1:5000 --username jarvis
--password vision --insecure --filter APP_ALARMS,APP_ALERTS
Successfully registered subscriber.
```

attribute	value
-----------	-------

```

| attribute |
value                                             |
+-----+
+-----+
| id       |
18                                             |
+-----+
+-----+
| handler  |
| http    |
+-----+
+-----+
| endpoint |
| https://127.0.0.1:5000                      |
+-----+
+-----+
| config   |
| {"cacert":"","filters":["APP_ALARMS","APP_ALERTS"],"insecure":true,          |
|                                     |
| `password":"vision","username":"jarvis"}                                         |
+-----+
+-----+
Notification Subscriber ID=18
--- Time Elapsed: 2.148580068s ---

```

The following example enables webhook subscriber with only DEVICE_EVENTS of type audit.

```

#efa notification subscribers add-https --url https://127.0.0.1:5000 -username jarvis
--password vision --insecure --filter DEVICE_EVENTS
--device-event audit-firmware,audit-security,audit-configuration
Successfully registered subscriber.

+-----+
+-----+
| attribute |
value                                             |
+-----+
+-----+
| id       |
18                                             |
+-----+
+-----+
| handler  |
| http    |
+-----+
+-----+
| endpoint |
| https://127.0.0.1:5000                      |
+-----+
+-----+
| config   |
| {"cacert":"","filters":["DEVICE_EVENTS"],                                     |
|                                     |
|                                     | "device-event":["audit-firmware",audit-
security","audit-configuration"],          |
|                                     | "insecure":true,"minimum-
severity":"","password":"vision",          |
|                                     |
| "username":"jarvis"}                                         |
+-----+
+-----+
Notification Subscriber ID=18
--- Time Elapsed: 2.148580069s ---

```

efa notification subscribers add-syslog-relp

Registers a syslog subscriber to receive notifications over Reliable Event Logging Protocol (RELP) from the XCO Notification service.

```
efa notification subscribers add-syslog-relp [ --address host:port ]
  [ --insecure ] [ --cacert local-path ] [ --conn-timeout seconds |
  --filter strings | --rfc5424 | --device-event strings | --minimum-
  severity string ]
```

By default, subscribers do not receive syslog notifications over RELP.

--address *host:port*

(Required) Specifies the address of the syslog server in *host:port* format. The default port is 514. Syslog server address can be IPv4 or IPv6.

--cacert *local-path*

(Optional) Specifies the local path to the *ca-cert.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.

--device-event *strings*

Comma separated filter values of allowed device event notifications. Possible values are "raslog", "audit-configuration", "audit-firmware" and "audit-security". If no sub-filters are provided it is equivalent to including all message types, e.g. **--device-event** raslog,audit-configuration,audit-firmware,audit-security. (Optional)

--filter *strings*

(Optional) Comma separated filter values. Possible values are "DEVICE_EVENTS" - RAS/auditlog events from devices, "APP_ALARMS" - fault alarms from application, "APP_ALERTS" - fault alerts from application, "APP_EVENTS" - task events from application. If no filters are provided it means all types. Example: **--filters** DEVICE_EVENTS,APP_ALARMS,APP_ALERTS,APP_EVENTS.

--insecure

(Optional) Indicates that insecure SSL connection and transfers are used for sending notifications. By default, the SSL connection and transfers are secure.

--minimum-severity *string*

Value of the minimum severity level value for all previously unfiltered DEVICE_EVENTS, APP_ALERTS, APP_EVENTS, and APP_ALARMS. Possible values are "critical", "error", "warning" and "info". If "info" or no sub-filter is provided then no filtering is done. Input of a higher level severity such as "critical", "error" or "warning" results in all messages of lower severity to be filtered out. (Optional)

--rfc5424

(Optional) Enable RFC5424 message format for syslog subscribers. (Default: non-RFC5424 format)

Any external server that is configured with RELP can be registered as a subscriber to XCO notifications. For more information, see "Notification Service" in the *ExtremeCloud Orchestrator Administration Guide*.

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 ---
```

The following example enables rsyslog subscriber with only APP_ALERTS and DEVICE_EVENTS.

```
#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure
--filter APP_ALERTS,DEVICE_EVENTS
Successfully registered subscriber.

+-----+-----+
| attribute | value |
+-----+-----+
| id        | 19    |
+-----+-----+
| handler   | relp  |
+-----+-----+
| endpoint  | 127.0.0.1:1601 |
+-----+-----+
| config    | {"cacert":"","conn-timeout":10,"filters":["APP_ALERTS","DEVICE_EVENTS"],"insecure":true}|
+-----+-----+
Notification Subscriber ID=19
--- Time Elapsed: 2.172557257s ---
```

The following example enables all notification types on rsyslog subscriber.

```
#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure
Successfully registered subscriber.
```

```

+-----+
| attribute | value |
+-----+
| id        | 20   |
+-----+
| handler   | relp |
+-----+
| endpoint  | 127.0.0.1:1601 |
+-----+
| config    | {"cacert":"","conn-timeout":10,"filters":[],"insecure":true} |
+-----+
Notification Subscriber ID=20
--- Time Elapsed: 2.042797885s ---

```

The following example enables RFC-5424 format.

```

#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure --
rfc5424
Successfully registered subscriber.

+-----+
+
| attribute | value |
|
+-----+
+
| id        | 7     |
|
+-----+
+
| handler   | relp  |
|
+-----+
+
| endpoint  | 134.141.21.190:1601 |
|
+-----+
+
| config    | {"cacert":"","conn-timeout":10,"filters":[],"insecure":true,"rfc5424":true} |
|
+-----+
+
Notification Subscriber ID=7

```

The following example enables rsyslog subscriber with only APP_ALARMS, APP_ALERTS, and DEVICE_EVENTS.

```

#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure
--filter APP_ALARMS,APP_ALERTS,DEVICE_EVENTS
Successfully registered subscriber.

+-----+
+-----+
| attribute | value |
+-----+
+-----+
| id        |      |
19 |
+-----+
+
| handler   |      |
| relp      |      |
+-----+
+-----+

```

```

| endpoint
| 127.0.0.1:1601
+-----+
+-----+
| config
| {"cacert":"","conn-timeout":10,"filters":["APP_ALARMS","APP_ALERTS",
|
| "DEVICE_EVENTS"], "insecure":true}
+-----+
+-----+
Notification Subscriber ID=19
--- Time Elapsed: 2.172557257s ---

```

The following example enables rsyslog subscriber with only notifications of warning or higher of APP_ALERTS and DEVICE_EVENTS, with DEVICE_EVENTS only including audit-security and audit-configuration.

```

#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601 --insecure
--filter APP_ALERTS,DEVICE_EVENTS --device-event audit-security,audit-configuration
--minimum-severity warning
Successfully registered subscriber.

+-----+
+-----+
| attribute |
value
+-----+
+-----+
| id      |
19
+-----+
+-----+
| handler
| relp
+-----+
+-----+
| endpoint
| 127.0.0.1:1601
+-----+
+-----+
| config
| {"cacert":"","conn-timeout":10,"filters":["APP_ALERTS","DEVICE_EVENTS"],
|
| "device-event":["audit-security","audit-
configuration"],
|
| } "minimum-
severity","warning","insecure":true}
+-----+
+-----+
Notification Subscriber ID=19
--- Time Elapsed: 2.172557260s ---

```

The following example enables all notification types on rsyslog subscriber of severity error or higher.

```

#efa notification subscribers add-syslog-relp --address 127.0.0.1:1601
--insecure --minimum-severity error
Successfully registered subscriber.

+-----+
+-----+
| attribute |
value
+-----+
+-----+
| id      |

```



```

20
+-----+
| handler
| relp
+-----+
+-----+
| endpoint
| 127.0.0.1:1601
+-----+
+-----+
| config
| {"cacert":"","conn-timeout":10,"filters":[],"minimum-severity","error",
|
|
"device-event":[],"insecure":true}
+-----+
+-----+
Notification Subscriber ID=20
--- Time Elapsed: 2.042797881s ---

```

efa notification subscribers delete

Deletes the specified subscriber from the Notification service.

efa notification subscribers delete *id*

id

Specifies the ID of the subscriber that you want to delete.

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.

efa notification subscribers get *id*

id

Specifies the ID of the subscriber that you want to review.

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.

efa notification subscribers list

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 policy community-list create

Create a community-list.

```
efa policy community-list create [ --name string | --type { standard | extended } | --rule stringArray ]
```

--name *string*

Specifies the name of the community list.

--rule *stringArray*

Specifies the rule in format seq[seq-num], action[permit/deny], std-value[{ **1-4294967295** } | (**AA:NN, AA and NN is 2 bytes**) | **internet** | **local-as** | **no-export** | **no-advertise** }] | ext-value[regular expression]

--type { **standard** | **extended** }

Specifies the type of community list. Valid options are standard, extended.

This example creates a community-list of type standard with std-value.

```
efa policy community-list create --name comm1 --type standard
--rule "seq[5],action[permit],std-value[100;11:22;local-as;no-export]"
```

This example creates a community list of type extended with ext-value.

```
efa policy community-list create --name stdext1 --type extended
--rule "seq[5],action[permit],ext-value[_2000_]"
```

efa policy community-list update

Update a community-list.

```
efa policy community-list update [ --name string | --type { standard | extended } | --rule stringArray | --operation string | --ip string ]
```

--ip *string*

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

--name *string*

Specifies the name of the community list.

--operation *string*

Valid options are update-rule, add-device, remove-device.

--rule *stringArray*

Specifies the rule in format seq[seq-num], action[permit/deny], std-value[{ 1-4294967295) | (AA:NN, AA and NN is 2 bytes) | internet | local-as | no-export | no-advertise }]| ext-value[regular expression]

--type { **standard** | **extended** }

Specifies the type of community list. Valid options are standard, extended.

This example configures a community-list on the specified devices.

```
efa policy community-list update --name stdext1 --type extended
--operation add-device --ip 10.20.246.29-30
```

Community List Name	Seq num	Action	Std Value	Ext Value
stdext1	4	deny		_1000_
stdext1	5	permit		_2000_
stdext1	7	deny		_3000_

Community List details

IP Address	Result	Reason	Rollback reason
10.20.246.29	Success		
10.20.246.30	Success		

Device Results

```
efa policy community-list update --name std1 --type standard
--operation add-device --ip 10.20.63.140-141,10.20.246.29-30
```

Community List Name	Seq num	Action	Std Value	Ext Value
std1	4	deny	local-as	
std1	5	deny	no-export	
std1	7	deny	no-advertise	

Community List details

IP Address	Result	Reason
10.20.63.140	Rollback	

```

| 10.20.63.141 | Rollback |
|
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.29 | Failed | Reason: For seq 5: netconf rpc [error] '%Error:
An IP | | |
| | | Community access-list with this name and instance number
| | |
| | | already exists'; Rollback community list rules from
this | | |
| | | device.
| | |
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.30 | Failed | Reason: For seq 4: netconf rpc [error] '%Error:
An IP | | |
| | | Community access-list with this name and instance number
| | |
| | | already exists'; Rollback community list rules from
this | | |
| | | device.
| | |
+-----+-----+-----+-----+-----+
+-----+
Device Results

```

This example removes the community-list from the specified devices.

```

efa policy community-list update --name comm1 --type standard
--operation remove-device --ip 10.20.63.140-141

+-----+-----+-----+-----+-----+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| comm1                | 3       | permit | 65:12    |           |
+-----+-----+-----+-----+-----+
Community List details
+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+
| 10.20.63.140 | Success |         |                  |
+-----+-----+-----+-----+-----+
| 10.20.63.141 | Success |         |                  |
+-----+-----+-----+-----+-----+
Device Results

```

This example updates a community-list by updating a rule.

```

efa policy community-list update --name commExt1 --type extended
--operation update-rule --rule "seq[1],action[permit],ext-value[_30000_]"

+-----+-----+-----+
| Community List Name | Seq num | Action |
+-----+-----+-----+
| commExt1            | 1       | permit |
+-----+-----+-----+
Community List details
+-----+-----+-----+-----+-----+
| IP Address | Result | Reason
| Rollback reason |
+-----+-----+-----+-----+-----+
+-----+
| 10.139.44.159 | Failed | Failed to create community list for commExt1 on the device
|
|
| 10.139.44.159. Reason: For seq 1: netconf rpc [error]

```

```

|
|
|           |           | '%Error: Same filter is already configured with sequence
|           |           | number 30.'
|
+-----+-----+-----+
+-----+
| 10.139.44.163 | Rollback |
|
+-----+-----+-----+
+-----+
Device Results

efa policy community-list update --name comm1 --type standard
--operation update-rule --rule "seq[5],action[permit],std-value[100;no-advertise]"
+-----+-----+-----+
+-----+
| IP Address | Result | Reason
| Rollback reason |
+-----+-----+-----+
+-----+
| 10.139.44.159 | Failed | Failed to create community list for comm1 on the device
|
|           |           | 10.139.44.159. Reason: For seq 5: netconf rpc [error]
|           |           | '%Error: Same filter is already configured with sequence
|           |           | number 30.'
|
+-----+-----+-----+
+-----+
| 10.139.44.163 | Rollback |
|
+-----+-----+-----+
+-----+
Device Results

```

efa policy community-list delete

Delete community list.

```
efa policy community-list delete [ --name string | --type { standard | extended } | --seq string | --oob ]
```

--name *string*

Specifies the name of the community list.

--oob

Optional command to delete out of band entries and remove configuration from associated devices.

--seq *string*

Sequence numbers. For example 5,10,20 or all.

--type { **standard** | **extended** }

Specifies the type of community list. Valid options are standard, extended.

This example deletes the standard community-list rules on all devices for the name, type, and seq number provided.

```
efa policy community-list delete --name commExt1 --seq all --type standard
+-----+-----+-----+
| Community List Name | Seq num | Action |
+-----+-----+-----+
| commExt1           | 1       | permit |
+-----+-----+-----+
| commExt1           | 2       | permit |
+-----+-----+-----+
| commExt1           | 3       | permit |
+-----+-----+-----+
Community List details
+-----+-----+-----+
+-----+
| IP Address | Result | Reason
| Rollback reason |
+-----+-----+-----+
+-----+
| 10.139.44.159 | Success |
|               |         |
+-----+-----+-----+
+-----+
| 10.139.44.163 | Success |
|               |         |
|               |         |
+-----+-----+-----+
+-----+
Device Results
```

efa policy community-list list

List community-list details

```
efa policy community-list list [ --ip string | --name | --type { standard
| extended } ]
```

--ip *string*

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

--name *string*

Specifies the name of the community list.

--type { **standard** | **extended** }

Specifies the type of community list. Valid options are standard, extended.

This example lists community-list details.

```
efa policy community-list list --type standard --ip 10.20.246.29-30

Community list details:

Name: clist1
Seq: 5
Action: deny
StdValue: 50:125 internet local-as no-advertise
ExtValue:
```

```

Name: clist1
Seq: 15
Action: deny
StdValue: 50:125 local-as
ExtValue:

IP Addresses:
+-----+-----+-----+-----+
| Name | Seq | IP Address | App State |
+-----+-----+-----+-----+
| clist1 | 5 | 10.20.246.29 | cfg-in-sync |
+-----+-----+-----+-----+
| clist1 | 5 | 10.20.246.30 | cfg-in-sync |
+-----+-----+-----+-----+
| clist1 | 15 | 10.20.246.29 | cfg-in-sync |
+-----+-----+-----+-----+
| clist1 | 15 | 10.20.246.30 | cfg-in-sync |
+-----+-----+-----+-----+

```

efa policy extcommunity-list

Extended community list commands

```
efa policy extcommunity-list create [ --name string | --type { standard | extended } | --rule stringArray ]
```

```
efa policy extcommunity-list update [ --name string | --type { standard | extended } | --rule stringArray | --operation string | --ip string ]
```

```
efa policy extcommunity-list delete [ --name string | --type { standard | extended } | --seq string ]
```

```
efa policy extcommunity-list list [ --ip string | --name string | --type { standard | extended } ]
```

--ip *string*

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

--name *string*

The name of the extended community list.

--operation *string*

Valid options are update-rule, add-device, remove-device.

--rule *stringArray*

Rule in format seq[seq-num], action[permit/deny], rt[ASN:NN | IpAddress:NN, ASN and NN is 2 or 4 bytes], soo[ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes] OR seq[num], action[permit/deny], ext-value[regular expression].

--seq *string*

Sequence numbers. For example 5,10,20 or all.

--type { **standard** | **extended** }

The type of extended community list. Valid options are standard, extended.

This example creates an extcommunity-list of type standard.

```
efa policy extcommunity-list create --name excommmlist-1 --type standard
--rule "seq[4],action[permit],soo[10.11.2.3:22]"
```

Extended community list name	Seq num	Action	Rt	Soo	Ext Value
excommmlist-1	4	permit		10.11.2.3:22	

Extended community list details

IP Address	Result	Reason	Rollback reason
------------	--------	--------	-----------------

Device Results

This example adds a device to an extcommunity-list.

```
efa policy extcommunity-list update --name excommmlist-1 --type standard
--operation add-device --ip 10.20.246.29,10.20.246.30
```

Extended community list name	Seq num	Action	Rt	Soo	Ext Value
excommmlist-1	4	permit		10.11.2.3:22	
excommmlist-1	5	deny	1:345		
excommmlist-1	6	permit	1:45	10.11.2.3:22	

Extended community list details

IP Address	Result	Reason	Rollback reason
10.20.246.29	Success		
10.20.246.30	Success		

Device Results

This example deletes a device from an extcommunity-list.

```

efa policy extcommunity-list update --name excommmlist-2 --type extended
--operation remove-device --ip 10.20.246.29,10.20.246.30
+-----+-----+-----+-----+-----+-----+
| Extended community list name | Seq num | Action | Rt | Soo | Ext Value |
+-----+-----+-----+-----+-----+-----+
| excommmlist-2                | 2       | permit |   |   | _15000_   |
+-----+-----+-----+-----+-----+-----+
| excommmlist-2                | 5       | deny   |   |   | _25000_   |
+-----+-----+-----+-----+-----+-----+

Extended community list details
+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+
| 10.20.246.29 | Success |      |                  |
+-----+-----+-----+-----+-----+
| 10.20.246.30 | Success |      |                  |
+-----+-----+-----+-----+-----+

Device Results
    
```

These examples update rules on an extcommunity-list.

```

efa policy extcommunity-list update --name excommmlist-1 --type standard
--operation update-rule --rule "seq[5],action[permit],rt[0:123],soo[0:12]"
+-----+-----+-----+-----+-----+-----+
| Extended community list name | Seq num | Action | Rt | Soo | Ext Value |
+-----+-----+-----+-----+-----+-----+
| excommmlist-1                | 5       | permit | 0:123 | 0:12 |          |
+-----+-----+-----+-----+-----+-----+

Extended community list details
+-----+-----+-----+-----+-----+
+-----+
| IP Address | Result | Reason
| Rollback reason |
+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.29 | Failed | Reason: For seq 5: netconf rpc [error] "'rt
0:123 soo 0:12" |          |
    
```

```

|          |          | is an invalid value.'
|          |
+-----+-----+-----+-----+-----+-----+
+-----+
| 10.20.246.30 | Failed | Reason: For seq 5: netconf rpc [error] '"rt
0:123 soo 0:12" |          |
|          |          | is an invalid value.'
|          |
+-----+-----+-----+-----+-----+-----+
+-----+

Device Results
    
```

This example deletes the extcommunity-list rules on all devices for the name, type, and seq provided.

```

efa policy extcommunity-list delete --name excommmlist-2 --type extended --seq all
+-----+-----+-----+-----+-----+-----+
| Extended community list name | Seq num | Action | Rt | Soo | Ext Value |
+-----+-----+-----+-----+-----+-----+
| excommmlist-2                | 2       | permit |   |   | _15000_   |
+-----+-----+-----+-----+-----+-----+
| excommmlist-2                | 5       | deny  |   |   | _25000_   |
+-----+-----+-----+-----+-----+-----+

Extended community list details
+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+
| 10.20.246.29 | Success |        |                  |
+-----+-----+-----+-----+-----+
| 10.20.246.30 | Success |        |                  |
+-----+-----+-----+-----+-----+

Device Results
    
```

This example lists extcommunity-list details.

```

efa policy extcommunity-list list
Extended community list details:
Name: excommmlist-1
Seq: 5
Action: permit
Route Target: 1:100 2:200 3:145 4:123
Site of Origin: 1.2.3.4:12 5:400 10.11.12.13:22
ExtValue:
    
```

```

Name: excommlist-1
Seq: 6
Action: permit
Route Target: 1:45
Site of Origin: 10.11.2.3:22
ExtValue:

Name: excommlist-1
Seq: 9
Action: deny
Route Target: 1:345
Site of Origin: 6:12
ExtValue:

Name: excommlist-2
Seq: 2
Action: permit
Route Target:
Site of Origin:
ExtValue: _15000_

```

efa policy extcommunity-list create

Create an extcommunity-list.

```
efa policy extcommunity-list create [ --name string | --type { standard | extended } | --rule stringArray ]
```

--name *string*

The name of the extended community list.

--rule *stringArray*

Rule in format seq[seq-num], action[permit/deny], rt[ASN:NN | IpAddress:NN, ASN and NN is 2 or 4 bytes], soo[ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes] OR seq[num], action[permit/deny], ext-value[regular expression].

--type { **standard** | **extended** }

The type of extended community list. Valid options are standard, extended.

This example creates an extcommunity-list of type standard.

```

efa policy extcommunity-list create --name excommlist-1 --type standard
--rule "seq[4],action[permit],soo[10.11.2.3:22]"
+-----+-----+-----+-----+-----+-----+
| Extended community list name | Seq num | Action | Rt |      Soo      | Ext Value |
+-----+-----+-----+-----+-----+-----+
| excommlist-1                |      4  | permit |   | 10.11.2.3:22 |           |
+-----+-----+-----+-----+-----+-----+

Extended community list details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |

```

```

+-----+-----+-----+-----+
Device Results

This example adds an extcommunity-list to the specified devices.

efa policy extcommunity-list update --name excommmlist-1 --type standard
--operation add-device --ip 10.20.246.29,10.20.246.30
+-----+-----+-----+-----+-----+-----+-----+
| Extended community list name | Seq num | Action | Rt |      Soo      | Ext Value |
+-----+-----+-----+-----+-----+-----+-----+
| excommmlist-1                | 4       | permit |   | 10.11.2.3:22 |          |
+-----+-----+-----+-----+-----+-----+-----+
| excommmlist-1                | 5       | deny   | 1:345 |              |          |
+-----+-----+-----+-----+-----+-----+-----+
| excommmlist-1                | 6       | permit | 1:45 | 10.11.2.3:22 |          |
+-----+-----+-----+-----+-----+-----+-----+

Extended community list details

+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.20.246.29 | Success |        |                  |
+-----+-----+-----+-----+
| 10.20.246.30 | Success |        |                  |
+-----+-----+-----+-----+

Device Results

```

efa policy extcommunity-list update

Update extcommunity-list.

```
efa policy extcommunity-list update [ --name string | --type { standard | extended } | --rule stringArray | --operation string | --ip string ]
```

--ip *string*

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

--name *string*

The name of the extended community list.

--operation *string*

Valid options are update-rule, add-device, remove-device.

--rule *stringArray*

Rule in format seq[seq-num], action[permit/deny], rt[ASN:NN | IpAddress:NN, ASN and NN is 2 or 4 bytes], soo[ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes] OR seq[num], action[permit/deny], ext-value[regular expression].

--type { **standard** | **extended** }

The type of extended community list. Valid options are standard, extended.

This example configures an extcommunity-list on the specified devices.

```
efa policy extcommunity-list update --name excommmlist-1 --type standard
--operation add-device --ip 10.20.246.29,10.20.246.30
```

Extended community list name	Seq num	Action	Rt	Soo	Ext Value
excommmlist-1	4	permit		10.11.2.3:22	
excommmlist-1	5	deny	1:345		
excommmlist-1	6	permit	1:45	10.11.2.3:22	

Extended community list details

IP Address	Result	Reason	Rollback reason
10.20.246.29	Rollback		
10.20.246.30	Rollback		

Device Results

This example removes the extcommunity-list from the specified devices.

```
efa policy extcommunity-list update --name excommmlist-2 --type extended
--operation remove-device --ip 10.20.246.29,10.20.246.30
```

Extended community list name	Seq num	Action	Rt	Soo	Ext Value
excommmlist-2	2	permit			_15000_

```

excommlist-2          | 5          | deny  |  |  | _25000_ |
+-----+-----+-----+-----+-----+-----+
Extended community list details
+-----+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+-----+
| 10.20.246.29 | Rollback|        |                  |
+-----+-----+-----+-----+-----+-----+
| 10.20.246.30 | Rollback|        |                  |
+-----+-----+-----+-----+-----+-----+
Device Results

```

This example update rules on an extcommunity-list.

```
efa policy extcommunity-list update --name excommlist-2 --type extended --operation
update-rule --rule "seq[5],action[deny],ext-value[_25000_]"
```

```

+-----+-----+-----+-----+-----+-----+
| Extended community list name | Seq num | Action | Rt | Soo | Ext Value |
+-----+-----+-----+-----+-----+-----+
| excommlist-2          | 5          | deny  |  |  | _25000_ |
+-----+-----+-----+-----+-----+-----+
Extended community list details
+-----+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+-----+
| 10.20.246.29 | Success |        |                  |
+-----+-----+-----+-----+-----+-----+
| 10.20.246.30 | Success |        |                  |
+-----+-----+-----+-----+-----+-----+
Device Results

```

efa policy extcommunity-list delete

Delete extended community list.

```
efa policy extcommunity-list delete [ --name string | --type { standard | extended } | --seq string | --oob ]
```

--name *string*

The name of the extended community list.

--oob

Optional command to delete out of band entries and remove configuratom from associated devices.

--seq *string*

Sequence numbers. For example 5,10,20 or all.

--type { **standard** | **extended** }

The type of extended community list. Valid options are standard, extended.

This example deletes the extcommunity-list rules on all devices for the name, type, and seq provided.

```
efa policy extcommunity-list delete --name excommlist-2 --type extended --seq all
```

Extended community list name	Seq num	Action	Rt	Soo	Ext Value
excommlist-2	2	permit			_15000_
excommlist-2	5	deny			_25000_

Extended community list details

IP Address	Result	Reason	Rollback reason
10.20.246.29	Success		
10.20.246.30	Success		

Device Results

efa policy extcommunity-list list

List extcommunity-list details.

```
efa policy extcommunity-list list [ --ip string | --name string | --type
  { standard | extended } ]
```

--ip *string*

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

--name *string*

The name of the extended community list.

--type { **standard** | **extended** }

The type of extended community list. Valid options are standard, extended.

This example lists extcommunity-list details.

```
efa policy extcommunity-list list
Extended community list details:
Name: excommlist-1
Seq: 5
Action: permit
```

```

Route Target: 1:100 2:200 3:145 4:123
Site of Origin: 1.2.3.4:12 5:400 10.11.12.13:22
ExtValue:

Name: excommlist-1
Seq: 6
Action: permit
Route Target: 1:45
Site of Origin: 10.11.2.3:22
ExtValue:

Name: excommlist-1
Seq: 9
Action: deny
Route Target: 1:345
Site of Origin: 6:12
ExtValue:

Name: excommlist-2
Seq: 2
Action: permit
Route Target:
Site of Origin:
ExtValue: _15000_

```

This example lists the extcommunity-list details for the specified device and extended community list.

```

efa policy extcommunity-list list --ip 10.20.246.29 --name excommlist-1

Extended community list details:

Name: excommlist-1
Seq: 6
Action: permit
Route Target: 1:45
Site of Origin: 10.11.2.3:22
ExtValue:

Name: excommlist-1
Seq: 9
Action: deny
Route Target: 1:345
Site of Origin: 6:12
ExtValue:

IP Addresses:
+-----+-----+-----+-----+
|   Name   | Seq | IP Address | App State |
+-----+-----+-----+-----+
| excommlist-1 | 6   | 10.20.246.29 | cfg-in-sync |
+-----+-----+-----+-----+
| excommlist-1 | 9   | 10.20.246.29 | cfg-in-sync |
+-----+-----+-----+-----+

```

efa policy large-community-list create

Create large community list in XCO.

```

efa policy large-community-list create [ --type string | --name string |
--rule stringArray ]

```


--name *string*

Specifies the name of large community list.

--rule *stringArray*

Rule in format for large community: seq[seq-num],action[permit/deny], value[value: <NN:NN:NN> NN is 4 bytes] (or) ext-value[regular expression].

Example: "seq[5],action[permit],std-value[10:20:30;50:60:70]" (or) "seq[4],action[deny],ext-value[^65000:.*_]"

--type *string*

Specifies the type of large community list to create. Valid types are standard and extended.

```
efa policy large-community-list create --name lgcomm1 --type standard
--rule "seq[5],action[permit],std-value[10:10:10;20:20:20]"

+-----+-----+-----+-----+-----+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcomm1            | 5       | permit | 10:10:10 20:20:20 |          |
+-----+-----+-----+-----+-----+
Community List details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
Device Results
```

```
efa policy large-community-list create --name lgcommExt1 --type extended
--rule "seq[5],action[permit],ext-value[_645XX:.*:.*]"

+-----+-----+-----+-----+-----+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcommExt1         | 5       | permit |          | _645XX:.*:.* |
+-----+-----+-----+-----+-----+
Community List details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
Device Results
```

```
efa policy large-community-list create --name lgcomm1 --type standard
--rule "seq[15],action[permit],std-value[10:10:10;20:20:20]"

+-----+-----+-----+-----+-----+
| Large Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcomm1                  | 15      | permit | 10:10:10 20:20:20 |          |
+-----+-----+-----+-----+-----+
Large Community List details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.139.44.159 | Failed | Policy lgcomm1 type large-community-list seq# 15 operation |
|               |       | failed on device 10.139.44.159 due to Reason: For seq |
|               |       | 15: netconf rpc [error] '%Error: Same filter is already |
|               |       | |
```

```

|
|
| | configured with sequence number 10.'
|
+-----+-----+-----+-----+-----+
+-----+
| 10.139.44.160 | Rollback |
|
+-----+-----+-----+-----+
+-----+
Device Results

```

efa policy large-community-list update

Update large community list on XCO.

```
efa policy large-community-list update [ --type string | --name string |
--ip string | --operation string ]
```

--ip *string*

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

--name *string*

Name of large community list.

--operation *string*

Valid options are `add-device`, `remove-device`.

--type *string*

Type of large community list. Valid types are `standard` and `extended`.

This example configures a Large Community List on the specified device.

```
efa policy large-community-list update --operation add-device --ip 10.139.44.159
--name lgcomm1 --type standard
```

```

+-----+-----+-----+-----+-----+
| Large Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcomm1 | 5 | permit | 10:10:10 20:20:20 | |
+-----+-----+-----+-----+-----+
| lgcomm1 | 10 | permit | 30:30:30 | |
+-----+-----+-----+-----+-----+
Community List details
+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+
| 10.139.44.159 | Success | | |
+-----+-----+-----+-----+-----+
Device Results

```

This example configures a Large Community List of type `extended` on the specified device.

```
efa policy large-community-list update --name lgcommExt1 --type extended
--operation add-device --ip 10.139.44.159
```

```

+-----+-----+-----+-----+-----+
| Community List Name | Seq num | Action | Std Value | Ext Value. |
+-----+-----+-----+-----+-----+
| lgcommExt1 | 5 | permit | | 645XX:.*.* |
+-----+-----+-----+-----+-----+

```

```
Community List details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.139.44.159 | Success | | |
+-----+-----+-----+-----+
Device Results
```

```
efa policy large-community-list update --operation add-device --ip 10.139.44.159-160
--name lgcomm1 --type standard

+-----+-----+-----+-----+-----+
| Large Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcomm1 | 5 | permit | 10:10:10 20:20:20 | |
+-----+-----+-----+-----+-----+
| lgcomm1 | 10 | permit | 30:30:30 | |
+-----+-----+-----+-----+-----+
```

```
Community List details
+-----+-----+-----+-----+
| IP Address | Result | Reason |
| Rollback reason |
+-----+-----+-----+-----+
| 10.139.44.159 | Success | |
| | |
+-----+-----+-----+-----+
| 10.139.44.160 | Failed | Policy lgcomm1 type large-community-list seq# 5 operation
| | |
| | | failed on device 10.139.44.160 due to Reason: For seq
10: | | |
| | | netconf rpc [error] '%Error: An IP Community access-list
| | |
| | | with this name and instance number already exists'
| | |
+-----+-----+-----+-----+
Device Results
```

This example removes a Large Community List from the specified device.

```
efa policy large-community-list update --name lgcomm1 --type standard
--operation remove-device --ip 10.139.44.159

+-----+-----+-----+-----+-----+
| Large Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcomm1 | 5 | permit | 10:10:10 20:20:20 | |
+-----+-----+-----+-----+-----+
| lgcomm1 | 10 | permit | 30:30:30 | |
+-----+-----+-----+-----+-----+
Community List details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.139.44.159 | Success | | |
+-----+-----+-----+-----+
Device Results
```

efa policy large-community-list delete

Delete large community list from devices and XCO.

```
efa policy large-community-list delete [ --type string | --name string |
--seq string | --oob ]
```

--name *string*

Name of large community list.

--oob

Optional command to delete out of band entries and remove configuration from associated devices.

--seq *string*

Sequence numbers. For example 5,10,20 or all.

--type *string*

Type of large community list. Valid types are standard and extended.

This example removes the large community rules from all devices of the type and seq specified.

```
efa policy large-community-list delete --name lgcomm1 --seq all --type standard
```

```
+-----+-----+-----+-----+-----+
| Community List Name | Seq num | Action | Std Value | Ext Value |
+-----+-----+-----+-----+-----+
| lgcomm1             | 5       | deny  | 10:10:10 20:20:20 |          |
+-----+-----+-----+-----+-----+
| lgcomm1             | 10      | permit | 30:30:30          |          |
+-----+-----+-----+-----+-----+
```

Community List details

```
+-----+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+-----+
| 10.139.44.159 | Success |        |                  |
+-----+-----+-----+-----+-----+
```

Device Results

efa policy large-community-list list

Lists the large community list for a list of devices.

```
efa policy large-community-list list [ --ip string --type string | --name
string | ]
```

--ip *string*

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

--name *string*

Specifies the name of large community list.

--type *string*

Specifies the type of large community list. Valid types are standard and extended. Optional. If not specified both standard and extended list will be listed.

```
efa policy large-community-list list

large community list details:

Name: lgcomm1
Seq: 5
Action: deny
StdValue: 10:10:10 20:20:20
ExtValue:

Name: lgcomm1
Seq: 10
Action: permit
StdValue: 30:30
ExtValue:

efa policy large-community-list list --type standard --ip 10.139.44.159

Large community list details:

Name: lgcomm1
Seq: 5
Action: deny
StdValue: 10:10:10 20:20:20
ExtValue:

Name: lgcomm1
Seq: 10
Action: permit
StdValue: 30:30:30
ExtValue:

IP Addresses:
+-----+-----+-----+-----+
| Name   | Seq | IP Address | App State |
+-----+-----+-----+-----+
| lgcomm1 | 5   | 10.139.44.159 | cfg-in-sync |
+-----+-----+-----+-----+
| lgcomm1 | 10  | 10.139.44.159 | cfg-in-sync |
+-----+-----+-----+-----+
```

efa policy prefix-list create

Creates a prefix list.

```
efa policy prefix-list create [ --type string | --name list name | --rule stringArray ]
```

--name *list name*

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|IPv6 prefix],ge[prefix-len],le[prefix-len].
Example: seq [5], action[permit], prefix [10.0.0.0/8|2001:db8: :/32],ge[10], le[24] .

--type *string*

Specifies the type of prefix list. Valid types are `ipv4` and `ipv6`.

Execute **efa policy prefix-list update --operation add-device** to push the configuration on to the device.

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 ---
```

This example creates an IPv6 prefix list.

```
efa policy prefix-list create --type ipv6 --name prefix_1_in
--rule "seq[11], action[permit], prefix[2001:db8::/48]"
```

```
Name: prefix_1_in
+-----+-----+-----+-----+-----+-----+
| Type | Seq num | Action | Prefix | Ge | Le | Status |
+-----+-----+-----+-----+-----+-----+
| ipv6 | 11      | permit | 2001:db8::/48 | | | Success |
+-----+-----+-----+-----+-----+-----+
Prefix-list details
```

efa policy prefix-list update

Configures prefix list rules on specified devices.

```
efa policy prefix-list update [ --type string | --name list name | --
operation { add-rule | remove-rule | add-device | remove-device } |
--rule stringArray | --ip string ]
```

--type *string*

Specifies the type of prefix list. Valid type is [`ipv4` | `ipv6`].

--name *list name*

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|IPv6 prefix],ge[**prefix-len**],le[**prefix-len**].
 Example: seq [5], action[permit], prefix [10.0.0.0/8|2001:db8: :/32],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.

- 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.

This example configures a prefix list on the specified devices.

```
efa policy prefix-list update --type ipv6 --name prefix_1_in --operation add-device
--ip 10.20.246.29-30
+-----+-----+-----+-----+-----+-----+
| Name | Type | Seq num | Action | Prefix | Ge | Le |
+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 11 | permit | 2001:db8::/48 | | |
+-----+-----+-----+-----+-----+-----+
Prefix-list details
+-----+-----+-----+
| IP Address | Result | Reason |
+-----+-----+-----+
| 10.20.246.29 | Success | |
+-----+-----+-----+
| 10.20.246.30 | Success | |
+-----+-----+-----+
Device Results
```

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 --type ipv6 --name prefix_1_in
--rule "seq[11],action[permit],prefix[2001:db8::/48],ge[64],le[128]" --operation add-rule
+-----+-----+-----+-----+-----+-----+
| Name | Type | Seq num | Action | Prefix | Ge | Le |
+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 13 | permit | 2001:db8::/32 | 48 | 128 |
+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 14 | permit | 2003:db8::/32 | 64 | 128 |
+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 15 | deny | 2003:db8::/63 | 64 | 128 |
+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 11 | permit | 2001:db8::/48 | 64 | 128 |
+-----+-----+-----+-----+-----+-----+
Prefix-list details
+-----+-----+-----+
| IP Address | Result | Reason |
+-----+-----+-----+
| 10.20.246.29 | Success | |
+-----+-----+-----+
```

```
| 10.20.246.30 | Success |          |
+-----+-----+-----+
Device Results
```

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]
```

This example configures an IPv6 prefix list on the specified devices.

```
efa policy prefix-list update --type ipv6 --name prefix_1_in
--operation add-device --ip 10.20.246.29-30
+-----+-----+-----+-----+-----+-----+-----+
| Name      | Type | Seq num | Action | Prefix      | Ge | Le |
+-----+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 11      | permit | 2001:db8::/48 |   |   |
+-----+-----+-----+-----+-----+-----+-----+
Prefix-list details
+-----+-----+-----+
| IP Address | Result | Reason |
+-----+-----+-----+
| 10.20.246.29 | Success |        |
+-----+-----+-----+
| 10.20.246.30 | Success |        |
+-----+-----+-----+
Device Results
```

This example adds a rule to an existing IPv6 prefix list.

```
efa policy prefix-list update --type ipv6 --name prefix_1_in
--rule "seq[11],action[permit],prefix[2001:db8::/48],ge[64],le[128]"
--operation add-rule
+-----+-----+-----+-----+-----+-----+-----+
| Name      | Type | Seq num | Action | Prefix      | Ge | Le |
+-----+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 13      | permit | 2001:db8::/32 | 48 | 128 |
+-----+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 14      | permit | 2003:db8::/32 | 64 | 128 |
+-----+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 15      | deny   | 2003:db8::/63 | 64 | 128 |
+-----+-----+-----+-----+-----+-----+-----+
| prefix_1_in | ipv6 | 11      | permit | 2001:db8::/48 | 64 | 128 |
+-----+-----+-----+-----+-----+-----+-----+
Prefix-list details
+-----+-----+-----+
| IP Address | Result | Reason |
+-----+-----+-----+
| 10.20.246.29 | Success |        |
+-----+-----+-----+
| 10.20.246.30 | Success |        |
+-----+-----+-----+
Device Results
```

efa policy prefix-list delete

Deletes prefix list on all devices and from XCO.

```
efa policy prefix-list delete [ --type string | --name string | --oob ]
--type string
```

Specifies the type of prefix list. Valid type is [ipv4|ipv6].

--name *string*

Specifies the name of the prefix list.

--oob

Optional command to delete out of band entries and remove configuratom from associated devices.

This example deletes an IPv4 prefix list named `prefix_v4`.

```
$ efa policy prefix-list delete --type ipv4 --name prefix_v4
```

This example deletes an IPv6 prefix list named `plist2`.

```
efa policy prefix-list delete --name plist2 --type ipv6
+-----+-----+-----+
| Name | Type | Status |
+-----+-----+-----+
| plist2 | ipv6 | Success |
+-----+-----+-----+
Prefix-list details
+-----+-----+-----+
| IP Address | Result | Reason |
+-----+-----+-----+
| 10.20.246.29 | Success | |
+-----+-----+-----+
| 10.20.246.30 | Success | |
+-----+-----+-----+
```

efa policy prefix-list list

Lists the prefix list on the specified devices.

```
efa policy prefix-list list [ --type string | --ip string ]
```

--type *string*

Specifies the type of prefix list. Valid type is [ipv4|ipv6].

--ip *string*

Specifies a comma-separated range of device IP addresses. For example:
10.11.13,10.11.50,10.11.101.

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 |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

This example lists the IPv6 prefix lists on the specified devices.

```
efa policy prefix-list list --type ipv6 --ip 10.20.246.29-30

Prefix-list details:
```

```
Name: prefix_1_in
+-----+-----+-----+-----+-----+-----+-----+-----+
| Type | Seq num | Action | Prefix | Ge | Le | DeviceIP | AppState |
+-----+-----+-----+-----+-----+-----+-----+-----+
| ipv6 | 11      | permit | 2001:db8::/48 |   |   | 10.20.246.29 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+-----+
| ipv6 | 11      | permit | 2001:db8::/48 |   |   | 10.20.246.30 | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+-----+-----+
```

This example lists all prefix lists of type IPv6.

```
efa policy prefix-list list --type ipv6

Prefix-list details:
Name: prefix_1_in
+-----+-----+-----+-----+-----+-----+
| Type | Seq num | Action | Prefix | Ge | Le |
+-----+-----+-----+-----+-----+-----+
| ipv6 | 11      | permit | 2001:db8::/48 |   |   |
+-----+-----+-----+-----+-----+-----+
```

efa policy qos map create

efa policy qos map create

```
efa policy qos map create [ --name string | --type string | --rule
  stringArray ]
```

--name *string*

Specifies the name of the QoS map.

--type *string*

Specifies the type of QoS map. Valid types are pcp-tc-map, dscp-tc-map, tc-pcp-map, and tc-dcsp-map.

--rule *stringArray*

Rule in format "dscp|pcp|tc[from dscp value(s) or range], dscp|pcp|tc[to dscp value(s) or range],dp[DP value]"

This example creates a QoS Map named tenant1PcpToTCMap, which will classify packets with PCP values 0-7 in different Traffic Classes. Values can be provided as a range.

```
efa policy qos map create --name tenant1PcpToTCMap --type pcp-tc-map
--rule "pcp[0],tc[7],dp[1]" --rule "pcp[1],tc[6]" --rule "pcp[2-5],tc[5],dp[2]"
+-----+-----+-----+-----+-----+-----+-----+-----+
| QoS Map Name | Map Type | Dscp | Pcp | Tc | Dp |
+-----+-----+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0 | 3 | 5 | 2 |
+-----+-----+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0 | 4 | 5 | 2 |
+-----+-----+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0 | 5 | 5 | 2 |
+-----+-----+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0 | 0 | 7 | 1 |
+-----+-----+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0 | 1 | 6 | 0 |
+-----+-----+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0 | 2 | 5 | 2 |
+-----+-----+-----+-----+-----+-----+-----+
QoS Map details
```

```
Success
--- Time Elapsed: 270.05958ms ---
```

This example creates a QoS Map named `tenant1DscpToTCMap`, which will classify packets with various DSCP values into different Traffic Classes. Values can be provided as a range. Values that are not specified will be classified into default classes.

```
efa policy qos map create --name tenant1DscpToTCMap --type dscp-tc-map
--rule "dscp[14-25],tc[7],dp[2]" --rule "dscp[8],tc[1]" --rule "dscp[10],tc[1],dp[1]"
+-----+-----+-----+-----+-----+-----+
|   QoS Map Name   | Map Type | Dscp | Pcp | Tc | Dp |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 24  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 19  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 21  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 14  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 15  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 10  | 0  | 1  | 1  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 23  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 25  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 16  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 17  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 20  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 22  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 18  | 0  | 7  | 2  |
+-----+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 8   | 0  | 1  | 0  |
+-----+-----+-----+-----+-----+-----+
QoS Map details
Success
--- Time Elapsed: 387.47453ms ---
```

efa policy qos map update

Update a QoS map.

```
efa policy qos map update [ --name string | --type string | --rule
stringArray ]
```

--name *string*

Specifies the name of the QoS map.

--type *string*

Specifies the type of QoS map. Valid options are `pcp-tc-map`, `dscp-tc-map`, `tc-pcp-map`, and `tc-dscp-map`.

--rule *stringArray*

Rule in format "`dscp|pcp|tc[from value(s) or range], dscp|pcp|tc[to value],dp[DP value]`" Example "`pcp[0;2-4;6],tc[5],dp[2]`" for `pcp-tc-map`, "`pcp[5],tc[1-3;5],dp[1]`" for

tc-pcp-map. And update apply only to the "to" values, like "tc" and "dp" in pcp-tc-map, "pcp" in tc-pcp-map

This example updates the mapping for QoS Map named tenant1PcpToTCMap for the PCP values specified in the command.

The original configuration prior to the update is:

```
efa policy qos map create --name tenant1PcpToTCMap --type pcp-tc-map --rule
"pcp[0],tc[7],dp[1]"
  --rule "pcp[1],tc[6]" --rule "pcp[2-5],tc[4],dp[2]"
```

To update map pcp 0 and pcp 1 to different traffic classes, run the following command:

```
efa policy qos map update --name tenant1PcpToTCMap --type pcp-tc-map
--rule "pcp[0],tc[6],dp[1]" --rule "pcp[1],tc[7]"
+-----+-----+-----+-----+-----+-----+
|   QoS Map Name   | Map Type | Dscp | Pcp | Tc | Dp |
+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0    | 0   | 6   | 1   |
+-----+-----+-----+-----+-----+
| tenant1PcpToTCMap | pcp-tc-map | 0    | 1   | 7   | 0   |
+-----+-----+-----+-----+-----+
QoS Map details
Success
--- Time Elapsed: 153.400707ms ---
```

This example updates the mapping for QoS Map named tenant1DscpToTCMap for the DSCP values specified in this command. Mappings that are not specified remain unchanged. Only existing mappings can be updated.

The original configuration prior to the update is:

```
efa policy qos map create --name tenant1DscpToTCMap --type dscp-tc-map
--rule "dscp[14-25],tc[7],dp[2]" --rule "dscp[8],tc[1]" --rule "dscp[10],tc[1],dp[1]"
```

To update DSCP mappings, run the following command:

```
efa policy qos map update --name tenant1DscpToTCMap --type dscp-tc-map
--rule "dscp[14-16],tc[1]" --rule "dscp[22-25],tc[5],dp[1]"
+-----+-----+-----+-----+-----+-----+
|   QoS Map Name   | Map Type | Dscp | Pcp | Tc | Dp |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 24   | 0   | 5   | 1   |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 25   | 0   | 5   | 1   |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 22   | 0   | 5   | 1   |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 16   | 0   | 1   | 0   |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 14   | 0   | 1   | 0   |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 15   | 0   | 1   | 0   |
+-----+-----+-----+-----+-----+
| tenant1DscpToTCMap | dscp-tc-map | 23   | 0   | 5   | 1   |
+-----+-----+-----+-----+-----+
QoS Map details
Success
--- Time Elapsed: 409.157029ms ---
```

efa policy qos map delete

Delete QoS map

```
efa policy qos map delete [ --name string | --type string | --rule
  stringArray | --oob ]
```

--name *string*

Specifies the name of the QoS map.

--type *string*

Specifies the type of QoS map. Valid options are pcp-tc-map, dscp-tc-map, tc-pcp-map, and tc-dscp-map.

--rule *stringArray*

Specifies rule to be deleted. Rule in format "dscp|pcp|tc][from value(s) or range],dp[DP value]" Example "pcp[0;2-4;6]" for pcp-tc-map, "tc[1-3;5],dp[1]" for tc-pcp-map.

--oob

Optional command to delete out of band entries and remove configuratom from associated devices.

This example deletes the QoS Map tenant1PcpToTCMap from XCO database.

```
efa policy qos map delete --name tenant1PcpToTCMap --type pcp-to-tc
```

This example deletes the user defined DSCP to TC mapping for DSCP values 0-7 in the QoS Map tenant1DscpToTCMap. This causes the traffic with DSCP values 0-7 to be classified based on system defaults.

```
efa policy qos map delete --name tenant1DscpToTCMap --type dscp-to-tc -rule "dscp[0-7]"
```

efa policy qos map list

List QoS maps

```
efa policy qos map list [ --ip device-ips | --name string | --type
  string ]
```

--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.
```

--name *string*

Specifies the name of the QoS map.

--type *string*

Specifies the type of QoS map. Valid options are pcp-tc-map, dscp-tc-map, tc-pcp-map, and tc-dscp-map.

This example lists the configuration details of the specified name and type.

```
efa policy qos map list --name tenant1DscpToTCMap10 --type dscp-tc-map
```

```
QoS Map details:
Name: tenant1DscpToTCMap10
+-----+-----+-----+
```

```

| DSCP | TC | DP |
+-----+-----+-----+
| 0 | 3 | 1 |
+-----+-----+-----+
| 1 | 3 | 1 |
+-----+-----+-----+
| 2 | 3 | 1 |
+-----+-----+-----+
| 3 | 3 | 1 |
+-----+-----+-----+
| 4 | 3 | 1 |
+-----+-----+-----+
| 5 | 3 | 1 |
+-----+-----+-----+
| 6 | 3 | 1 |
+-----+-----+-----+
| 7 | 3 | 1 |
+-----+-----+-----+
| 8 | 1 | 0 |
+-----+-----+-----+
| 10 | 1 | 1 |
+-----+-----+-----+
    
```

This example lists the configuration details of the QoS Map tenant1DscpToTCMap, along with the app state on the specified devices.

```
efa policy qos map list --name tenant1DscpToTCMap10 --type dscp-tc-map --ip 10.20.245.1-2
```

QoS Map details:
Name: tenant1DscpToTCMap10

```

+-----+-----+-----+
| DSCP | TC | DP |
+-----+-----+-----+
| 0 | 3 | 1 |
+-----+-----+-----+
| 1 | 3 | 1 |
+-----+-----+-----+
| 2 | 3 | 1 |
+-----+-----+-----+
| 3 | 3 | 1 |
+-----+-----+-----+
| 4 | 3 | 1 |
+-----+-----+-----+
| 5 | 3 | 1 |
+-----+-----+-----+
| 6 | 3 | 1 |
+-----+-----+-----+
| 7 | 3 | 1 |
+-----+-----+-----+
| 8 | 1 | 0 |
+-----+-----+-----+
| 10 | 1 | 1 |
+-----+-----+-----+
    
```

Device Bindings:

```

+-----+-----+-----+-----+
| Name | DSCP | IP Address | App State |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 0 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 1 | 10.20.245.1 | cfg-refreshed |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 2 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
    
```

```

+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 3 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 4 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 5 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 6 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 7 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 8 | 10.20.245.1 | cfg-refreshed |
+-----+-----+-----+-----+
| tenant1DscpToTCMap10 | 10 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+

```

This example lists configuration details of all the QoS Maps of type dscp-to-tc, along with the app state on the specified device.

```
efa policy qos map list --type dscp-tc-map --ip 10.20.245.1
```

```

QoS Map details:
Name: tenant1DscpToTCMap15
+-----+-----+-----+
| DSCP | TC | DP |
+-----+-----+-----+
| 10 | 3 | 1 |
+-----+-----+-----+
| 15 | 3 | 1 |
+-----+-----+-----+

Name: oobDscpToTCMap1
+-----+-----+-----+
| DSCP | TC | DP |
+-----+-----+-----+
| 20 | 5 | 1 |
+-----+-----+-----+
| 25 | 5 | 1 |
+-----+-----+-----+

Device Bindings:
+-----+-----+-----+-----+
| Name | DSCP | IP Address | App State |
+-----+-----+-----+-----+
| tenant1DscpToTCMap15 | 10 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+-----+
| tenant1DscpToTCMap15 | 15 | 10.20.245.1 | cfg-refreshed |
+-----+-----+-----+-----+
| oobDscpToTCMap1 | 20 | 10.20.245.1 | cfg-not-managed |
+-----+-----+-----+-----+
| oobDscpToTCMap1 | 25 | 10.20.245.1 | cfg-not-managed |
+-----+-----+-----+-----+

```

efa policy qos profile create

Create QoS profile

```
efa policy qos profile create [ --name string | --trust string | --pcp-tc string | --dscp-tc string | --tc-pcp string | --tc-dscp string | --service-policy string ]
```

```
--name string
```

Specifies the name of the QoS profile.

--trust *string*

{ **auto** | **dscp** }. Selecting auto trusts PCP for switched traffic and DSCP for routed traffic. Default is auto.

--pcp-tc *string*

Specifies the PCP to TC QoS map name.

--dscp-tc *string*

Specifies the DSCP to TC QoS map name.

--tc-pcp *string*

Specifies the TC to PCP QoS map name.

--tc-dscp *string*

Specifies the TC to DSCP QoS map name.

--service-policy *string*>

Specifies the Service Policy Map Name along with direction in format "name[servicePolicy],dir[in|out]"

This example creates a QoS profile named "fabricProfile1", which defines dscp trust, DSCP to TC maps, and service policy map "servicePolicy1".

```
efa policy qos profile create --name fabricProfile1 --trust[dscp] --dscp-tc[map2] --tc-
dscp[map3]
--service-policy "name[servicePolicy1],dir[out]"
```

efa policy qos profile update

Update or add value to QoS profile

```
efa policy qos profile update [ --name string | --trust string | --pcp-
tc string | --dscp-tc string | --tc-pcp string | --tc-dscp string |
--service-policy string ]
```

--name *string*

Specifies the name of the QoS profile.

--trust *string*

{ **auto** | **dscp** }. Selecting auto trusts PCP for switched traffic and DSCP for routed traffic. Default is auto.

--pcp-tc *string*

Specifies the PCP to TC QoS map name.

--dscp-tc *string*

Specifies the DSCP to TC QoS map name.

--tc-pcp *string*

Specifies the TC to PCP QoS map name.

--tc-dscp *string*

Specifies the TC to DSCP QoS map name.

--service-policy *string*>

Specifies the Service Policy Map Name along with direction in format
"name[servicePolicy],dir[in|out]"

This example updates a QoS profile named "qosProfile1", enabling system default for qos dscp trust, and applying service policy map "servicePolicy2".

```
efa policy qos profile update --name fabricProfile1 --trust[auto]
--service-policy "name[servicePolicy2],dir[out]"
```

efa policy qos profile delete

Delete or set to default value for QoS profile

```
efa policy qos profile delete [ --name string | --trust string | --pcp-
tc string | --dscp-tc string | --tc-pcp string | --tc-dscp string |
--service-policy string ]
```

--name *string*

Specifies the name of the QoS profile.

--trust *string*

{ **auto** | **dscp** }. Selecting auto trusts PCP for switched traffic and DSCP for routed traffic. Default is auto.

--pcp-tc *string*

Specifies the PCP to TC QoS map name.

--dscp-tc *string*

Specifies the DSCP to TC QoS map name.

--tc-pcp *string*

Specifies the TC to PCP QoS map name.

--tc-dscp *string*

Specifies the TC to DSCP QoS map name.

--service-policy *string*>

Specifies the Service Policy Map Name along with direction in format
"name[servicePolicy],dir[in|out]"

This example removes the PCP to TC map and service policy map associated with QoS profile.

```
eefa policy qos profile delete --name qosProfile1 --pcp-tc pcp2tcMap
--service-policy servicePolicyMap
```

This example deletes the entire specified profile.

```
efa policy qos profile delete --name qosProfile1
```

efa policy qos profile bind

Apply QoS profile settings

```
efa policy qos profile bind [ --name string | --fabric fabricName | --
tenant tenantName | --port stringArray | --po po-list ]
```

--name *string*

Specifies the name of the QoS profile.

--fabric *fabricName*

Specifies the name of the fabric to bind.

--tenant *tenantName*

Specifies the name of the tenant to bind.

--port *stringArray*

Optional. Used to bind fabric-internal ports or a list of physical ports of a tenant to bind. Example: "fabric-internal" or "SW1_IP[0/1,0/2,0/5:1-2], SW2_IP[0/3,0/6,0/9-20]".

--po *po-list*

Optional. Used to specify the list of port-channels of a tenant to bind. Example: po3,po5,po9,po12.

This example applies the configuration of QoS profile "fabricProfile1" to all devices and interfaces in fabric "fabric1".

```
efa policy qos profile bind --name fabricProfile1 --fabric fabric1
```

This example applies the configuration of QoS profile "fabricProfile2" to all internal ports in fabric "fabric1", but excludes applying the configuration on any edge interfaces. These ports were previously bound with "fabricProfile1" in the above example.

```
efa policy qos profile bind --name fabricProfile2 --fabric fabric1 --port "fabric-internal"
```

This example applies the configuration of a QoS profile named "tenantProfile1" to all devices or interfaces in "tenant1".

```
efa policy qos profile bind --name tenantProfile1 --tenant tenant1
```

This example applies the configuration of a QoS profile named "tenantProfile2" to 2 interfaces and 3 port-channels within "tenant1".

```
efa policy qos profile bind --name tenantProfile2 --tenant tenant1 --port "1.1.1.1[0/1,0/2]" --po po3,po4,po22
```

efa policy qos profile unbind

Remove QoS profile settings

```
efa policy qos profile unbind [ --name string | --fabric fabricName | --tenant tenantName | --port stringArray | --po po-list ]
```

--name *string*

Specifies the name of the QoS profile.

--fabric *string*

Specifies the name of the fabric to unbind.

--tenant *tenantName*

Specifies the name of the tenant to bind.

--port *stringArray*

Optional. Used to bind fabric-internal ports or a list of physical ports of a tenant to bind. Example: "fabric-internal" or "SW1_IP[0/1,0/2,0/5:1-2], SW2_IP[0/3,0/6,0/9-20]".

--po *po-list*

Optional. Used to specify the list of port-channels of a tenant to bind. Example:
po3,po5,po9,po12.

This example removes the configuration of a specified QoS profile from all devices and interfaces on a specified fabric.

```
efa policy qos profile unbind --name fabricProfile1 --fabric fabric1
```

This example removes the configuration of a QoS profile named "tenantProfile2" from all devices or interfaces on "tenant1" on which the configuration has been applied.

```
efa policy qos profile unbind --name tenantProfile1 --tenant tenant1
```

This example removes the configuration of a QoS profile named "tenantProfile2" from device IP 1.1.1.1 ethernet 0/1 (ethernet 0/1 is a member of tenant1)

```
efa policy qos profile unbind --name tenantProfile2 --tenant tenant1 --port "1.1.1.1[0/1]" --po po3,po4
```

efa policy qos profile list

List QoS profiles

```
efa policy qos profile list [ --ip device-ips | --name string | --interface interface-name ]
```

--ip *Device-ips*

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

--name *string*

Specifies the name of the QoS profile.

--interface *interface-name*

Specifies the name of the interface. Example "Ethernet 0/3,0/10-16,0/18:2-4" or "Port-channel 1-3,64"

This example lists configuration details of the QoS Profile fabricProfile1 given a range of devices. The output now includes service policy map name.

```
efa policy qos profile list --name fabricProfile1 --ip 10.20.245.1-3
```

```
QoS Profile details:
Name: qosProfile10
Trust: auto
Dscp->TC: Map2
TC->Dscp: Map3
Service Policy Map: servicePolicy1, out
Device Bindings:
+-----+-----+-----+
| Name          | IP Address  | App State  |
+-----+-----+-----+
| fabricProfile1 | 10.20.245.1 | cfg-in-sync |
+-----+-----+-----+
|                | 10.20.245.2 | cfg-in-sync |
+-----+-----+-----+
|                | 10.20.245.3 | cfg-in-sync |
+-----+-----+-----+
```

efa policy qos service-policy-map create

Create service policy map

```
efa policy qos service-policy-map create [ --name string | --rule
stringArray ]
```

--name *string*

Specifies the name of the QoS service policy map.

--rule *stringArray*

Rule in format "strict-priority[3],dwrr[25;25;25;25;0],class-map[default]"

Table 1: Valid values for strict-priority

Value	Definition
0	No strict priority queue; all are DWRR
1	Traffic Class 7 strict priority queue; rest are DWRR.
2	Traffic Class 6 through 7 strict priority queues; rest are DWRR.
3	Traffic Class 5 through 7 strict priority queues; rest are DWRR.
4	Traffic Class 4 through 7 strict priority queues; rest are DWRR.
5	Traffic Class 3 through 7 strict priority queues; rest are DWRR.
6	Traffic Class 2 through 7 strict priority queues; rest are DWRR.
7	Traffic Class 1 through 7 strict priority queues; rest are DWRR.

default is the only value supported for *class-map*.

The sum of DWRR weights must equal 100.

This example creates a QoS service policy map named `servicePolicy1`, specifying strict priority for traffic classes 5-7 and specifies DWRR weights for traffic classes 0-4.

```
efa policy qos service-policy-map create --name servicePolicy1
--rule "strict-priority[3],dwrr[25;25;25;25;0],class-map[default]"
```

efa policy qos service-policy-map delete

Delete QoS service policy map

```
efa policy qos service-policy-map delete [ --name string | --oob ]
```

--name *string*

Specifies the name of the QoS service policy map to delete.

--oob

Optional command to delete out of band entries and remove configuration from associated devices.

This example deletes a QoS service policy map named servicePolicy1.

```
efa policy qos service-policy-map delete --name servicePolicy1
```

efa policy qos service-policy-map list

Show QoS service policy map configurations.

```
efa policy qos service-policy-map list [ --name string | --ip ip-addr ]
```

--name *string*

Specifies the name of the QoS service policy map.

--ip *ip-addr*

Specifies the IP address of the device.

This example lists the configuration details of the QoS service policy map named servicePolicy1

```
efa policy qos service-policy-map list --name servicePolicy1
```

QoS Service Policy Map details:

Name: servicePolicy1

Strict Priority	DWRR weights	Class
2	50,25,25,0,0,0	default

This example lists the configuration details of the QoS service policy map named servicePolicy1, along with the app state on the specified device.

```
efa policy qos service-policy-map list --name servicePolicy1 --ip 10.20.245.1-2
```

QoS Service Policy Map details:

Name: servicePolicy1

Strict Priority	DWRR weights	Class
2	50,25,25,0,0,0	default

Device Bindings:

IP Address	App State
10.20.245.1	cfg-in-sync
10.20.245.2	cfg-in-sync

efa policy qos service-policy-map update

Update QoS service policy map

```
efa policy qos service-policy-map update [ --name string | --rule stringArray ]
```

--name *string*

Specifies the name of the QoS service policy map.

--rule *stringArray*

Rule in format "strict-priority[3],dwrr[25;25;25;25;0],class-map[default]"

Table 2: Valid values for strict-priority

Value	Definition
0	No strict priority queue; all are DWRR
1	Traffic Class 7 strict priority queue; rest are DWRR.
2	Traffic Class 6 through 7 strict priority queues; rest are DWRR.
3	Traffic Class 5 through 7 strict priority queues; rest are DWRR.
4	Traffic Class 4 through 7 strict priority queues; rest are DWRR.
5	Traffic Class 3 through 7 strict priority queues; rest are DWRR.
6	Traffic Class 2 through 7 strict priority queues; rest are DWRR.
7	Traffic Class 1 through 7 strict priority queues; rest are DWRR.

default is the only value supported for *class-map*.

The sum of DWRR weights must equal 100.

Update of a policy map's content could be restricted if the map is referenced by an actively bound QoS profile.

This example updates a QoS service policy map named `servicePolicy1`, for strict priority value of 2 and new weights as follows.

```
efa policy qos service-policy-map update --name servicePolicy1
--rule "strict-priority[2],dwrr[50;25;25;0;0;0],class-map[default]"
```

efa policy route-map create

Creates a route map with one or more rules.

efa policy route-map create [**--name** *string* | **--rule** *stringArray*]

--name *string*

Specifies the name of the route map.

--rule *stringArray*

Specifies the rule in the following format: `seq[seq-num],action[permit/deny]`.

Execute **efa policy route-map update --operation add-device** to push the configuration initially on to the device. Any subsequent changes to existing configurations are automatically applied to devices where the route-map is present.

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.

```
efa policy route-map update [ --name string | --rule string | --operation
  { add-device | remove-device | update-action } | --ip string]
```

--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.

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.

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 XCO:

- `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.

```
efa policy route-map delete [ --name string | --seq sequence-numbers |
--oob ]
```

--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.

--oob

Optional command to delete out of band entries and remove configuration from associated devices.

The command removes the route map rule from the XCO 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.

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.

efa policy route-map list [**--ip** *string* | **--name** *string*]

--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.

--name *string*

Specifies the name of the route map.

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`.

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.

efa policy route-map-match create [**--name** *route-map name* | **--rule** *string*
| **--match-ipv4-prefix** *string* | **--match-ipv6-prefix** *string* | **--match-**
community-list *string* | **--match-extcommunity-list** *string* | **--match-**
largecommunity-list *string*]

--name *route-map name*

Specifies the name of the route map.

--rule *string*

Specifies the rule in the following format: seq[seq-num],action[permit/deny].

Example: seq[5],action[permit].

--match-ipv4-prefix *string*

Specifies the name of the IPv4 prefix list that contains the rule that you want to match.

--match-ipv6-prefix *string*

Specifies the name of the IPv6 prefix list that contains the rule that you want to match.

--match-community-list *string*

Specifies the community list name.

--match-extcommunity-list *string*

Specifies the extcommunity list name.

--match-largcommunity-list *string*

Specifies the large community list name.

There is no separate CLI to display the route-map match criteria. The information is shown in **efa policy route-map list**.

```
efa policy route-map-match create --name rmap_1
--rule seq[5],action[permit] --match-ipv4-prefix prefix_1
```

```
efa policy route-map-match create --name telco_3
_in --rule seq[14],action[permit] --match-ipv4-prefix prefix_1_in
```

```
efa policy route-map-match create --name telco_3_in
--rule seq[14],action[permit] --match-community-list commList1 --match-extcommunity-list
excommList1
```

```
efa policy route-map-match create --name foo
--rule seq[10],action[permit] --match-extcommunity-list extFoo
```

```
efa policy route-map-match create --name rmap_test
--rule "seq[11],action[permit]" --match-ipv6-prefix prefix_1_in
```

```
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap_test      | 11      | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+-----+
| IP Address  | Result  | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.20.246.29 | Success |        |                  |
+-----+-----+-----+-----+
| 10.20.246.30 | Success |        |                  |
+-----+-----+-----+-----+
Device Results
```

```
efa policy route-map-match create --name rmap_test
--rule "seq[11],action[permit]" --match-ipv6-prefix prefix_1_in
```

```
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap_test      | 11      | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+-----+
```

```

| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.20.246.29 | Success | | |
+-----+-----+-----+-----+
| 10.20.246.30 | Success | | |
+-----+-----+-----+-----+
Device Results

```

```

efa policy route-map-match delete --name rmap_test
--rule "seq[10],action[permit]" --match-ipv6-prefix prefix_1_in
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap_test      | 10      | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+-----+
| IP Address | Result | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.20.246.29 | Success | | |
+-----+-----+-----+-----+
| 10.20.246.30 | Success | | |
+-----+-----+-----+-----+
Device Results

```

efa policy route-map-match delete

Deletes the match criterion for a route map rule.

```

efa policy route-map-match delete [ --name string | --rule string |
--match-ipv4-prefix string | --match-ipv6-prefix string | --match-
community-list string | --match-extcommunity-list string | --match-
largecommunity-list string ]

```

--name *string*

Specifies the name of the route map.

--rule *string*

Specifies the rule in the following format: seq[seq-num],action[permit/deny].

Example: seq[5],action[permit].

--match-ipv4-prefix *string*

Specifies the name of the IPv4 prefix list that contains the rule that you want to delete.

--match-ipv6-prefix *string*

Specifies the name of the IPv6 prefix list that contains the rule that you want to match.

--match-community-list *string*

Specifies the community-list name.

--match-extcommunity-list *string*

Specifies the excommunity-list name.

--match-largecommunity-list *string*

Specifies the large community list name.

```
efa policy route-map-match delete --name telco_3_in
--rule seq[14],action[permit] --match-community-list commList1

efa policy route-map-match delete --name telco_3_in
--rule seq[14],action[permit] --match-extcommunity-list excommList1

efa policy route-map-match delete --name telco_3_in
--rule seq[14],action[permit] --match-ipv4-prefix prefix_1_in
efa policy route-map-match delete --name foo
--rule seq[10],action[permit] --match-extcommunity-list extFoo

efa policy route-map-match delete --name rmap_test --rule "seq[10],action[permit]"
--match-ipv6-prefix prefix_1_in
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap_test      | 10      | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+-----+
| IP Address  | Result  | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.20.246.29 | Success |        |                  |
+-----+-----+-----+-----+
| 10.20.246.30 | Success |        |                  |
+-----+-----+-----+-----+
Device Results

Name: rmap_test
Seq: 10
Action: permit
Matches:
  match: matchIPv6PrefixList prefix_1_in
Sets:
Name: rmap_test
Seq: 11
Action: permit
Matches:
  match: matchIPv6PrefixList prefix_1_in
Sets:
```

efa policy route-map-set create

Create route map set.

```
efa policy route-map-set create [ --name string | --rule string
  | --set-community string | --set-extcommunity-rt string | --set-
extcommunity-soo string | --set-communitylist-delete string | --set-
large-community string | --set-large-communitylist-delete string |
--name string
  Specifies the name of the route map.
--rule string
  Rule in format seq[seq-num],action[permit/deny].
--set-community string
```

In format [1-4294967295 | AA:NN, AA and NN is 2 bytes | internet | local-as | no-export | no-advertise]. Example: 6550:125,local-as,internet.

--set-extcommunity-rt *string*

In format [ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes]. Example: 2:300,12.12.13.33:24.

--set-ext-community-soo *string*

In format [ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes]. Example: 32:124.

--set-communitylist-delete *string*

Specifies the name of the community list to delete.

--set-large-community *string*

In format value[value: <NN:NN:NN | additive> NN is 4 bytes]. Example: 10:20:30,40:50:60,additive.

--set-large-communitylist-delete *string*

Specifies the name of the large community list to delete.

```
efa policy route-map-set create --name rmap1 --rule "seq[1],action[permit]"
--set-largecommunitylist-delete lgcomm1
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap1          | 1       | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+-----+
| IP Address   | Result  | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.139.44.159 | Success |        |                  |
+-----+-----+-----+-----+
Device Results

efa policy route-map-set create --name rmap1 --rule "seq[1],action[permit]"
--set-largecommunity 50:50:50,additive
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap1          | 1       | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+-----+
| IP Address   | Result  | Reason | Rollback reason |
+-----+-----+-----+-----+
| 10.139.44.159 | Success |        |                  |
+-----+-----+-----+-----+
Device Results
```

efa policy route-map-set delete

Delete route map set.

```
efa policy route-map-set delete [ --name string | --rule string
| --set-community string | --set-extcommunity-rt string | --set-
extcommunity-soo string | --set-communitylist-delete string | --set-
large-community string | --set-large-communitylist-delete string |
```

- name** *string*
Specifies the name of the route map.
- rule** *string*
Rule in format seq[seq-num],action[permit/deny].
- set-community** *string*
In format [1-4294967295 | AA:NN, AA and NN is 2 bytes | internet | local-as | no-export | no-advertise]. Example: 6550:125,local-as,internet.
- set-extcommunity-rt** *string*
In format [ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes]. Example: 2:300,12.12.13.33:24.
- set-extcommunity-soo** *string*
In format [ASN:NN|IpAddress:NN, ASN and NN is 2 or 4 bytes]. Example: 32:124.
- set-communitylist-delete** *string*
Specifies the name of the community list to delete.
- set-large-community** *string*
In format value[value: <NN:NN:NN | additive> NN is 4 bytes].
- set-large-communitylist-delete** *string*
Specifies the name of the large community list to delete.

```
efa policy route-map-set delete --name rmap1 --rule "seq[1],action[permit]"
--set-largecommunity 50:50:50,additive --set-largecommunitylist-delete lg1
+-----+-----+-----+
| Route Map Name | Seq num | Action |
+-----+-----+-----+
| rmap1          | 1       | permit |
+-----+-----+-----+
Route Map details
+-----+-----+-----+
| IP Address    | Result  | Reason | Rollback reason |
+-----+-----+-----+
| 10.139.44.159 | Success |        |                  |
+-----+-----+-----+
Device Results
```

efa rbac execution

Displays or deletes the execution logs of the Auth and RBAC services.

```
efa rbac execution show [ --id id | --limit exec-limit | --status
{ failed | succeeded | all }]
```

```
efa rbac execution delete [ --day days ]
```

--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 XCO.

efa rbac role show [**--name** *role-name*]

--name *role-name*

Specifies name of the role.

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 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 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.

```
efa scvmm delete [ --host ip | --cleanup-eggs ]
```

```
--host ip
```

Specifies the IP address or host name of the SCVMM server to be deleted.

```
--cleanup-eggs
```

Cleans up EPGs from the tenant service when the SCVMM server is deleted.

efa scvmm links physical

Configures SCVMM physical links.

```
efa scvmm links physical [ --host ip-host | --hyperv hostname ]
```

```
--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.

```
efa scvmm links virtual [ --hyperv hostname ]
```

```
--hyperv hostname
```

Specifies the name of the Hyper-V host.

efa scvmm list

Lists all SCVMM servers and their summary information.

```
efa scvmm list [ --host string]
```

```
--host string
```

Specifies the name or IP address of the SCVMM server.

efa scvmm register

Registers an SCVMM server.

```
efa scvmm register [ --host ip | --username user name | --password password | --tenant tenant ]
```

--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.

```
efa scvmm settings show
```

efa scvmm settings update

Updates SCVMM service settings.

```
efa scvmm settings update [ --poll-frequency frequency ]
```

--poll-frequency *frequency*

Specifies the polling frequency in hours. Valid values range from 1 to 24.

efa scvmm update

Updates the SCVMM server.

```
efa scvmm update [ --host ip-host | --username user name | --password password ]
```

--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. Traps received can have source IPv4/IPv6 address of Active Node or Virtual IPv4/IPv6 address for multi node deployment.

efa snmp subscriber list

efa snmp subscriber add [**--host** *ipv4/ipv6* | **--type** *type* | **--community** *community* | **--v3user** *user* | **--auth** *auth_protocol* | **--authpassword** *password* | **--priv** *priv_protocol* | **--privpassword** *password*]

efa snmp subscriber unregister [**--host** *ipv4/ipv6*]

--host *ipv4/ipv6*

(Required) Specifies the IP address or host name of the SCVMM server you want to update. Can be either IPv4 or IPv6 address.

--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>**".

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 XCO configurations.

efa show-running-config

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 XCO 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.

efa status

Roles are either Active or Standby.

The following example shows output for a single node installation.

```
xmcdev@xmcdev-virtual-machine:~$ efa status
+-----+-----+-----+-----+
| Node Name           | Role  | Status | IP           |
+-----+-----+-----+-----+
| xmcdev-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 alarm acknowledge

Acknowledges the alarm instance referenced by resource path with query parameters and optional comment.

```
efa system alarm acknowledge [ --resource string | --instance instance-id
  --comment string ]
```

--resource *string*

Specifies the health resource path, with possible query parameters to acknowledge the raised alarm. Value is enclosed in quotes. Example: `--resource "/App/System/Security/Certificate?type=app_server_cert"`.

--instance *instance-id*

Specifies the run-time generated IDs of alarms. Instance id can be single, sequential, or multiple. Example: `--instance 3,4,5-6` or `all`.

--comment *string*

Brief description that is tracked with the acknowledgment. Value is enclosed in quotes. Example: `--comment "Will investigate, ticket 4412"`

This command is available only to users with the admin role.

The following example acknowledges the alarm instance referenced by resource path with query parameters and optional comment.

```
# efa system alarm acknowledge --resource "/App/System/Security/Certificate?
type=app_server_cert"
--comment "Will investigate, ticket 12345"
```

The following example acknowledges the alarm by run-time generated id and no additional comment.

```
# efa system alarm acknowledge --instance 1
```

efa system alarm close

Closes an open alarm.

```
efa system alarm close [ --resource string | --instance string --comment
  string ]
```

--resource *string*

Specifies the health resource path, with possible query parameters to close a raised alarm. Example: `--resource "/App/System/Security/Certificate?type=app_server_cert"`.

--instance *string*

Specifies the run-time generated IDs of alarms. Instance id can be single, sequential, or multiple. Example: `--instance 3,4,5-6` or `all`.

--comment *string*

Brief description that is tracked with the close operation. Example: `--comment "Will investigate, ticket 4412"`.

This command is available only to users with the admin role.

You can close an alarm, but cannot clear it. Only the system can clear a raised alarm.

You cannot close or acknowledge a cleared alarm until the system reopens it.

The alarm close operation is stored in the user state change list. When the user state change list is full, the oldest user state changes are overwritten.

The following example closes the alarm instance referenced by resource path with query parameters and no additional comment.

```
# efa system alarm close --resource "/App/System/Security/Certificate?
type=app_server_cert"
```

The following example closes the alarm by run-time generated id with a comment.

```
# efa system alarm close --instance 1 --comment "Ticket 12345 has been investigated,
partially resolved, and determined that the alarm can be closed."
```

efa system alarm inventory show

Displays the list of alarms supported by XCO.

```
efa system alarm inventory show [ --name alarm-name | --id alarm-id |
--type alarm-category | --resource resource-path | --detail ]
```

--name *alarm-name*

Specifies the alarm inventory name. Example: --name CertificateExpiration.

--id *alarm-id*

Specifies the alarm inventory id. Example: --id 32000.

--type *alarm-category*

Specifies the alarm inventory type. Example: --type security.

--resource *resource-path*

Specifies the health resource path string. Example: --resource "/App/System/Security/Certificate".

--detail

Displays detailed system alarm information.

This command is available only to users with the admin role.

The following example displays the list of supported alarms in the inventory.

```
# efa system alarm inventory show [ --detail ]
+-----+-----+-----+-----+-----+-----+
| Name | ID | Type | Resource | Will Clear | Severity |
|-----+-----+-----+-----+-----+-----+
| CertificateExpiration | 32000 | Security | /App/System/Security/Certificate | true | Warning,
Critical | Notify an application certificate is about to expire or has |
| | | | | | |
| | | | | | |
| | | | | | |
| | | | | | |
|-----+-----+-----+-----+-----+-----+
| DeviceCertificateExpiration | 32001 | Security | /App/System/Security/Certificate | true | Warning,
```

Critical	Notify a managed device certificate is about to expire or				
	has expired.				

LoginAuthentication	32010	Security		false	
Warning	Notify suspicious login activity on the application.				

StorageUtilization	32040	Processing	/App/System/Storage	true	Warning,
Minor, Major, Critical	Notify when storage utilization for the application has				
	reached a certain threshold or is full				

The following example displays an alarm and its configuration in the inventory by the alarm name.

```
# efa system alarm inventory show --name CertificateExpiration --detail
```

The following example displays an alarm and its configuration in the inventory by the alarm id.

```
# efa system alarm inventory show --id 1 --detail
```

The following example displays the set of alarms and their configurations in the inventory by the alarm type.

```
# efa system alarm inventory show --type security --detail
```

The following example displays the set of alarms and their configurations in the inventory by the associated resource path.

```
# efa system alarm inventory show --resource /App/System/Security/Certificate --detail
```

efa system alarm purge

Purges all closed alarms.

efa system alarm purge

This command is available only to users with the admin role.

An alarm purge releases all resources related to the stateful alarm entity.

The following example purges all closed alarms from the system.

```
# efa system alarm purge
```

efa system alarm show

Displays the list of alarm history.

```
efa system alarm show [ --unacked | --acked | --cleared | --closed | --
  name alarm-name | --id alarm-id | --type alarm-category | --resource
  resource-path | --instance run-time-instance-id | --detail ]
```

--unacked

Specifies the unacknowledged alarms.

--acked

Specifies the acknowledged alarms.

--cleared

Specifies the cleared alarms.

--closed

Specifies the closed alarms.

--name *alarm-name*

Specifies the name of the alarm. Example: --name CertificateExpiration

--id *alarm-id*

Specifies the alarm inventory id. Example: --id 32000.

--type *alarm-category*

Specifies the alarm type. Example: --type security.

--resource *resource-path*

Specifies the health resource path string. Example: --resource /App/System/Security/Certificate

--instance *run-time-instance-id*

Specifies the run-time generated ID for the alarm. Example: --instance 3.

--detail

Displays the detailed system alarm information.

This command is available only to users with the admin role.

If parameters are not provided, this command lists unclosed alarms.

Perform the alarms order on server-side to have consistent ordering on both REST and CLI.

The following example displays alarms that are currently open.

```
# efa system alarm show --detail

+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| Instance | Name | ID | Type | |
Resource | | Severity | Is Cleared | |
| | Time Created | Last Raised | Last Changed |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 2 | DeviceCertificateExpiration | 32001 | Security | /App/System/Security/
Certificate?device_ip=10.20.30.43&type=https_server_cert | Critical | false | The HTTPS Server Certificate on device "10.20.30.43"
has | 2022-08-18 09:46:49.954727778 -0700 PDT | 2022-08-18 09:46:49.954728325 -0700 PDT | 2022-08-18 09:46:49.954728607 -0700 PDT |
| | | | |
| | expired on "2022-08-18 09:46:49.954548665 -0700
PDT" | | | |
+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+
| 4 | StorageUtilization | 32040 | Processing | /App/
System/Storage?node_ip=10.10.10.1&mount_point=%2F | Major | true | The Node IP "10.10.10.1" mount point "/" is at
a safe | 2022-11-11 09:46:49.954728804 -0800 PST | 2022-12-12 09:46:49.954729031 -0800 PST | 2022-12-12 09:46:49.954729303 -0800 PST |
| | | | |
| | storage utilization of 55% with 4.371 GB free.
| | | | |
```



```

+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| 3 | | LoginAuthentication |
| 32010 | Security |
| Warning | false | Authentication failed for user "extreme"
| 2022-12-16 09:46:49.954729461 -0800 PST | 2022-12-16 09:46:49.954729617 -0800 PST | 2022-12-16 09:46:49.954729774 -0800 PST |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
| 1 | | CertificateExpiration | 32000 | Security | /App/
System/Security/Certificate?type=app_server_cert | Warning | false | The App Server Certificate on the application will
expire | 2022-10-08 09:46:49.954740405 -0700 PDT | 2022-10-08 09:46:49.954740671 -0700 PDT | 2022-10-08 09:46:49.954740891 -0700 PDT |
| | |
| | | soon on "2023-01-08 09:46:49.95472993 -0800
PST" | |
+-----+-----+-----+-----+
+-----+-----+-----+-----+
+-----+-----+-----+-----+

```

The following example displays alarms that are currently open and unacknowledged.

```
# efa system alarm show --unacked --detail
```

The following example displays alarms that are currently open and acknowledged.

```
# efa system alarm show --acked --detail
```

The following example displays alarms that are currently open and cleared by the system.

```
# efa system alarm show --cleared --detail
```

The following example displays alarms that are closed by the user.

```
# efa system alarm show --closed --detail
```

The following example displays one or more instances of the alarm by the given inventory name even if the alarm is closed.

```
# efa system alarm show --name CertificateExpiration --detail
```

The following example displays one or more instances of the alarm by the given inventory id even if the alarm is closed.

```
# efa system alarm show --id 32000 --detail
```

The following example displays one or more instances of the alarm by the given inventory type even if the alarm is closed.

```
# efa system alarm show --type security --detail
```

The following example displays one or more instances of the alarm by the resource path (as referenced in the alarm inventory) even if the alarm is closed.

```
# efa system alarm show --resource "/App/System/Security/Certificate" --detail
```

The following example displays the instance of the alarm by the resource path with query parameters even if the alarm is closed.

```
# efa system alarm show --resource "/App/System/Security/Certificate?
type=app_server_cert" --detail
```

The following example displays instance of the alarm by the run-time generated id even if the alarm is closed.

```
# efa system alarm show --instance 1 --detail
```

efa system alarm summary

Displays alarm counts and statistics grouped by severity.

efa system alarm summary

This command is available only to users with the admin role.

The following example displays the alarm summary.

```
efa system alarm summary
```

efa system alert

Alert commands.

```
efa system alert show [ --id int32 | --resource string | --severity string | --sequence-id int32 | --limit int32 | --before-timestamp string | --after-timestamp string ]
```

```
efa system alert inventory show [ --id int32 | --resource string ]
```

efa system alert clear

--id *int32*

Specifies the ID identifying XCO alert. Example: --id 3100.

--resource *string*

Specifies the health resource path associated to the alert being sent. Example: --resource /App/System/Security/Certificate

--severity *string*

Specifies the severity of XCO alert. Valid values are <critical|major|minor|warning|info>. Example: --severity warning.

--sequence-id *int32*

Tracks the sequence that messages are submitted. Example --sequence-id 41.

--limit *int32*

Limits the number of alerts to be displayed. Example: --limit 2.

--before-timestamp *string*

Alerts being sent before a particular time stamp. Example: --before-timestamp 2022-10-11T22:14:15.003Z

--after-timestamp *string*

Alerts being sent after a particular time stamp. Example: --after-timestamp 2022-10-11T22:14:15.003Z

efa system backup

Backs up the XCO system, including the database and certificates.

```
efa system backup [ --device-ip stringArray | --remote | --fabric-all |
--fabric-name string ]
```

--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. Use the [efa system settings update](#) on page 219 command to configure the remote server.

--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.

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.

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 XCO backup files.

efa system backup-list

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.

```
efa system cleanup [ --archive-type { backup | support } | --name
  string ]
```

```
--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.

```
efa system feature show
```

Use the **efa system feature update** command to enable the feature.

For more information, see the "In-flight Transaction Recovery" topic in the [ExtremeCloud Orchestrator CLI Administration Guide, 3.7.0](#).

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.

```
efa system feature update [--inflight-transaction-auto-recovery { enable
  | disable } | --tenant-api-concurrency { enable | disable }]
```

```
--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.

After a service restart or high-availability failover, XCO 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 *ExtremeCloud Orchestrator Administration Guide*.

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 logging set

Configure logging of a service.

```
efa system logging set [ --service service-name | --type logging-type |
--size file-size | --maximum-files number-of-files | --level logging-
level ]
```

--service *service-name*

Comma separated list of service names. Valid values are [all|auth|fabric|faultmanager|hyperv|inventory|notification|policy|raslog|rbac|system|tenant|vcenter].

--type *logging-type*

Logging type to modify. Valid values are [all|error]. Default all.

--size *file-size*

Logging file size in megabytes. Default 100 for type base. Default 10 for type error. Range 1-100.

--maximum-files *number-of-files*

Maximum number of Logging files. Default 10. Range 1-100.

--level *logging-level*

Detail level of logging. Allow logging between panic and selected level. Valid values are [panic|fatal|error|warning|info|debugging]. Default debugging for type base. Default error for type error.

The following examples configures logging on the specified services.

```
efa system logging set --service inventory --size 200 --level info
+-----+-----+-----+-----+-----+-----+-----+
```

```

| Service | Type | Size | Maximum Files | Level | Status | Reason |
+-----+-----+-----+-----+-----+-----+-----+
| inventory | all | 200 | | info | Success | |
+-----+-----+-----+-----+-----+-----+-----+
Logging set

efa system logging set --service inventory,fabric --type error --maximum-files 20
+-----+-----+-----+-----+-----+-----+-----+
| Service | Type | Size | Maximum Files | Level | Status | Reason |
+-----+-----+-----+-----+-----+-----+-----+
| inventory | error | | 20 | | Success | |
+-----+-----+-----+-----+-----+-----+-----+
| fabric | error | | 20 | | Success | |
+-----+-----+-----+-----+-----+-----+-----+
Logging set

efa system logging set --service inventory --type error --level fatal
+-----+-----+-----+-----+-----+-----+-----+
| Service | Type | Size | Maximum Files | Level | Status | Reason |
+-----+-----+-----+-----+-----+-----+-----+
| inventory | error | | | fatal | Success | |
+-----+-----+-----+-----+-----+-----+-----+
Logging set

```

efa system logging unset

Unconfigure logging for a service.

```
efa system logging unset [ --service service name | --type logging type ]
```

--service *service name*

Comma separated list of service names. Valid values are [all|auth|fabric|faultmanager|hyperv|inventory| notification|policy|raslog|rbac| system|tenant|vcenter].

--type *logging type*

Logging type to modify. Valid values are [all|error]. Default all.

The following examples set logging configuration to defaults with the unset command.

```

efa system logging unset --service inventory
+-----+-----+-----+-----+
| Service | Type | Status | Reason |
+-----+-----+-----+-----+
| inventory | all | Success | |
+-----+-----+-----+-----+
Logging unset

efa system logging unset --service inventory,fabric --type error
+-----+-----+-----+-----+
| Service | Type | Status | Reason |
+-----+-----+-----+-----+
| inventory | error | Success | |
+-----+-----+-----+-----+
| fabric | error | Success | |
+-----+-----+-----+-----+
Logging unset

```

efa system logging show

Display the current logging configuration.

```
efa system logging show [ --service service name | --type logging type ]
```

--service *service name*

Comma separated list of service names. Valid values

are [all|auth|fabric|faultmanager|hyperv|inventory| notification|policy|raslog|rbac| system|tenant|vcenter].

--type *logging type*

Logging type to modify. Valid values are [all|error]. Default all.

This example shows the current logging customization configuration.

```
efa system logging show
+-----+-----+-----+-----+-----+
| Service | Type | File Size | Number of Files | Detail Level |
+-----+-----+-----+-----+-----+
| auth    | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| fabric  | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| fabricmanager | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| hyperv  | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| inventory | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| notification | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| policy  | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| raslog  | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| rbac    | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| system  | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
+-----+-----+-----+-----+-----+
| tenant  | all  | 100 MB   | 10                | debugging    |
+-----+-----+-----+-----+-----+
|         | error| 10 MB    | 10                | error        |
```



```

+-----+-----+-----+-----+-----+
|  vcenter  |  all  | 100 MB  |      10  | debugging |
+-----+-----+-----+-----+-----+
|           | error |  10 MB  |      10  |    error  |
+-----+-----+-----+-----+-----+
Logging show

efa system logging show --service inventory --type all
+-----+-----+-----+-----+-----+
| Service | Type | File Size | Number of Files | Detail Level |
+-----+-----+-----+-----+-----+
| inventory | all | 100 MB | 10 | debug |
+-----+-----+-----+-----+-----+
Logging show

```

efa system restore

Restores the XCO system from a backup.

```
efa system restore [ --backup-tar filename ]
```

```
--backup-tar filename
```

Specifies the file name for the backup tar file.

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.

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 service enable

Enable and start a service.

```
efa system service enable [ --name { hyperv | vcenter | snmp | notification } ]
```

```
--name { hyperv | vcenter | snmp | notification }
```

Specifies the service to enable and start.

efa system service disable

Disable and stop a service.

```
efa system service disable [ --name { hyperv | vcenter | snmp | notification } ]
```

```
--name { hyperv | vcenter | snmp | notification }
```

Specifies the service to disable and stop.

efa system settings reset

Resets system settings to default values.

```
efa system settings reset [ --backup-schedule | -max-backup-files | --max-supportsave-files | --remote-server-settings | --periodic-device-config-backup ]
```

--backup-schedule

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

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**, **private-key**, **passphrase**, **certificate-key**, **ca-public-key**, and the **remote-server-protocol** setting to **scp**.

--periodic-device-config-backup

By specifying this flag, you turn off the backup.

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.

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.

```
efa system settings show
```

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         | 10.32.85.24 |
+-----+
| Remote Server Username   | user        |
+-----+
| Remote Server Password   | *****    |
+-----+
| Remote Server Directory  | /home/user  |
+-----+
| Remote Transfer Protocol | scp         |
+-----+
| Periodic Device Config Backup | Disabled   |
+-----+
| Show User Creation On Startup | true       |
+-----+
| Private Key               | *****    |
+-----+
| Certificate Key           | *****    |
+-----+
| CA Public Key             | *****    |
+-----+
| Passphrase                | *****    |
+-----+
```

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.

```
efa system settings update [--backup-schedule time |--max-backup-files
num |--max-supportsave-files num |--remote-server-ip ip/ipv6 |--
remote-transfer-protocol { ftp | scp } | --remote-server-username
name |--remote-server-password password |--remote-server-directory
directory |--periodic-device-config-backup { Enabled | Disabled } |
--private-key string | --passphrase string | --certificate-key string
| --ca-public-key string]
```

--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).

--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. IP address can be either IPv4 or IPv6. Example: 10.1.10.1 or 2000::1

--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.

--private-key *string*

Private Key file path to connect to remote server.

--passphrase *string*

Provide private key passphrase to connect to remote server.

--certificate-key *string*

Signed Certificate file path to connect to remote server.

--ca-public-key *string*

Public Key file path to connect to remote server.

The settings described here enable you to configure details of the remote server where the backup is copied. Use the [efa system settings reset](#) on page 218 command to remove settings.



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 also specify values for **all the others**:

- remote-server-ip
- remote-server-username
- remote-server-password
- remote-server-directory
- private-key
- certificate-key
- ca-public-key
- passphrase (optional)

If the above values have already been set, they are retained in the database. If you then modify any one of them, the other values are used.

```
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 ---
```

Configuring the remote settings with password using cli:

```
efa system settings update --remote-server-ip 10.32.85.5 --remote-server-username root --
remote-server-    directory /home/user/Downloads --remote-password password
```

Configuring the remote settings without password using cli:

```
efa system settings update -remote-server-password password --remote-server-ip
10.32.85.5 --remote-server-username root --remote-server-directory /home/user/Downloads --
certificate-key /home/user/id_rsa-cert.pub --private-key id_rsa --passphrase pkeyabc --
ca-publickey /home/user/ca.pub
```

efa system supportsave

Collects the system supportsave of the Inventory, Tenant, and Fabric service logs, and their associated database dumps.

```
efa system supportsave [--device-ip stringArray | --fabric-all | --  
fabric-name string ]
```

```
--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.

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.

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 system supportsave-list

Displays a list of all available EFA supportsave files.

```
efa system supportsave-list
```

This example shows two supportsave files available.

```
(efa:user)user@devqa-server:~$ efa system supportsave-list.  
  
Fetching supportsave-list of EFA...  
  
+-----+-----+-----+-----+  
| ID | File | Date |  
+-----+-----+-----+-----+
```

```
| 1 | efa_2022-06-30T17-36-57.278.logs.zip | 2022-06-30 |
+-----+-----+-----+-----+-----+
| 2 | efa_2022-06-30T17-44-09.225.logs.zip | 2022-06-30 |
+-----+-----+-----+-----+-----+
```

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.

```
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 ]
```

--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*

List of physical ports of devices which will be reserved as mirror destination ports.
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].

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.

```
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 ]
```

```
--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.

You can also use the API to set up debug level and debug for the module.

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.

```
efa tenant delete [ --name tenant-name | --force ]
```

```
--name tenant-name
```

Specifies a tenant and a tenant name.

```
--force
```

Forces the deletion on the tenant if the option is provided.

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.

```
efa tenant epg configure [ --name name | --tenant name ]
```

```
--name name
```

Specifies the name of the endpoint group.

```
--tenant name
```

Specifies the name of the tenant.

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.

```
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 | port-profile } | --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 |
  --dhcpv4-relay-address-ip ipv4 | --dhcpv6-relay-address-ip ipv6 |
  --dhcpv4-relay-gateway-ip ipv4 | --dhcpv4-relay-gateway-ip-interface
  ipv4 | --dhcpv6-relay-gateway-ip-interface ipv6 | --dhcpv4-relay-
  gateway-interface ipv4 | --dhcpv6-relay-gateway-interface ipv6 | --
  dhcpv6-relay-gateway-interface-ip ipv6 | --ip-icmp-redirect ctag:ip-
icmp-redirect | --ipv6-icmp-redirect ctag:ipv6-icmp-redirect | --
  help ]
```

--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** | **port-profile** }
- Configures the BGP service type. Valid values are `l3-hand-off`, `port-profile`, 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:l2-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:l2-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:false`.

--suppress-nd *value*

Sets suppress-nd flag to this network. The format is `ctag:suppress-nd`. Example: `1002:false`.

--pp-mac-acl-in *ext-mac-permit-any-mirror-acl*

Apply MAC ACL in ingress direction on ethernet / portchannel interfaces. The only supported ACL name is `ext-mac-permit-any-mirror-acl` and only supported ACL type is `extended`. Format `--pp-mac-acl-in <acl-name>` Example: `--pp-mac-acl-in ext-mac-permit-any-mirror-acl`.

--pp-mac-acl-out *ext-mac-permit-any-mirror-acl*

Apply MAC ACL in egress direction on ethernet / portchannel interfaces. The only supported ACL name is ext-mac-permit-any-mirror-acl and only supported ACL type is extended. Format --pp-mac-acl-out <acl-name>. Example: --pp-mac-acl-out ext-mac-permit-any-mirror-acl.

--pp-ip-acl-in *ext-ip-permit-any-mirror-acl*

Apply IP ACL in ingress direction on ethernet / portchannel interfaces. The only supported ACL name is ext-ip-permit-any-mirror-acl and only supported ACL type is extended. Format --pp-ip-acl-in <acl-name>. Example: --pp-ip-acl-in ext-ip-permit-any-mirror-acl.

--pp-ip-acl-out *ext-ip-permit-any-mirror-acl*

Apply IP ACL in egress direction on ethernet / portchannel interfaces. The only supported ACL name is ext-ip-permit-any-mirror-acl and only supported ACL type is extended. Format --pp-ip-acl-out <acl-name>. Example: --pp-ip-acl-out ext-ip-permit-any-mirror-acl.

--pp-ipv6-acl-in *ext-ipv6-permit-any-mirror-acl*

Apply IPv6 ACL in ingress direction on ethernet / portchannel interfaces. The only supported ACL name is ext-ipv6-permit-any-mirror-acl and only supported ACL type is extended. Format --pp-ipv6-acl-in <acl-name>. Example: --pp-ipv6-acl-in ext-ipv6-permit-any-mirror-acl.

--np-mac-acl-in *ctag:ext-mac-permit-any-mirror-acl*

Apply MAC ACL in ingress direction on vlan. The only supported ACL name is ext-mac-permit-any-mirror-acl and only supported ACL type is extended. Format --np-mac-acl-in <ctag:acl-name>. Example: --np-mac-acl-in 101:ext-mac-permit-any-mirror-acl.

--np-mac-acl-out *ctag:ext-mac-permit-any-mirror-acl*

Apply MAC ACL in egress direction on vlan. The only supported ACL name is ext-mac-permit-any-mirror-acl and only supported ACL type is extended. Format --np-mac-acl-out <ctag:acl-name>. Example: --np-mac-acl-out 101:ext-mac-permit-any-mirror-acl.

--np-ip-acl-in *ctag:ext-ip-permit-any-mirror-acl*

Apply IP ACL in ingress direction on ve interface. The only supported ACL name is ext-ip-permit-any-mirror-acl and only supported ACL type is extended. Format --np-ip-acl-in <ctag:acl-name>. Example: --np-ip-acl-in 101:ext-ip-permit-any-mirror-acl.

--np-ip-acl-out *ctag:ext-ip-permit-any-mirror-acl*

Apply IP ACL in egress direction on ve interface. The only supported ACL name is ext-ip-permit-any-mirror-acl and only supported ACL type is extended. Format --np-ip-acl-out <ctag:acl-name>. Example: --np-ip-acl-out 101:ext-ip-permit-any-mirror-acl.

--np-ipv6-acl-in *ctag:ext-ipv6-permit-any-mirror-acl*

Apply IPv6 ACL in ingress direction on ve interface. The only supported ACL name is ext-ipv6-permit-any-mirror-acl and only supported ACL type is extended. Format `--np-ipv6-acl-in <ctag:acl-name>`. Example: `--np-ipv6-acl-in 101:ext-ipv6-permit-any-mirror-acl`.

--dhcpv4-relay-address-ip *ipv4*

DHCP Server IPv4 Address

--dhcpv6-relay-address-ip *ipv6*

DHCP Server IPv6 Address

--dhcpv4-relay-gateway-ip *ipv4*

DHCP ipv4 relay gateway.

--dhcpv4-relay-gateway-ip-interface *ipv4*

DHCP ipv4 relay gateway ip interface.

--dhcpv6-relay-gateway-ip-interface *ipv6*

DHCP ipv6 relay gateway interface.

--dhcpv4-relay-gateway-interface *ipv4*

DHCP ipv4 relay gateway interface.

--dhcpv6-relay-gateway-interface *ipv6*

DHCP ipv6 relay gateway interface.

--dhcpv6-relay-gateway-interface-ip *ipv6*

DHCP ipv6 relay gateway interface ip.

--ip-icmp-redirect *ctag:ip-icmp-redirect*

Sets IPv4 icmp redirect flag in the format `ctag:icmp-redirect`. Example: `1002:true`.

--ipv6-icmp-redirect *ctag:ipv6-icmp-redirect*

Sets IPv6 icmpv6 redirect flag in the format `ctag:icmpv6-redirect`. Example: `1002:true`.

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.

XCO 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>`.

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 ten1vrf1 --ip-mtu 11:7900 --ip-mtu 12:8900
```

This example creates an endpoint group with ICMP redirect.

```
efa tenant epg create --tenant "t1" --name "epg1" --type extension
--switchport-mode trunk --single-homed-bfd-session-type auto
--po po1
--vrf vrf1 --ctag-range 19
--l3-vni 5001 --anycast-ip 19:3.33.3.3/24 --bridge-domain 19:Auto-BD-2 --ctag-description
"19:Tenant L3 Extended BD" --l2-vni 19:2 --ip-mtu 19:1600
--ip-icmp-redirect 19:true
--ipv6-icmp-redirect 19:true
```

efa tenant epg delete

Deletes an endpoint group.

```
efa tenant epg delete [ --force | --name epg-name | --tenant tenant-name ]
```

--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.

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 XCO to return an error. Using the `-force` option overrides this error and deletes the underlying endpoint group configuration on the device.

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.

```
efa tenant epg detach [ --tenant name | --source-epg epg-  
name | --destination-epg epg-name | --destination-epg-description | --  
destination-epg-ctag-range ctag-range ]
```

--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.

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 | | |
| Tagging | | | | Mode |
+-----+-----+-----+-----+-----+-----+
+-----+
| 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 |
| Mode | Tagging |
+-----+-----+-----+-----+-----+-----+
+-----+
| 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 |
| Mode | Tagging |
+-----+-----+-----+-----+-----+-----+
+-----+
| 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.

```
efa tenant epg error show [ --name epg-name | --tenant tenant-name | --summary ]
```

--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.

```

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.

```
efa tenant epg show [ --name epg-name | --detail | --tenant tenant-name | --vrf vrf-name ]
```

--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.

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
Network property flags		X
IPv6 ND prefix flags		X

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

Update an endpoint group.

```
efa tenant epg update [--name epg-name | --tenant tenant-name |
  --operation { port-group-add | port-group-delete | port-property-
  update | ctag-range-add | ctag-range-delete | vrf-add | vrf-delete
  | local-ip-add | local-ip-delete | anycast-ip-add | anycast-ip-
  delete | network-property-update | port-property-add | port-property-
  delete | port-property-update | network-property-add | network-
  property-delete | network-property-update } | --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 | --dhcpv4-relay-address-
  ip ipv4 | --dhcpv6-relay-address-ip ipv6 | --dhcpv4-relay-gateway-
  ip ipv4 | --dhcpv4-relay-gateway-ip-interface ipv4 | --dhcpv6-relay-
  gateway-ip-interface ipv6 | --dhcpv4-relay-gateway-interface ipv4
  | --dhcpv6-relay-gateway-interface ipv6 | --dhcpv6-relay-gateway-
  interface-ip ipv6 | --ip-icmp-redirect ctag:ip-icmp-redirect | --
  ipv6-icmp-redirect ctag:ipv6-icmp-redirect | --help ]
```

```
--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*.
- 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:l2-vni*.
- dhcpv4-relay-address-ip** *ipv4*
DHCP Server IPv4 Address
- dhcpv6-relay-address-ip** *ipv6*
DHCP Server IPv6 Address
- dhcpv4-relay-gateway-ip** *ipv4*
DHCP ipv4 relay gateway.
- dhcpv4-relay-gateway-ip-interface** *ipv4*
DHCP ipv4 relay gateway ip interface.
- dhcpv6-relay-gateway-ip-interface** *ipv6*
DHCP ipv6 relay gateway interface.
- dhcpv4-relay-gateway-interface** *ipv4*
DHCP ipv4 relay gateway interface.
- dhcpv6-relay-gateway-interface** *ipv6*
DHCP ipv6 relay gateway interface.
- dhcpv6-relay-gateway-interface-ip** *ipv6*
DHCP ipv6 relay gateway interface ip.
- ip-icmp-redirect** *ctag:ip-icmp-redirect*
Sets IPv4 icmp redirect flag in the format *ctag:icmp-redirect*. Example: 1002:true.
- ipv6-icmp-redirect** *ctag:ipv6-icmp-redirect*
Sets IPv6 icmpv6 redirect flag in the format *ctag:icmpv6-redirect*. Example: 1002:true.
- 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*.
- 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:l2-vni.

--np-mac-acl-in *ctag:ext-mac-permit-any-mirror-acl*

Apply MAC ACL for mirror action in ingress direction on vlan. The only supported ACL name is ext-mac-permit-any-mirror-acl and only supported ACL type is extended. Format --np-mac-acl-in <ctag:acl-name>. Example: --np-mac-acl-in 101:ext-mac-permit-any-mirror-acl.

--np-mac-acl-out *ctag:ext-mac-permit-any-mirror-acl*

Apply MAC ACL for mirror action in egress direction on vlan. The only supported ACL name is ext-mac-permit-any-mirror-acl and only supported ACL type is extended. Format --np-mac-acl-out <ctag:acl-name>. Example: --np-mac-acl-out 101:ext-mac-permit-any-mirror-acl.

--np-ip-acl-in *ctag:ext-ip-permit-any-mirror-acl*

Apply IP ACL for mirror action in ingress direction on ve interface. The only supported ACL name is ext-ip-permit-any-mirror-acl and only supported ACL type is extended. Format --np-ip-acl-in <ctag:acl-name>. Example: --np-ip-acl-in 101:ext-ip-permit-any-mirror-acl.

--np-ip-acl-out *ctag:ext-ip-permit-any-mirror-acl*

Apply IP ACL for mirror action in egress direction on ve interface. The only supported ACL name is ext-ip-permit-any-mirror-acl and only supported ACL type is extended. Format --np-ip-acl-out <ctag:acl-name>. Example: --np-ip-acl-out 101:ext-ip-permit-any-mirror-acl.

--np-ipv6-acl-in *ctag:ext-ipv6-permit-any-mirror-acl*

Apply IPv6 ACL for mirror action in ingress direction on ve interface. The only supported ACL name is ext-ipv6-permit-any-mirror-acl and only supported ACL type is extended. Format --np-ipv6-acl-in <ctag:acl-name>. Example: --np-ipv6-acl-in 101:ext-ipv6-permit-any-mirror-acl.

--name *epg-name*

Specifies the name of the endpoint group.

--operation { **port-group-add** | **port-group-delete** | **port-property-update** | **ctag-range-add** | **ctag-range-delete** | **vrf-add**, **vrf-delete** | **local-ip-add** | **local-ip-delete** | **anycast-ip-add** | **anycast-ip-delete** | **port-property-add** | **port-property-delete** | **port-property-update** | **network-property-add** | **network-property-delete** | **network-property-update** }

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.

--pp-mac-acl-in *ext-mac-permit-any-mirror-acl*

Apply MAC ACL for mirror action in ingress direction on ethernet / portchannel interfaces. The only supported ACL name is `ext-mac-permit-any-mirror-acl` and only supported ACL type is `extended`. Format `--pp-mac-acl-in <acl-name>`. Example: `--pp-mac-acl-in ext-mac-permit-any-mirror-acl`.

--pp-mac-acl-out *ext-mac-permit-any-mirror-acl*

Apply MAC ACL for mirror action in egress direction on ethernet / portchannel interfaces. The only supported ACL name is `ext-mac-permit-any-mirror-acl` and only supported ACL type is `extended`. Format `--pp-mac-acl-out <acl-name>`. Example: `--pp-mac-acl-out ext-mac-permit-any-mirror-acl`.

--pp-ip-acl-in *ext-ip-permit-any-mirror-acl*

Apply IP ACL for mirror action in ingress direction on ethernet / portchannel interfaces. The only supported ACL name is `ext-ip-permit-any-mirror-acl` and only supported ACL type is `extended`. Format `--pp-ip-acl-in <acl-name>`. Example: `--pp-ip-acl-in ext-ip-permit-any-mirror-acl`.

--pp-ip-acl-out *ext-ip-permit-any-mirror-acl*

Apply IP ACL for mirror action in egress direction on ethernet / portchannel interfaces. The only supported ACL name is `ext-ip-permit-any-mirror-acl` and only supported ACL type is `extended`. Format `--pp-ip-acl-out <acl-name>`. Example: `--pp-ip-acl-out ext-ip-permit-any-mirror-acl`.

--pp-ipv6-acl-in *ext-ipv6-permit-any-mirror-acl*

Apply IPv6 ACL for mirror action in ingress direction on ethernet / portchannel interfaces. The only supported ACL name is `ext-ipv6-permit-any-mirror-acl` and only supported ACL type is `extended`. Format `--pp-ipv6-acl-in <acl-name>`. Example: `--pp-ipv6-acl-in ext-ipv6-permit-any-mirror-acl`.

--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.

--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`.

--suppress-arp *array*

Sets `suppress-arp` flag to this network. Format **ctag:suppress-arp**. Example:

--suppress-arp 1002:true --suppress-arp 1003:false.

--suppress-nd *array*

Sets suppress-nd flag to this network. Format **ctag:suppress-arp**. Example: **--suppress-arp 1002:true --suppress-arp 1003:false**.

--tenant *tenant-name*

Specifies the name of the associated tenant.

--vrf *vrf-name*

Specifies the VRF to which these networks are attached.

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.

XCO 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>**.

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 tenlvrf1 --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 212 --anycast-ip 213:33.1.1.1/24 --anycast-ipv6 213:12::1/127 --ip-mtu
213:6990
--ip-icmp-redirect 213:true --ipv6-icmp-redirect 213:true
```

efa tenant execution

Displays the list of runs and event histories and deletes entries older than the specified number of days.

```
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 ]
```

```
--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.

```
efa tenant po configure [ --name po-name | --tenant name ]
```

```
--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 pol --tenant tenant11
```

```
PortChannel: pol configured successfully.
```

efa tenant po create

Creates a port channel for a tenant.

```
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 ]
```

```
--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.

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.

efa tenant po delete [**--name** *name* | **--force** | **--tenant** *name*]

--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.

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.

```
efa tenant po show [--name po-name | --tenant tenant-name | --detail ]
```

--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.

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 | Ports | 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.

```

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 ]

```

--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.

This example updates the min-link-count of an existing PO.

```
(efa:extreme)extreme@node-1:~$ efa tenant po update --name po11 --tenant tenant21
--operation min-link-count --min-link-count 1
```

```
PortChannel: po11 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 po11 --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.

efa tenant service bgp peer configure [**--name** *peer-name* | **--tenant** *tenant-name*]

--name *peer-name*

Specifies the name of the BGP peer instance.

--tenant *tenant-name*

Specifies the name of the tenant.

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.

efa tenant service bgp peer create [**--name** *peer-name* | **--tenant** *tenantname* | **--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-send-community string | --ipv4-uc-nbr-next-hop-self string | --ipv4-uc-nbr-remove-private-as | --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} | --ipv4-uc-nbr-prefix-list stringArray | --ipv6-uc-nbr-prefix-list stringArray | --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 | --ipv6-uc-nbr-send-community stringArray | --ipv4-uc-nbr-add-path-capability string | --ipv4-uc-nbr-add-path-advertise-all string | --ipv4-uc-nbr-add-path-advertise-group-best string | --ipv4-uc-nbr-add-path-advertise-best string | --ipv6-uc-nbr-add-path-capability string | --ipv6-uc-nbr-add-path-advertise-all string | --ipv6-uc-nbr-add-path-advertise-group-best string | --ipv6-uc-nbr-add-path-advertise-best string | --ipv4-uc-nbr-activate-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-default-originate-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-route-map-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-prefix-list-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-send-community-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af stringArray ]

```

--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-send-community *string*

Identifies the send-community in the following format: `device-ip,vrf-name:neighbor-ip,(all | both | extended | large | standard | large-and-standard | large-and-extended) .`

--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=='`

--ipv4-uc-nbr-prefix-list *stringArray*

IPv4 Unicast Neighbor prefix list in the format `device-ip,vrf-name:ipv4-neighbor,prefix-list-name,direction(in/out)`

--ipv6-uc-nbr-prefix-list *stringArray*

IPv6 Unicast Neighbor prefix list in the format device-ip,vrf-name:ipv6-neighbor,prefix-list-name,direction(in/out)

--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`.

--ipv6-uc-nbr-send-community *string*

Identifies the send-community in the following format: `device-ip,vrf-name:neighbor-ip,(all | both | extended | large | standard | large-and-standard | large-and-extended) .`

--ipv4-uc-nbr-add-path-capability *string*

Identifies the BGP neighbor add-path capability in the format `device-ip,vrf-name:ipv4-neighbor,add-path-capability`. Valid values are (send / receive / both).

--ipv4-uc-nbr-add-path-advertise-all *string*

Identifies the BGP neighbor add-path advertise all in the format `device-ip,vrf-name:ipv4-neighbor,add-path-advertise-all`(true/false).

--ipv4-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP neighbor add-path advertise group best in the format `device-ip,vrf-name:ipv4-neighbor,add-path-advertise-group-best`(true/false).

--ipv4-uc-nbr-add-path-advertise-best *string*

Identifies the BGP neighbor add-path advertise best in the format `device-ip,vrf-name:ipv4-neighbor,add-path-advertise-best`. Valid values are 2-16.

--ipv6-uc-nbr-add-path-capability *string*

Identifies the BGP neighbor add-path capability in the format `device-ip,vrf-name:ipv6-neighbor,add-path-capability`. Valid values are (send / receive / both).

--ipv6-uc-nbr-add-path-advertise-all *string*

Identifies the BGP neighbor add-path advertise all in the format `device-ip,vrf-name:ipv6-neighbor,add-path-advertise-all`(true/false).

--ipv6-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP neighbor add-path advertise group best in the format `device-ip,vrf-name:ipv6-neighbor,add-path-advertise-group-best`(true/false).

--ipv6-uc-nbr-add-path-advertise-best *string*

Identifies the BGP neighbor add-path advertise best in the format `device-ip,vrf-name:ipv6-neighbor,add-path-advertise-best`. Valid values are 2-16.

--ipv4-uc-nbr-activate-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor in the format device-ip,vrf-name:ipv4-neighbor,activate(true/false).

--ipv4-uc-nbr-default-originate-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor default originate in the format device-ip,vrf-name:ipv4-neighbor,default-originate(true/false).

--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor default originate route map in the format device-ip,vrf-name:ipv4-neighbor,route-map-name.

--ipv4-uc-nbr-route-map-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor prefix list in the format device-ip,vrf-name:ipv4-neighbor,prefix-list-name,direction(in/out).

--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor route map in the format device-ip,vrf-name:ipv4-neighbor,route-map-name,direction(in/out).

--ipv4-uc-nbr-send-community-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor send-community in the format device-ip,vrf-name:ipv4-neighbor,send-community. Valid values are (all - large, extended and standard / both - extended and standard / large / extended / standard / large-and-standard - large and standard / large-and-extended - large and extended).

--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af *stringArray*

BGP neighbor add-path capability in the format device-ip,vrf-name:ipv4-neighbor,add-path-capability. Valid values are (send / receive / both).

--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af *stringArray*

BGP neighbor add-path advertise all in the format device-ip,vrf-name:ipv4-neighbor,add-path-advertise-all(true/false).

--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af *stringArray*

BGP neighbor add-path advertise best in the format device-ip,vrf-name:ipv4-neighbor,add-path-advertise-best. Valid values are 2-16.

--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af *stringArray*

BGP neighbor add-path advertise group best in the format device-ip,vrf-name:ipv4-neighbor,add-path-advertise-group-best(true/false).

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 [ExtremeCloud Orchestrator Configuration Guide, 3.7.0](#).

```
$ 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,tenlvrf1:10.20.30.40,true
--ipv4-uc-nbr-default-originate 10.20.246.16,tenlvrf1:10.20.30.40,true
--ipv4-uc-nbr-default-originate-route-map 10.20.246.16,tenlvrf1:10.20.30.40,rmap1
```

```
efa tenant service bgp peer create --name tenlbgppeer1 --tenant ten1
--ipv4-uc-nbr-bfd 10.20.246.15,tenlvrf1:25.1.1.2,true
--ipv6-uc-nbr 10.20.246.15,tenlvrf1:25:1::3,5900
--ipv6-uc-nbr-bfd 10.20.246.15,tenlvrf1:25:1::3,true
--ipv6-uc-nbr-md5-password 10.20.246.15,tenlvrf1:25:1::3,password
--ipv4-uc-nbr-md5-password 10.20.246.15,tenlvrf1:25.1.1.2,password
--ipv4-uc-nbr-send-community 10.20.246.15,tenlvrf1:25.1.1.2,all
--ipv6-uc-nbr-send-community 10.20.246.15,tenlvrf1:1::1,both
--ipv4-uc-nbr-add-path-capability 10.20.246.15,tenlvrf1:25.1.1.2,send
--ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,tenlvrf1:25.1.1.2,true
--ipv4-uc-nbr 10.20.246.15,tenlvrf1:25.1.1.2,5900
--ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,tenlvrf1:25.1.1.2,true
--ipv6-uc-nbr 10.20.246.15,tenlvrf1:1::1,5900
--ipv6-uc-nbr-add-path-advertise-best 10.20.246.15,tenlvrf1:1::1,2
--ipv6-uc-nbr-add-path-capability 10.20.246.15,tenlvrf1:1::1,both
--ipv6-uc-nbr-add-path-advertise-best 10.20.246.15,tenlvrf1:1::1,2
--ipv6-uc-nbr-add-path-advertise-group-best 10.20.246.15,tenlvrf1:1::1,true
```

```
(efa:root)root@Server41:~# efa tenant service bgp peer create --name customer_1 --tenant
tv3 \
--ipv4-uc-nbr 10.20.49.119,v1:10.10.10.11,95001 \ --ipv4-uc-nbr-md5-password
10.20.49.119,v1:10.10.10.11,password \ --ipv4-uc-nbr-bfd
10.20.49.119,v1:10.10.10.11,true,50,5000,50 \
--ipv4-uc-nbr-next-hop-self 10.20.49.119,v1:10.10.10.11,always \ --ipv4-uc-nbr-update-
source-ip
10.20.49.119,v1:10.10.10.11,10.11.12.13 \ --ipv4-uc-nbr-remove-private-as
10.20.49.119,v1:10.10.10.11,true \ --ipv4-uc-nbr-default-originate
10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate-route-map 10.20.49.119,v1:10.10.10.11,rt1 \ --ipv4-uc-nbr-
route-map
10.20.49.119,v1:10.10.10.11,customer_1_v4_in,in \ --ipv4-uc-nbr-prefix-list
10.20.49.119,v1:10.10.10.11,customer_1_v4_out,out \ --ipv4-uc-nbr-send-community
10.20.49.119,v1:10.10.10.11,all \ --ipv4-uc-nbr-add-path-capability
10.20.49.119,v1:10.10.10.11,both \
--ipv4-uc-nbr-add-path-advertise-all 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-add-path-advertise-best 10.20.49.119,v1:10.10.10.11,16 \
--ipv4-uc-nbr-add-path-advertise-group-best 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-activate-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,rt1 \
--ipv4-uc-nbr-send-community-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,all \
--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,both \
--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,16 \
--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af
10.20.49.119,v1:10.10.10.11,true \
--ipv4-uc-nbr-route-map-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,customer_1_v6_in,in \
--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af 10.20.49.119,v1:10.10.10.11,customer_1_v6_out,out
(efa:root)root@Server41:~# efa tenant service bgp peer update --name customer_1 --tenant
tv3
--operation peer-add \ --ipv4-uc-nbr 10.20.49.118,v1:10.10.10.12,95001 \ --ipv4-uc-nbr-
md5-password
10.20.49.118,v1:10.10.10.12,password \ --ipv4-uc-nbr-bfd
10.20.49.118,v1:10.10.10.12,true,50,5000,50 \
--ipv4-uc-nbr-next-hop-self 10.20.49.118,v1:10.10.10.12,always \ --ipv4-uc-nbr-update-
source-ip
10.20.49.118,v1:10.10.10.12,10.11.12.12 \ --ipv4-uc-nbr-remove-private-as
10.20.49.118,v1:10.10.10.12,true \ --ipv4-uc-nbr-default-originate
10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate-route-map 10.20.49.118,v1:10.10.10.12,rt1 \ --ipv4-uc-nbr-
```



```

route-map
10.20.49.118,v1:10.10.10.12,customer_1_v4_in,in \ --ipv4-uc-nbr-prefix-list
10.20.49.118,v1:10.10.10.12,customer_1_v4_out,out \ --ipv4-uc-nbr-send-community
10.20.49.118,v1:10.10.10.12,all \ --ipv4-uc-nbr-add-path-capability
10.20.49.118,v1:10.10.10.12,both \
--ipv4-uc-nbr-add-path-advertise-all 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-add-path-advertise-best 10.20.49.118,v1:10.10.10.12,16 \
--ipv4-uc-nbr-add-path-advertise-group-best 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-activate-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,rt1 \
--ipv4-uc-nbr-send-community-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,all \
--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,both \
--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,16 \
--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af
10.20.49.118,v1:10.10.10.12,true \
--ipv4-uc-nbr-route-map-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,customer_1_v6_in,in \
--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af 10.20.49.118,v1:10.10.10.12,customer_1_v6_out,out

```

```
(efa:root)root@Server41:~# efa tenant service bgp peer show --detail
```

```

=====
===== Name : customer_1 Tenant : tv3 State : bgp-peer-created Description : Static
Peer ----- Device IP : 10.20.49.119 VRF : v1 AFI : ipv4 SAFI : unicast Remote IP :
10.10.10.11 Remote ASN : 95001 Activate : true Next Hop Self : always Update Source IP :
10.11.12.13 BFD Enabled : true BFD Interval : 50 BFD Rx : 5000 BFD Multiplier : 50 MD5
Password : $9$MCgKGaNt6OASX68/7TC6Lw== Remove Private AS : true Default
Originate : true Default Originate Route Map : rt1 Send Community : all Prefix List In :
Prefix List Out : customer_1_v4_out Route Map In : customer_1_v4_in Route Map Out :
Add Path Capability : both Add Path Advertise : All, Group Best, Best 16 Dev State :
provisioned App State : cfg-in-sync --Multi protocol capability-- AFI : ipv6 SAFI : unicast
Activate : true Default Originate : true Default Originate Route Map : rt1 Send
Community : all Prefix List In : Prefix List Out : customer_1_v6_out Route Map In :
customer_1_v6_in Route Map Out : Add Path Capability : both Add Path Advertise : All,
Group Best, Best 16 Dev State : provisioned App State : cfg-in-sync Device IP :
10.20.49.118 VRF : v1 AFI : ipv4 SAFI : unicast Remote IP : 10.10.10.12 Remote ASN : 95001
Activate : true Next Hop Self : always Update Source IP : 10.11.12.12 BFD Enabled : true
BFD Interval : 50 BFD Rx : 5000 BFD Multiplier : 50 MD5 Password :
$9$MCgKGaNt6OASX68/7TC6Lw== Remove Private AS : true Default Originate : true
Default Originate Route Map : rt1 Send Community : all Prefix List In : Prefix List Out :
customer_1_v4_out Route Map In : customer_1_v4_in Route Map Out : Add Path
Capability : both Add Path Advertise : All, Group Best, Best 16 Dev State : provisioned
App State : cfg-in-sync --Multi protocol capability-- AFI : ipv6 SAFI : unicast Activate :
true Default Originate : true Default Originate Route Map : rt1 Send Community : all
Prefix List In : Prefix List Out : customer_1_v6_out Route Map In : customer_1_v6_in
Route Map Out : Add Path Capability : both Add Path Advertise : All, Group Best, Best 16
Dev State : provisioned App State : cfg-in-sync Dynamic Peer ----- 0 Records
=====
=====

```

efa tenant service bgp peer delete

Deletes the BGP neighbors for a given VRF on the fabric device.

```
efa tenant service bgp peer delete [ --force | --name service-name | --tenant tenant-name ]
```

--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.

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.

```
efa tenant service bgp peer operational show [--tenant name | --vrf name ]
```

--tenant *name*

Specifies the name of the tenant.

--vrf *name*

Specifies the name of the VRF.

```
efa tenant service bgp peer operational show --tenant tenant2 --vrf vrf101
+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+
| 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 | 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      |
+-----+
+-----+
| 10.20.246.14 | default-vrf |
| 64512        | 66000       | 172.31.254.239 | 10.10.10.12
|          | ipv4      | unicast  | ESTAB      |
+-----+
+-----+
| 10.20.246.14 | default-vrf |
| 64512        | 66000       | 172.31.254.239 | 10.10.10.14
|          | ipv4      | unicast  | ESTAB      |
+-----+
+-----+
| 10.20.246.14 | default-vrf |
| 64512        | 65000       | 172.31.254.239 | 10.10.10.0
|          | l2vpn     | evpn     | ESTAB      |
+-----+
+-----+
| 10.20.246.14 | default-vrf |
| 64512        | 65000       | 172.31.254.239 | 10.10.10.2
|          | l2vpn     | evpn     | ESTAB      |
+-----+
+-----+
| 10.20.246.14 | default-vrf |
| 64512        | 66000       | 172.31.254.239 | 10.10.10.4
|          | l2vpn     | evpn     | ESTAB      |
+-----+
+-----+
| 10.20.246.14 | default-vrf |
| 64512        | 66000       | 172.31.254.239 | 10.10.10.6
|          | 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.

efa tenant service bgp peer show [**--name** *peer-name* | **--tenant** *tenant-name* | **--detail**]

--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.

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 | Activate | 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 | Activate | Next Hop |
| | | | | | | Self |
+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv4 | unicast | 1.1.1.11 | 95001 | true | always |
+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv6 | unicast | 20a1:a::10 | 95001 | true | 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     : pgl
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 | Activate | Next Hop |
|           |       |       |       |           |           |         | Self     |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv4 | unicast | 1.1.1.11 | 95001 | true | always |
+-----+-----+-----+-----+-----+-----+-----+-----+
| 10.20.216.16 | blue11 | ipv6 | unicast | 20a1:a::10 | 95001 | true | 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 |
| Listen | Peer Group | Dev State | App State | | Limit |
| | | | | | |
+-----+-----+-----+-----+-----+-----+-----+
+-----+
Dynamic Peer Details
=====
Name          : B2
Tenant        : tenant11
State         : bs-state-created

+-----+-----+-----+-----+-----+-----+-----+
| Device IP | VRF | AFI | SAFI | Remote IP | Remote ASN | Activate | 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.

```

efa tenant service bgp peer update [ --name peer-name | --operation
  {peer-add | peer-delete | desc-update} | --tenant tenant-name | --description
  description | --ipv4-uc-dyn-nbr string | --ipv4-uc-nbr string |
  --ipv4-uc-nbr-bfd string | --ipv4-uc-nbr-send-community 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 | --md5-password-prompt-enable
{ true | false } | --ipv4-uc-nbr-prefix-list stringArray |
--ipv6-uc-nbr-prefix-list stringArray | --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 | --ipv6-uc-nbr-send-community stringArray
| | --ipv4-uc-nbr-add-path-capability string | --ipv4-uc-nbr-add-
path-advertise-all string | --ipv4-uc-nbr-add-path-advertise-group-
best string | --ipv4-uc-nbr-add-path-advertise-best string | --ipv6-
uc-nbr-add-path-capability string | --ipv6-uc-nbr-add-path-advertise-
all string | --ipv6-uc-nbr-add-path-advertise-group-best string
| --ipv6-uc-nbr-add-path-advertise-best string | --ipv4-uc-nbr-
activate-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-default-originate-
in-ipv6-uc-af stringArray | --ipv4-uc-nbr-default-originate-route-
map-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-route-map-in-ipv6-uc-
af stringArray | --ipv4-uc-nbr-prefix-list-in-ipv6-uc-af stringArray
| --ipv4-uc-nbr-send-community-in-ipv6-uc-af stringArray | --ipv4-
uc-nbr-add-path-capability-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-
add-path-advertise-all-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-add-
path-advertise-best-in-ipv6-uc-af stringArray | --ipv4-uc-nbr-add-
path-advertise-group-best-in-ipv6-uc-af stringArray ]

```

--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-send-community *string*

Updates the send-community in the following format: device-ip,vrf-name:neighbor-ip, (**all** | **both** | **extended** | **large** | **standard** | **large-and-standard** | **large-and-extended**).

--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=='`

--ipv4-uc-nbr-prefix-list *stringArray*

IPv4 Unicast Neighbor prefix list in the format device-ip,vrf-name:ipv4-neighbor,prefix-list-name,direction(in/out)

--ipv6-uc-nbr-prefix-list *stringArray*

IPv6 Unicast Neighbor prefix list in the format device-ip,vrf-name:ipv6-neighbor,prefix-list-name,direction(in/out)

--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*.

--ipv6-us-nbr-send-community *stringArray*

Updates the send-community in the following format: device-ip,vrf-name:neighbor-ip, (**all** | **both** | **extended** | **large** | **standard** | **large-and-standard** | **large-and-extended**).

--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, listenlimit. The listen-limit value is optional.

--ipv6-uc-nbr *string*

Identifies the IPv6 unicast neighbor with a string in the following format: device-ip, vrfname: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, bfdenable (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, vrfname: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 default originate route map in the following format: device-ip, vrf-name:ipv6-neighbor, route-map-name.

Example for encrypted/password containing special characters: `--ipv6-uc-nbr-md5-password device-ip, vrf-name:ipv6neighbor, '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,vrfname: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:ipv6neighbor,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`.

--ipv4-uc-nbr-add-path-capability *string*

Identifies the BGP neighbor add-path capability in the format `device-ip,vrf-name:ipv4-neighbor,add-path-capability`. Valid values are (send / receive / both).

--ipv4-uc-nbr-add-path-advertise-all *string*

Identifies the BGP neighbor add-path advertise all in the format `device-ip,vrf-name:ipv4-neighbor,add-path-advertise-all(true/false)`.

--ipv4-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP neighbor add-path advertise group best in the format `device-ip,vrf-name:ipv4-neighbor,add-path-advertise-group-best(true/false)`.

--ipv4-uc-nbr-add-path-advertise-best *string*

Identifies the BGP neighbor add-path advertise best in the format `device-ip,vrf-name:ipv4-neighbor,add-path-advertise-best`. Valid values are 2-16.

--ipv6-uc-nbr-add-path-capability *string*

Identifies the BGP neighbor add-path capability in the format `device-ip,vrf-name:ipv6-neighbor,add-path-capability`. Valid values are (send / receive / both).

--ipv6-uc-nbr-add-path-advertise-all *string*

Identifies the BGP neighbor add-path advertise all in the format `device-ip,vrf-name:ipv6-neighbor,add-path-advertise-all(true/false)`.

--ipv6-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP neighbor add-path advertise group best in the format `device-ip,vrf-name:ipv6-neighbor,add-path-advertise-group-best(true/false)`.

--ipv6-uc-nbr-add-path-advertise-best *string*

Identifies the BGP neighbor add-path advertise best in the format `device-ip,vrf-name:ipv6-neighbor,add-path-advertise-best`. Valid values are 2-16.

--ipv4-uc-nbr-activate-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor in the format device-ip,vrf-name:ipv4-neighbor,activate(true/false).

--ipv4-uc-nbr-default-originate-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor default originate in the format device-ip,vrf-name:ipv4-neighbor,default-originate(true/false).

--ipv4-uc-nbr-default-originate-route-map-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor default originate route map in the format device-ip,vrf-name:ipv4-neighbor,route-map-name.

--ipv4-uc-nbr-route-map-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor prefix list in the format device-ip,vrf-name:ipv4-neighbor,prefix-list-name,direction(in/out).

--ipv4-uc-nbr-prefix-list-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor route map in the format device-ip,vrf-name:ipv4-neighbor,route-map-name,direction(in/out).

--ipv4-uc-nbr-send-community-in-ipv6-uc-af *stringArray*

IPv4 Unicast Neighbor send-community in the format device-ip,vrf-name:ipv4-neighbor,send-community. Valid values are (all - large, extended and standard / both - extended and standard / large / extended / standard / large-and-standard - large and standard / large-and-extended - large and extended).

--ipv4-uc-nbr-add-path-capability-in-ipv6-uc-af *stringArray*

BGP neighbor add-path capability in the format device-ip,vrf-name:ipv4-neighbor,add-path-capability. Valid values are (send / receive / both).

--ipv4-uc-nbr-add-path-advertise-all-in-ipv6-uc-af *stringArray*

BGP neighbor add-path advertise all in the format device-ip,vrf-name:ipv4-neighbor,add-path-advertise-all(true/false).

--ipv4-uc-nbr-add-path-advertise-best-in-ipv6-uc-af *stringArray*

BGP neighbor add-path advertise best in the format device-ip,vrf-name:ipv4-neighbor,add-path-advertise-best. Valid values are 2-16.

--ipv4-uc-nbr-add-path-advertise-group-best-in-ipv6-uc-af *stringArray*

BGP neighbor add-path advertise group best in the format device-ip,vrf-name:ipv4-neighbor,add-path-advertise-group-best(true/false).

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 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-add-path-capability 10.20.246.15,ten1vrf1:10.20.30.50,send
--ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,ten1vrf1:10.20.30.50,true
--ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,ten1vrf1:10.20.30.50,8
.....
--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-add-path-capability 10.20.246.16,ten1vrf1:10.20.30.50,receive
--ipv4-uc-nbr-add-path-advertise-all 10.20.246.16,ten1vrf1:10.20.30.50,false
--ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.16,ten1vrf1:10.20.30.50,false
--ipv4-uc-nbr-add-path-advertise-best 10.20.246.16,ten1vrf1:10.20.30.50,4
```

efa tenant service bgp peer-group configure

Pushes or removes a partial configuration for a peer-group.

```
efa tenant service bgp peer-group configure [ --name peer-group-name |
--tenant tenant-name ]
```

--name *peer-group-name*

Specifies the name of the BGP peer group instance.

--tenant *tenant-name*

Specifies the name of the tenant.

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.

```
efa tenant service bgp peer-group create [ --name peer-group-name |
--tenant tenant-name | --description string | --pg-name string |
--pg-asn string | --pg-bfd string | --pg-bfd-enable string | --
pg-remove-private-as stringArray | --pg-next-hop-self string | --
pg-update-source-ip string | --pg-md5-password string | --pg-md5-
password-prompt-enable { true | false } | --pg-ipv4-uc-nbr-prefix-
list string | --pg-ipv4-uc-nbr-route-map string | --pg-ipv4-uc-nbr-
send-community string | --pg-ipv6-uc-nbr-prefix-list string | --pg-
ipv6-uc-nbr-route-map string | --pg-ipv6-uc-nbr-send-community string
| --pg-ipv4-uc-nbr-add-path-capability string | --pg-ipv4-uc-nbr-
add-path-advertise-all string | --pg-ipv4-uc-nbr-add-path-advertise-
group-best string | --pg-ipv4-uc-nbr-add-path-advertise-best string
| --pg-ipv6-uc-nbr-add-path-capability string | --pg-ipv6-uc-nbr-
add-path-advertise-all string | --pg-ipv6-uc-nbr-add-path-advertise-
group-best string | --pg-ipv6-uc-nbr-add-path-advertise-best string |
--pg-ipv6-uc-nbr-activate stringArray ]
```

--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.

--pg-ipv4-uc-nbr-prefix-list *string*

Identifies the IPv4 unicast neighbor with a string in the following format: `device-ip,pgname:prefix-list-name,direction(in | out)`.

--pg-ipv4-uc-nbr-route-map *string*

Identifies the IPv4 unicast neighbor with a string in the following format: `device-ip,pgname:route-map-name,direction(in | out)`.

--pg-ipv4-uc-nbr-send-community *string*

Identifies the IPv4 send-community in the following format: `device-ip,pgname:send-community(all | both | extended | large | standard | large-and-standard | large-and-extended)`.

--pg-ipv6-uc-nbr-prefix-list *string*

Identifies the IPv6 unicast neighbor with a string in the following format: `device-ip,pgname:prefix-list-name,direction(in | out)`.

--pg-ipv6-uc-nbr-route-map *string*

Identifies the IPv6 unicast neighbor with a string in the following format: `device-ip,pgname:route-map-name,direction(in | out)`.

--pg-ipv6-uc-nbr-send-community *string*

Identifies the IPv6 send-community in the following format: device-ip,pgname:send-community (**all** | **both** | **extended** | **large** | **standard** | **large-and-standard** | **large-and-extended**).

--pg-ipv4-uc-nbr-add-path-capability *string*

Identifies the BGP peer-group add-path capability in the format device-ip,peer-group-name:add-path-capability. Valid values are (send / receive / both).

--pg-ipv4-uc-nbr-add-path-advertise-all *string*

Identifies the BGP peer-group add-path advertise all in the format device-ip,peer-group-name:add-path-advertise-all(true/false).

--pg-ipv4-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP peer-group add-path advertise group best in the format device-ip,peer-group-name:add-path-advertise-group-best(true/false).

--pg-ipv4-uc-nbr-add-path-advertise-best *string*

Identifies the BGP peer-group add-path advertise best in the format device-ip,peer-group-name:add-path-advertise-best. Valid values are 2-16.

--pg-ipv6-uc-nbr-add-path-capability *string*

Identifies the BGP peer-group add-path capability in the format device-ip,peer-group-name:add-path-capability. Valid values are (send / receive / both).

--pg-ipv6-uc-nbr-add-path-advertise-all *string*

Identifies the BGP peer-group add-path advertise best in the format device-ip,peer-group-name:add-path-advertise-best. Valid values are 2-16.

--pg-ipv6-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP peer-group add-path advertise group best in the format device-ip,peer-group-name:add-path-advertise-group-best(true/false).

--pg-ipv6-uc-nbr-add-path-advertise-best *string*

Identifies the BGP peer-group add-path advertise best in the format device-ip,peer-group-name:add-path-advertise-best. Valid values are 2-16.

--pg-ipv6-uc-nbr-activate *stringArray*

BGP peer-group Activate in the format device-ip,peer-group-name:activate(true/false).

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 [ExtremeCloud Orchestrator Configuration Guide, 3.7.0](#).

```
$ 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 create --name ten1bgppg1 --tenant ten1
--pg-name 10.20.246.15:pg1
--pg-bfd-enable 10.20.246.15,pg1:true
--pg-ipv4-uc-nbr-add-path-capability 10.20.246.15,pg1:both
--pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,pg1:true
--pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,pg1:true
--pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,pg1:10
--pg-ipv6-uc-nbr-add-path-capability 10.20.246.15,pg1:both
--pg-ipv6-uc-nbr-add-path-advertise-all 10.20.246.15,pg1:true
--pg-ipv6-uc-nbr-add-path-advertise-group-best 10.20.246.15,pg1:true
--pg-ipv6-uc-nbr-add-path-advertise-best 10.20.246.15,pg1:10
```

```
efa tenant service bgp peer-group create --tenant "t1" --name "vs"
--pg-name 10.20.246.30:v1 --pg-asn 10.20.246.30,v1:5200 --pg-ipv6-uc-nbr-activate
10.20.246.30,v1:true
--pg-ipv4-uc-nbr-route-map 10.20.246.30,v1:customer_1_in,in --pg-ipv4-uc-nbr-prefix-list
10.20.246.30,v1:customer_1_in,in
--pg-name 10.20.246.30:v2 --pg-asn 10.20.246.30,v2:5201
--pg-name 10.20.246.30:v3 --pg-asn 10.20.246.30,v3:5203
--pg-name 10.20.246.30:v4 --pg-asn 10.20.246.30,v4:5204
```

```
--pg-name 10.20.246.29:v1 --pg-asn 10.20.246.29,v1:5200 --pg-ipv6-uc-nbr-activate
10.20.246.29,v1:true
--pg-ipv4-uc-nbr-route-map 10.20.246.29,v1:customer_1_in,in --pg-ipv4-uc-nbr-prefix-list
10.20.246.29,v1:customer_1_in,in
--pg-name 10.20.246.29:v2 --pg-asn 10.20.246.29,v2:5201
--pg-name 10.20.246.29:v3 --pg-asn 10.20.246.29,v3:5203
--pg-name 10.20.246.29:v4 --pg-asn 10.20.246.29,v4:5204
```

```
efa tenant service bgp peer create --tenant "t1" --name "vs"
--ipv4-uc-dyn-nbr 10.20.246.29,vs:15.16.16.0/28,v1,20 --ipv4-uc-dyn-nbr
10.20.246.29,vs:15.16.17.0/28,v3,20
--ipv4-uc-dyn-nbr 10.20.246.29,vs:15.16.18.0/28,v4,20 --ipv6-uc-dyn-nbr
10.20.246.29,vs:14::/127,v2,10
```

```
--ipv6-uc-dyn-nbr 10.20.246.29,vs:15::/127,v3,10 --ipv4-uc-dyn-nbr
10.20.246.30,vs:15.16.16.0/28,v1,20
--ipv4-uc-dyn-nbr 10.20.246.30,vs:15.16.17.0/28,v3,20 --ipv4-uc-dyn-nbr
10.20.246.30,vs:15.16.18.0/28,v4,20
--ipv6-uc-dyn-nbr 10.20.246.30,vs:14::/127,v2,10 --ipv6-uc-dyn-nbr
10.20.246.30,vs:15::/127,v3,10
```

This example shows the use of the `--pg-update-source-ip` flag with IPv6 addressing.

```
efa tenant service bgp peer-group create --tenant "tenant1" --name "bgpPeerGroup2"
--pg-name 10.20.246.15:bgppg2 --pg-asn 10.20.246.15,bgppg2:100 --pg-
update-source-ip 10.20.246.15,bgppg2:10::10
```

```
efa tenant service bgp peer-group update --tenant "tenant1" --name "bgpPeerGroup2" --
operation peer-group-add
--pg-name 10.20.246.15:bgppg1 --pg-asn 10.20.246.15,bgppg1:200 --pg-
update-source-ip 10.20.246.15,bgppg1:10::20
```

```
efa tenant service bgp peer-group show --tenant tenant1 -detail
```

```

=====
Name           : bgpPeerGroup2
Tenant        : tenant1
State         : bgp-pg-created
Description    :

Peer Group
-----
Device IP      : 10.20.246.15
Peer Group    : bgppg2
Remote ASN     : 100
Next Hop Self  : false
Update Source IP : 10::10
BFD Enabled    : false
BFD Interval   :
BFD Rx        :
BFD Multiplier :
MD5 Password   :
Remove Private AS : false
Activate      : Activate (afi)
                true (ipv4)
                false (ipv6)

Prefix List In :
Prefix List Out :
Route Map In    :
Route Map Out   :
Send Community  :
Add Path Capability :
Add Path Advertise :
Dev State       : provisioned
App State       : cfg-in-sync

Device IP      : 10.20.246.15
Peer Group    : bgppg1
Remote ASN     : 200
Next Hop Self  : false
Update Source IP : 10::20
BFD Enabled    : false
BFD Interval   :
BFD Rx        :
BFD Multiplier :
MD5 Password   :
Remove Private AS : false
Activate      : Activate (afi)
                true (ipv4)
                false (ipv6)

Prefix List In :
Prefix List Out :
Route Map In    :
Route Map Out   :
Send Community  :
Add Path Capability :
Add Path Advertise :
Dev State       : provisioned
App State       : cfg-in-sync
=====

```

efa tenant service bgp peer-group delete

Deletes the specified BGP peer group.

```
efa tenant service bgp peer-group delete [ --force | --name peer-group-name | --tenant tenant-name ]
```

--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.

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.

```
efa tenant service bgp peer-group show [ --name peer-group-name | --
tenant tenant-name | --detail ]
```

--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.

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 for a specified tenant.

```

efa tenant service bgp peer-group show --tenant tenant1

=====
Name : ten1BgpPG1
Tenant : tenant1
State : bgp-pg-created
+-----+-----+-----+-----+-----+-----+
| Device IP | Peer Group | Remote | Next Hop | Update | BFD |
|           |            | ASN   | Self    | Source IP | Enabled |
|           | BFD      | Dev State | App State |
| [Interval,Rx,Multiplier] |          |          |
+-----+-----+-----+-----+-----+-----+
| 10.20.246.15 | pg7      | 6001 | true    | 10.20.30.40 | true |
|           | provisioned | cfg-in-sync |
+-----+-----+-----+-----+-----+-----+
BGP PeerGroup Details
    
```

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.

```

efa tenant service bgp peer-group update [ --name peer-group-name |
--tenant tenant-name | --description description | --operation { peer-
group-add | peer-group-delete | desc-update } | --pg-name string
| --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-next-hop-self string
| --pg-update-source-ip string | --pg-ipv4-uc-nbr-prefix-list string
| --pg-ipv4-uc-nbr-route-map string | --pg-ipv6-uc-nbr-prefix-list
string | --pg-ipv6-uc-nbr-route-map string | --pg-ipv4-uc-nbr-send-
community string | --pg-ipv6-uc-nbr-send-community string | --pg-
ipv4-uc-nbr-add-path-capability string | --pg-ipv4-uc-nbr-add-path-
advertise-all string | --pg-ipv4-uc-nbr-add-path-advertise-group-best
string | --pg-ipv4-uc-nbr-add-path-advertise-best string | --pg-
ipv6-uc-nbr-add-path-capability string | --pg-ipv6-uc-nbr-add-path-
advertise-all string | --pg-ipv6-uc-nbr-add-path-advertise-group-best

```

- string* | **--pg-ipv6-uc-nbr-add-path-advertise-best** *string* | **--pg-ipv6-uc-nbr-activate** *stringArray*]
- 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.

--pg-ipv4-uc-nbr-prefix-list *string*

Identifies the IPv4 unicast neighbor with a string in the following format: device-ip,pgname:prefix-list-name,direction(in | out).

--pg-ipv4-uc-nbr-route-map *string*

Identifies the IPv4 unicast neighbor with a string in the following format: device-ip,pgname:route-map-name,direction(in | out).

--pg-ipv6-uc-nbr-prefix-list *string*

Identifies the IPv6 unicast neighbor with a string in the following format: device-ip,pgname:prefix-list-name,direction(in | out).

--pg-ipv6-uc-nbr-route-map *string*

Identifies the IPv6 unicast neighbor with a string in the following format: device-ip,pgname:route-map-name,direction(in | out).

--pg-ipv4-uc-nbr-send-community *stringArray*

Identifies the IPv4 send-community in the following format: device-ip,pgname:send-community (**all** | **both** | **extended** | **large** | **standard** | **large-and-standard** | **large-and-extended**).

--pg-ipv6-uc-nbr-send-community *stringArray*

Identifies the IPv6 send-community in the following format: device-ip,pgname:send-community (**all** | **both** | **extended** | **large** | **standard** | **large-and-standard** | **large-and-extended**).

--pg-ipv4-uc-nbr-add-path-capability *string*

Identifies the BGP peer-group add-path capability in the format device-ip,peer-group-name:add-path-capability. Valid values are (send / receive / both).

--pg-ipv4-uc-nbr-add-path-advertise-all *string*

Identifies the BGP peer-group add-path advertise all in the format device-ip,peer-group-name:add-path-advertise-all(true/false).

--pg-ipv4-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP peer-group add-path advertise group best in the format device-ip,peer-group-name:add-path-advertise-group-best(true/false).

--pg-ipv4-uc-nbr-add-path-advertise-best *string*

Identifies the BGP peer-group add-path advertise best in the format device-ip,peer-group-name:add-path-advertise-best. Valid values are 2-16.

--pg-ipv6-uc-nbr-add-path-capability *string*

Identifies the BGP peer-group add-path capability in the format device-ip,peer-group-name:add-path-capability. Valid values are (send / receive / both).

--pg-ipv6-uc-nbr-add-path-advertise-all *string*

Identifies the BGP peer-group add-path advertise all in the format device-ip,peer-group-name:add-path-advertise-all(true/false).

--pg-ipv6-uc-nbr-add-path-advertise-group-best *string*

Identifies the BGP peer-group add-path advertise group best in the format device-ip,peer-group-name:add-path-advertise-group-best(true/false).

--pg-ipv6-uc-nbr-add-path-advertise-best *string*

BGP peer-group add-path advertise best in the format device-ip,peer-group-name:add-path-advertise-best.

--pg-ipv6-uc-nbr-activate *stringArray*

BGP peer-group Activate in the format device-ip,peer-group-name:activate(true/false).

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 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-ipv4-uc-nbr-add-path-capability 10.20.246.15,pg2:send
--pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.15,pg2:true
--pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.15,pg2:true
--pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.15,pg2:8
.....
--pg-name 10.20.246.16:pg2 --pg-asn 10.20.246.16,pg2:55002
--pg-bfd-enable 10.20.246.16,pg2:true
--pg-ipv4-uc-nbr-add-path-capability 10.20.246.16,pg2:receive
--pg-ipv4-uc-nbr-add-path-advertise-all 10.20.246.16,pg2:false
--pg-ipv4-uc-nbr-add-path-advertise-group-best 10.20.246.16,pg2:false
--pg-ipv4-uc-nbr-add-path-advertise-best 10.20.246.16,pg2:4
```

This example shows the use of the --pg-update-source-ip flag with IPv6 addressing.

```
efa tenant service bgp peer-group create --tenant "tenant1" --name "bgpPeerGroup2"
--pg-name 10.20.246.15:bgppg2 --pg-asn 10.20.246.15,bgppg2:100 --pg-
update-source-ip 10.20.246.15,bgppg2:10::10

efa tenant service bgp peer-group update --tenant "tenant1" --name "bgpPeerGroup2" --
operation peer-group-add
--pg-name 10.20.246.15:bgppg1 --pg-asn 10.20.246.15,bgppg1:200 --pg-
update-source-ip 10.20.246.15,bgppg1:10::20
```

```

efa tenant service bgp peer-group show --tenant tenant1 -detail
=====
Name           : bgpPeerGroup2
Tenant         : tenant1
State          : bgp-pg-created
Description    :

Peer Group
-----
Device IP      : 10.20.246.15
Peer Group     : bgppg2
Remote ASN     : 100
Next Hop Self  : false
Update Source IP : 10::10
BFD Enabled    : false
BFD Interval   :
BFD Rx        :
BFD Multiplier :
MD5 Password   :
Remove Private AS : false
Activate      : Activate (afi)
                true (ipv4)
                false (ipv6)

Prefix List In :
Prefix List Out :
Route Map In   :
Route Map Out  :
Send Community :
Add Path Capability :
Add Path Advertise :
Dev State      : provisioned
App State      : cfg-in-sync

Device IP      : 10.20.246.15
Peer Group     : bgppg1
Remote ASN     : 200
Next Hop Self  : false
Update Source IP : 10::20
BFD Enabled    : false
BFD Interval   :
BFD Rx        :
BFD Multiplier :
MD5 Password   :
Remove Private AS : false
Activate      : Activate (afi)
                true (ipv4)
                false (ipv6)

Prefix List In :
Prefix List Out :
Route Map In   :
Route Map Out  :
Send Community :
Add Path Capability :
Add Path Advertise :
Dev State      : provisioned
App State      : cfg-in-sync
=====

```

efa tenant service mirror session create

Creates a mirror session for a specified tenant.

```
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 ]
```

-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**.

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.

```
efa tenant service mirror session delete [ --tenant tenant-name | --name
mirror-session-name | --force ]
```

--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.

```
efa tenant service mirror session configure [ --tenant tenant-name | --name mirror-session-name ]
```

```
--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.

```
efa tenant service mirror session show [ --tenant tenant-name | --name mirror-session-name ]
```

```
--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.

```
efa tenant show [ --name tenant-name | --detail ]
```

```
--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.

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.

```
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 ]
--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
    List of physical ports of devices which will be reserved as mirror destination ports.
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].
--force
    Forces the deletion on the tenant if the option is provided.

```

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.

```
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 | 128 } | --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 | --layer3-extension-enable { true | false } | --next-hop-recursion-enable { true | false } | --default-information-originate-enable { true | false } ]
```

--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 128.

--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** | **128** }

Specifies the maximum number of resilient hashing paths allowed per tenant VRF. Valid values are 8, 16, 64, or 128. The default is 0 and unless specified, 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

--layer3-extension-enable { **true** | **false** }

Enable layer3-extension. Example: --layer3-extension-enable true | --layer3-extension-enable false

--next-hop-recursion-enable { **true** | **false** }

Enable next hop recursion for vrf. Example: --next-hop-recursion-enable true | --next-hop-recursion--enable false

--default-information-originate-enable { **true** | **false** }

Enables or disables default-information-originate per tenant VRF. Default is disabled.

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.

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 create --name ten1vrf2 --tenant ten1 --max-path 128 --rh-max-path 128
```

```
efa tenant vrf show --detail
```

```

=====
Name                : ten1vrf1
Tenant              : ten1
Routing Type        : distributed
Centralized Routers :
Enable Layer3 Extension : true
Redistribute        : connected
Max Path            : 8
Local Asn           :
L3VNI               :
EVPN IRB BD         :
EVPN IRB VE         :
BR VNI              :
BR BD               :
BR VE               :
RH Max Path         :
Enable RH ECMP      : false
Enable Graceful Restart : false
Enable NextHop Recursion: false
Route Target        :
Static Route        :
Static Route BFD    :
Network Route Address :
Static Network      :
Aggregate Address   :
VRF Type            : private
State               : vrf-created
Dev State           : not-provisioned
App State           : cfg-ready
=====

```

```

=====
Name                : ten1vrf2
Tenant              : ten1
Routing Type        : distributed
Centralized Routers :
Enable Layer3 Extension : true
Redistribute        : connected
Max Path            : 128
Local Asn           :
L3VNI               :
EVPN IRB BD         :
EVPN IRB VE         :
BR VNI              :
BR BD               :
BR VE               :
RH Max Path         : 128
Enable RH ECMP      : false
Enable Graceful Restart : false
Enable NextHop Recursion: false
=====

```



```

Route Target      :
Static Route      :
Static Route BFD  :
Network Route Address :
Static Network    :
Aggregate Address :
VRF Type          : private
State             : vrf-created
Dev State         : not-provisioned
App State         : cfg-ready
=====

```

```
efa tenant vrf create --name vs --tenant t1 -- default-information-originate-enable true
```

```
efa tenant vrf show --tenant t1 --name vs --detail
```

```

Name              : vs
Tenant            : t1
Routing Type      : distributed
Centralized Routers :
Enable Layer3 Extension : true
Redistribute      : connected
Max Path          : 8
Local Asn         :
L3VNI             :
EVPN IRB BD       :
EVPN IRB VE       :
BR VNI           :
BR BD            :
BR VE            :
RH Max Path       :
Enable RH ECMP    : false
Enable Graceful Restart : false
Enable NextHop Recursion : false
Default Information Originate : true
Route Target      :
Static Route      :
Static Route BFD  :
Network Route Address :
Static Network    :
Aggregate Address :
VRF Type          : private
State             : vrf-created
Dev State         : not-provisioned
App State         : cfg-ready

```

efa tenant vrf delete

Deletes a VRF.

```
efa tenant vrf delete [--name vrf name | --tenant tenant name ]
```

--name *vrf name*

Specifies the name of the VRF.

--tenant *tenant name*

Specifies the name of the tenant.

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.

```
efa tenant vrf error show [ --name vrf-name | --tenant tenant-name ]
```

--name *vrf-name*

Specifies the name of the VRF.

--tenant *tenant-name*

Specifies the name of the tenant.

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.

```
efa tenant vrf show [ --name vrf-name | --tenant tenant-name | --detail ]
```

--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.

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 t1 --name vs --detail
```

```

Name : vs

Tenant : t1

Routing Type : distributed

Centralized Routers :

Enable Layer3 Extension : true

Redistribute : connected

Max Path : 8

Local Asn :

L3VNI :

EVPN IRB BD :

EVPN IRB VE :

BR VNI :

BR BD :

BR VE :

RH Max Path :

Enable RH ECMP : false

Enable Graceful Restart : false

Enable NextHop Recursion : true

Route Target :

Static Route :

Static Route BFD :

Network Route Address :

```

```

Static Network :

Aggregate Address :

VRF Type : private

State : vrf-created

Dev State : not-provisioned

App State : cfg-ready

```

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
Enable NextHop Recursion: 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.

```

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 | next-hop-recursion-update | network-add | network-delete
  | static-network-add | static-network-delete | aggregate-address-add
  | aggregate-address-delete | default-information-originate-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 | 128 } | --rh-ecmp-enable { true
  | false } | --centralized-router | --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 | --next-hop-recursion-
  update { true | false } | --default-information-originate-enable
  { true | false } ]

--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 | next-hop-recursion-update |
network-add | network-delete | static-network-add | static-network-
delete | aggregate-address-add | aggregate-address-delete | default-
information-originate-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
- next-hop-recursion-update
- network-add
- network-delete
- static-network-add
- static-network-delete
- aggregate-address-add
- aggregate-address-delete
- default-information-originate-update



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 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 128.

--redistribute { **static** | **connected** }

Specifies the redistribute type for routes. Valid values are static or connected.

--rh-max-path { **8** | **16** | **64** | **128** }

Specifies the maximum number of resilient hashing paths allowed per tenant VRF. Valid values are 8, 16, 64, or 128.

--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. Format: 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.
- next-hop-recursion-update** { **true** | **false** }
Enable next hop recursion for vrf. Example: --next-hop-recursion-enable true | --next-hop-recursion--enable false
- default-information-originate-enable** { **true** | **false** }
Enables or disables default-information-originate per tenant VRF. Default is disabled.

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.

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 tenant vrf update --name tenlvrf1 --tenant ten1 --operation max-path-add --max-path
128
efa tenant vrf update --name tenlvrf1 --tenant ten1 --operation rh-max-path-add --rh-max-
path 128
```

```
efa tenant vrf show --detail
```

```
=====
Name                : tenlvrf1
Tenant              : ten1
Routing Type       : distributed
Centralized Routers :
Enable Layer3 Extension : true
Redistribute       : connected
Max Path           : 128
Local Asn          :
L3VNI              : 8192
EVPN IRB BD        : 4096
EVPN IRB VE        : 8192
BR VNI             : 4096
BR BD              :
BR VE              :
RH Max Path        : 128
Enable RH ECMP     : false
Enable Graceful Restart : false
Enable NextHop Recursion: false
Route Target       : import 101:101
                   : export 101:101
Static Route       :
Static Route BFD   :
Network Route Address :
Static Network     :
Aggregate Address   :
VRF Type           : private
```

```

State                : vrf-device-created
Dev State            : provisioned
App State            : cfg-in-sync
=====
=
Name                 : ten1vrf2
Tenant               : ten1
Routing Type        : distributed
Centralized Routers :
Enable Layer3 Extension : true
Redistribute         : connected
Max Path             : 128
Local Asn            :
L3VNI                : 8191
EVPN IRB BD         : 4095
EVPN IRB VE         : 8191
BR VNI              : 4096
BR BD                :
BR VE                :
RH Max Path         : 128
Enable RH ECMP      : false
Enable Graceful Restart : false
Enable NextHop Recursion: false
Route Target        : import 102:102
                   : export 102:102
Static Route        :
Static Route BFD    :
Network Route Address :
Static Network      :
Aggregate Address   :
VRF Type            : private
State               : vrf-device-created
Dev State           : provisioned
App State           : cfg-in-sync
=====

```

efa vcenter debug

A set of debug commands for troubleshooting XCO vCenter issues. This command also sets configuration parameters for vCenter tenants.

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 XCO 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 XCO 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 XCO vCenter service. Debug levels can be set to receiving information level messages or complete debugging messages.

This command has no defaults.

--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 XCO 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 XCO vCenter service database.

--level *[info | debug]*

Sets the debug level for this XCO 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.

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 | 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: | |
+-----+-----+-----+-----+-----+

```

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 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

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 XCO 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 XCO instance.

```
efa vcenter delete --host string [ --cleanup-tenants ]
```

```
--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.

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.

```
efa vcenter links [ physical | unconnected ] [ --host string | --esxi string ]
```

```
efa vcenter links virtual --esxi string
```

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.

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.

```
efa vcenter list [ --host string ]
```

--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.

This example lists a VMware vCenter Server that has just been added as a tenant with this XCO 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 | 0 |
+-----+-----+-----+-----+
+-----+-----+-----+-----+

```

This example lists a VMware vCenter Server with its status updated to this XCO 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           |
+-----+-----+-----+-----+-----+
+-----+-----+

```

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 XCO instance.

```
efa vcenter register [ --host string | --username string | --password string | --tenant string ]
```

This command has no defaults. All flags are mandatory.

Tenant details must be registered before running the command.

--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.

This example registers a VMware vCenter Server as a tenant with XCO 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.

```
efa vcenter update [ --host string | --username string | --password string ]
```

This command has no defaults. All flags are mandatory.

--host *string*

IP address or host name of the VMware vCenter server from which the local XCO 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.

This example updates the local XCO 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 XCO that is installed.

efa version

The following example shows typical output for the command.

```
(efa:user)user@xco-64-160:~$ efa version
Version : 3.7.0
Build: 34
Time Stamp: 24-12-02:19:24:36
Mode: Secure
Deployment Type: single-node
Deployment Platform: SERVER
Deployment Suite: Fabric Automation
Deployment IP Mode: ipv4
Node IP: 10.64.64.160
--- Time Elapsed: 1.280732ms ---
(efa:user)user@xco-64-160:~$
```

efa-change-hostname

Changes the host name in the XCO system after XCO is deployed.

efa-change-hostname *old-host-name*

old-host-name

Specifies the host name that you are changing.

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 *ExtremeCloud Orchestrator 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, XCO is not operational while this command is running. In a multi-node deployment, XCO remains operational if the command is running on the standby node. XCO is not operational if the command is running on the active node.

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 XCO is deployed.

efa-change-ip

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/`.

XCO is not operational while this command is running.

After the IP address is changed, either run `source /etc/profile` or open a new XCO session to log in.

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 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 XCO is deployed.

efa-change-vip *new-vip*

new-vip

Specifies the new VIP for the multi-node deployment.

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/`.

XCO is not operational while this command is running.

After the VIP is changed, run `source /etc/profile` or open a new XCO session to log in.

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 XCO.

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

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
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
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
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
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 XCO 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-lk6r9  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/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/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-ffccd5bf6-755ld  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-nntl1b  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	TYPE	CLUSTER-IP	EXTERNAL-
IP	PORT(S)	AGE	SELECTOR
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/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	SELECTOR
deployment.apps/efa-api-docs	2/2	2	2
46m	efa-api-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/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/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 XCO 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 XCO 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
d (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.

```